

CS 309: Autonomous Robots

FRI I

Introduction to Artificial Intelligence

Instructor: Justin Hart

http://justinhart.net/teaching/2020_spring_cs309/

What is Intelligence?

A very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience. It is not merely book learning, a narrow academic skill, or test-taking smarts. Rather, it reflects a broader and deeper capability for comprehending our surroundings— "catching on," "making sense" of things, or "figuring out" what to do.

-“Mainstream Science on Intelligence”

Frustratingly, this definition uses words like “mental” and “think” to define “intelligence.”

Concepts of "intelligence" are attempts to clarify and organize this complex set of phenomena. Although considerable clarity has been achieved in some areas, no such conceptualization has yet answered all the important questions, and none commands universal assent. Indeed, when two dozen prominent theorists were recently asked to define intelligence, they gave two dozen, somewhat different, definitions.

-“Intelligence: Knowns and Unknowns”

“Goal-directed adaptive behavior.”

-Sternberg & Salter

“The ability to deal with cognitive complexity.”

-Linda Gottfredson

“Judgement, otherwise called 'good sense,' 'practical sense,' 'initiative,' the faculty of adapting one's self to circumstances ..

auto-critique.

-Alfred Binet

Artificial Intelligence
has historically been a moving target

- 1951 – Strachey & Prinz
 - Programs for **checkers & chess**

- 1965 – Dartmouth Conference

- Founding event of Artificial Intelligence
 - Lasted 8 weeks
 - Gathered top figures at the time
- They talked a lot about solving **checkers and chess**



- 1968 – 2001: A Spacy Odyssey

- Essentially as big as every movie that came out this year rolled into 1
- A talking AI drives a spaceship & runs all life support
- Version 2.0 talks about dreaming and being afraid of death in the sequel
- But the only reason we know it is smart is because it can **play chess**

- 1990 – Ray Kurzweil
 - Predicts that a computer will beat a world chess champion by 1998
 - His model becomes a foundational principle for his book “The Singularity is Near” and transhumanism.

- 1997 – IBM’s Deep Blue
 - Beats world champion Garry Kasparov

- 2016 – DeepMind’s AlphaGo
 - Beats Lee Sedol – A 9-dan Go player
 - In 4 of 5 games
 - The only human to have ever beaten AlphaGo

- AlphaZero learns only by playing against itself
 - Superhuman performance after 24 hours
- Top Chess and Go programs are rated by playing against each other



For 46 years
winning chess a major theme

So...

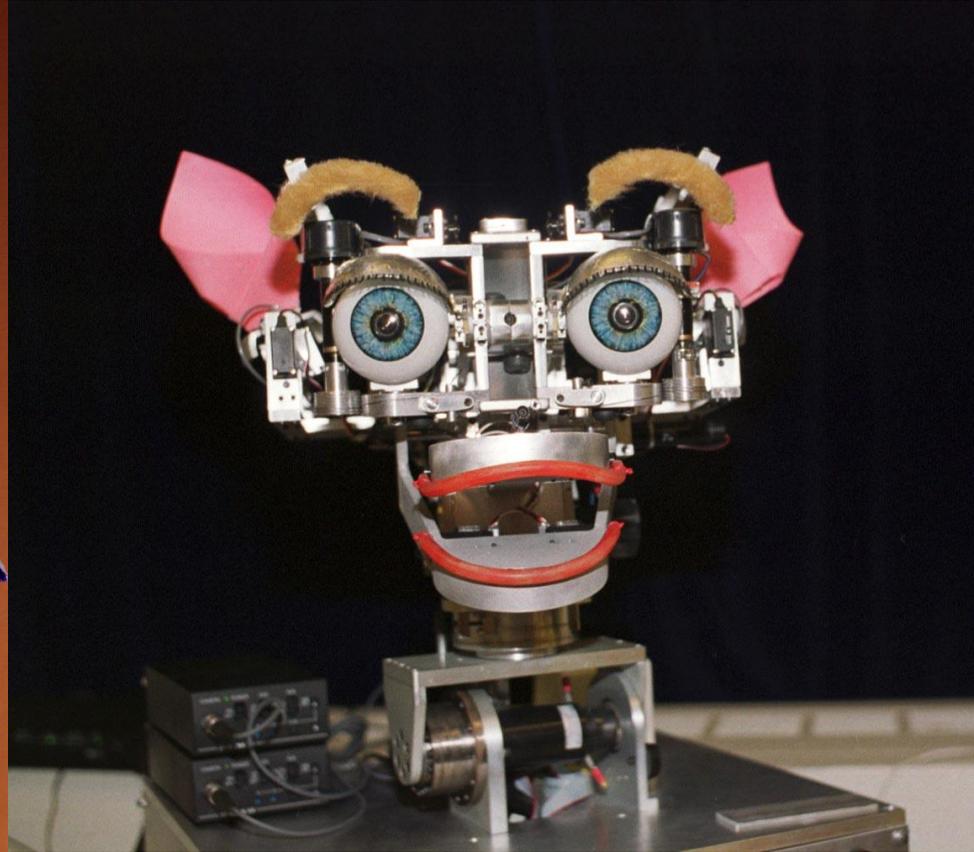
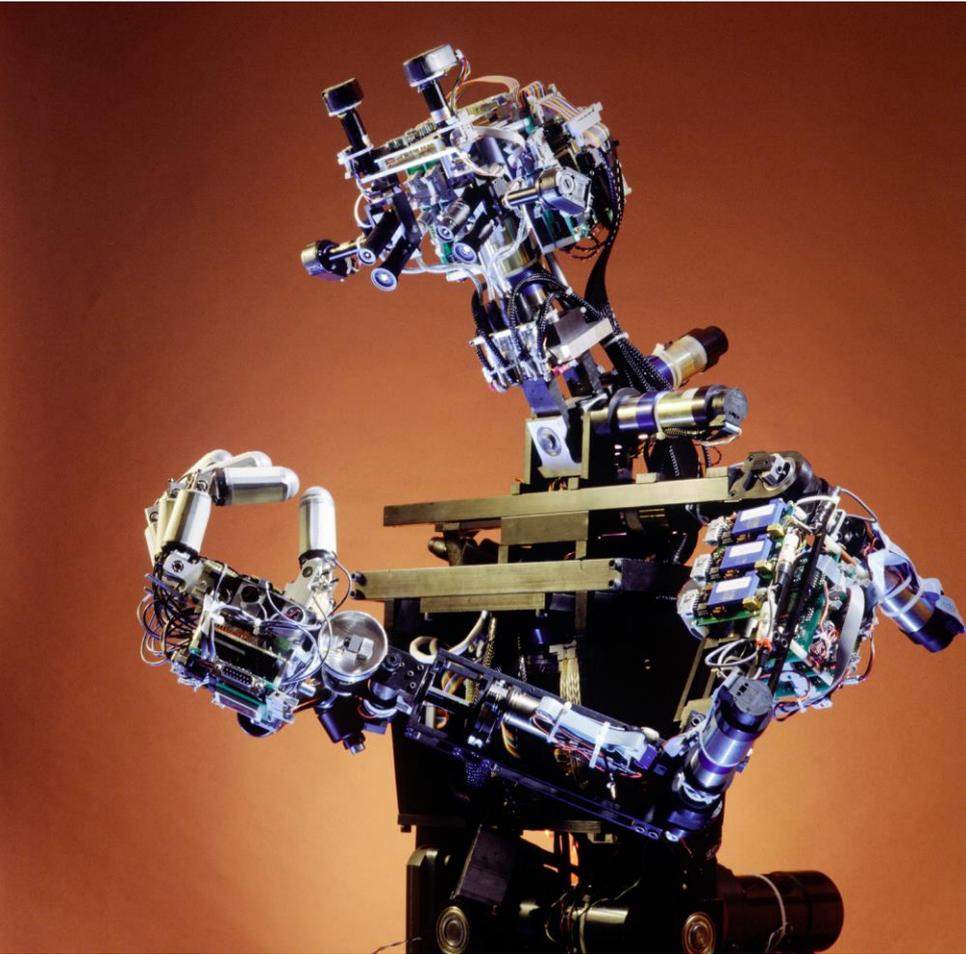
What's been up for the past 23 years?

RoboCup started in 1997

The goal is to beat the world champion team by 2050

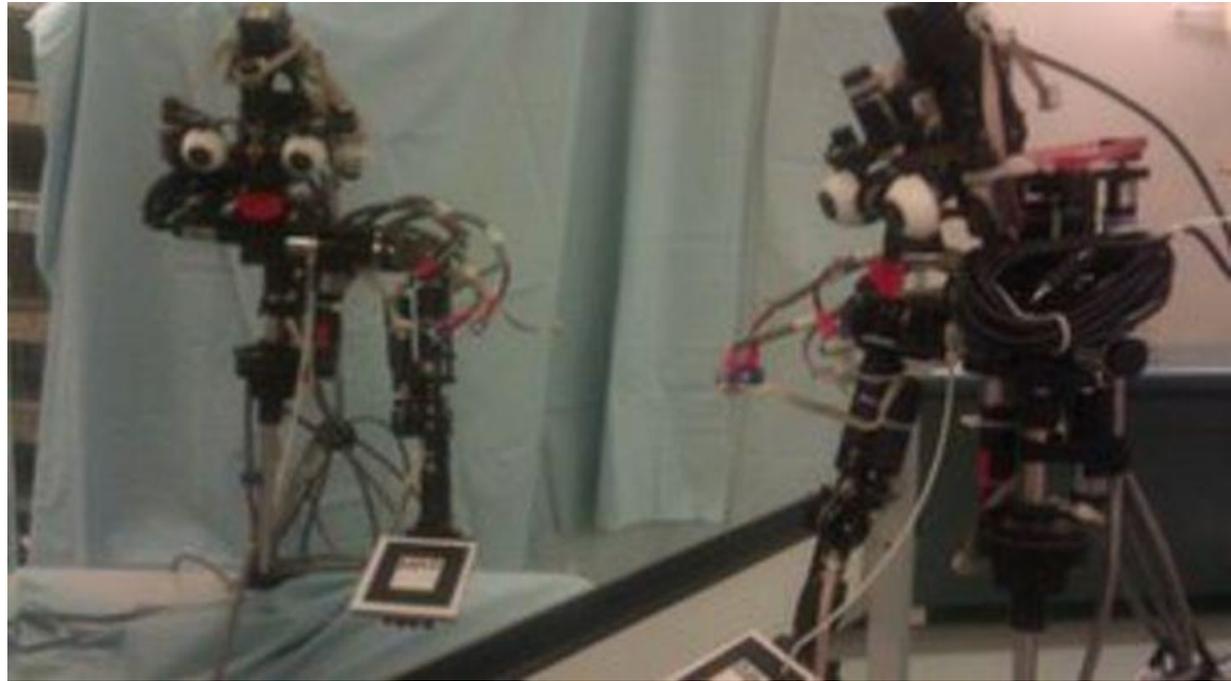


Developmental Robotics and Human-Robot Interaction



Developmental Robotics

- Emulate human cognitive processes
 - Better understanding of the cognitive process
 - Robot capabilities based on human capabilities



- Eventual goal
 - Build robots which learn like people do
 - Develop human-like AI

Human-Robot Interaction

Two general schools of thought

- HRI as understanding and implementing human behavior through the use of robots or on robots
- HRI as design



Androids



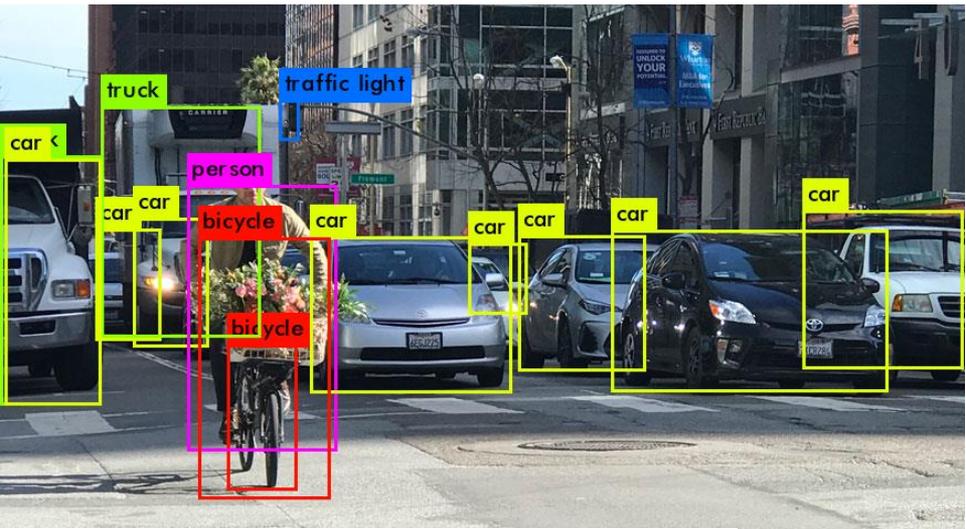
Siri, Alexa, Google Voice, Bixby...



Autonomous vehicles

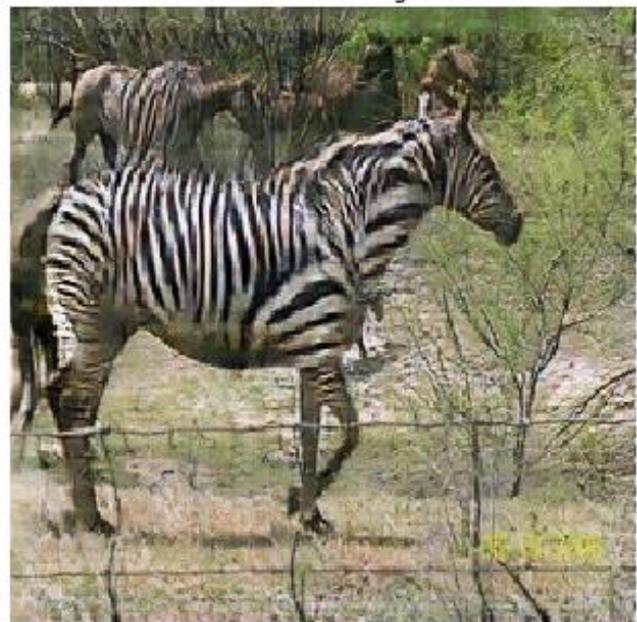


Deep Learning



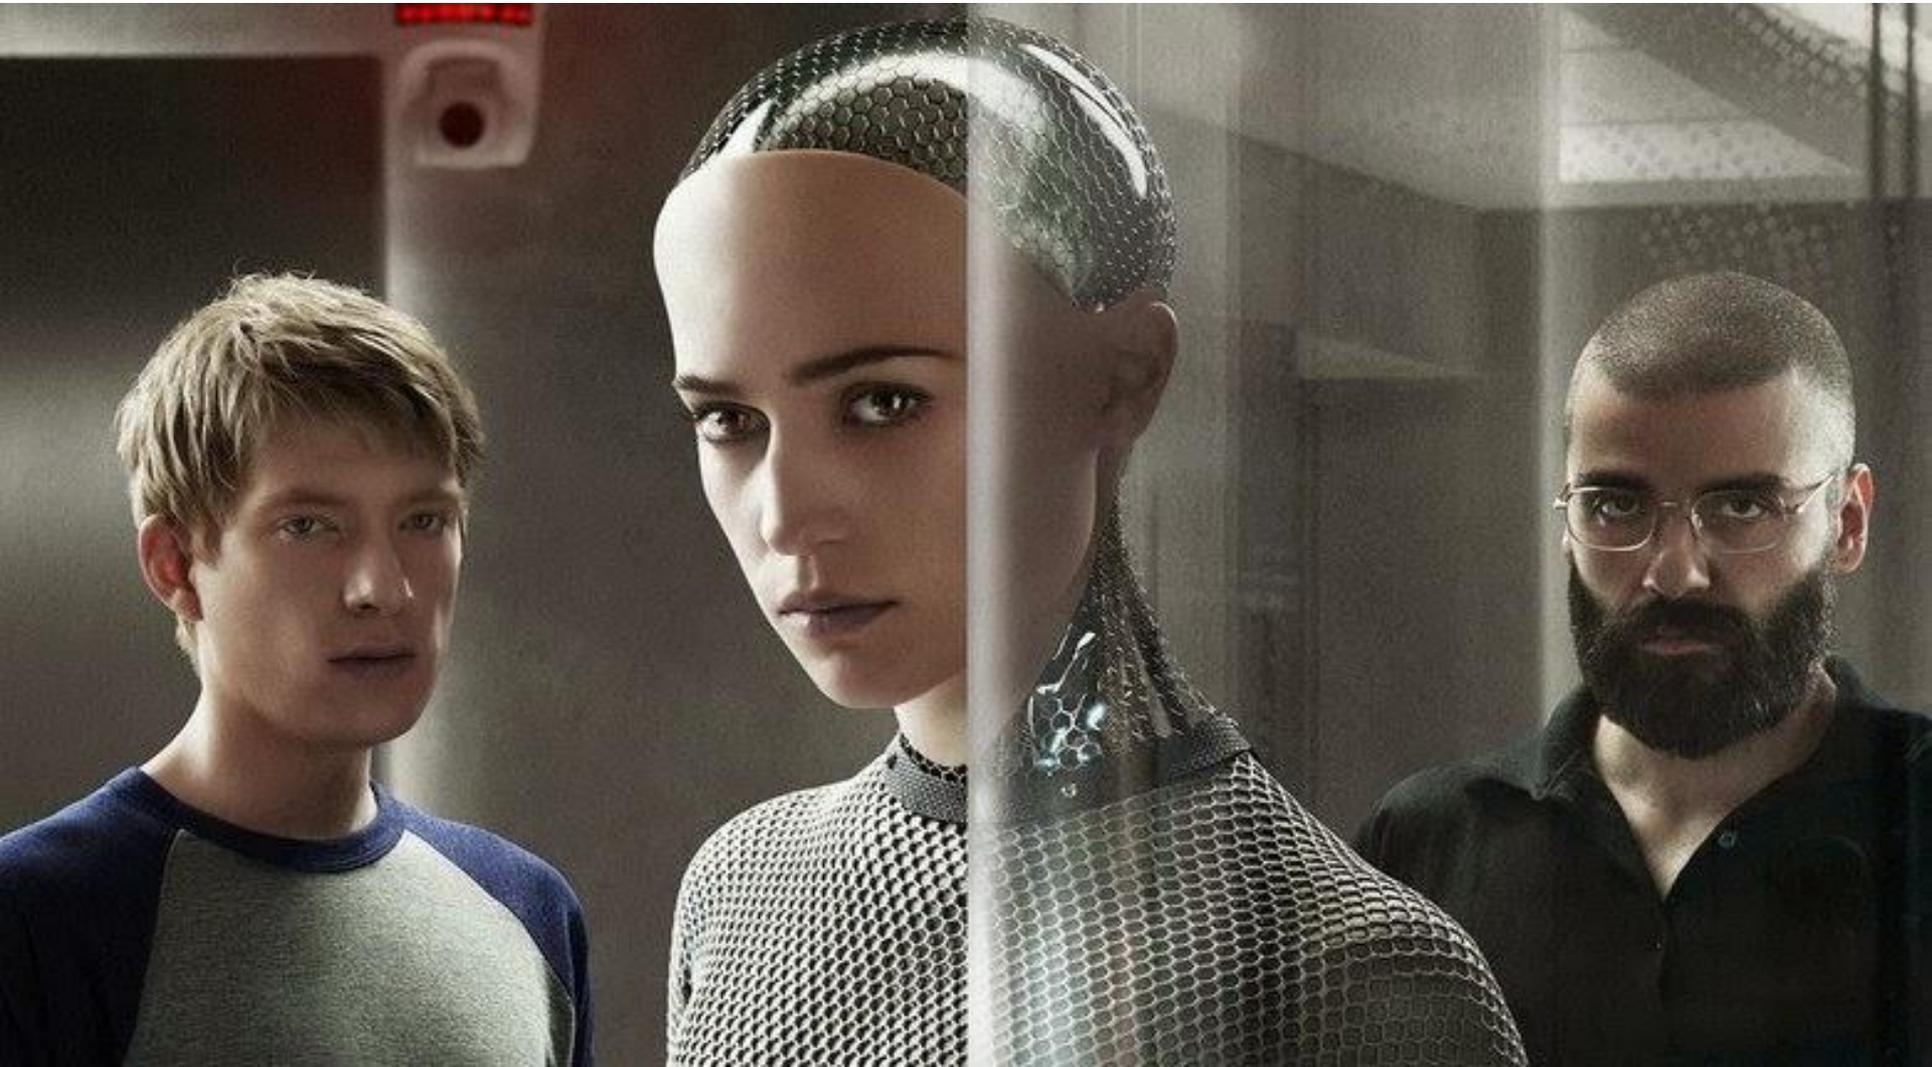
Input Image

Predicted Image



Hard sci-fi tries to track recent research trends..

Recent fictional robots demonstrate intelligence through human-robot interaction.





IJCAI



All of this used to be contrasted with mainstream AI, and yet now represents some of the top work.

- Pre-historic - Medieval times

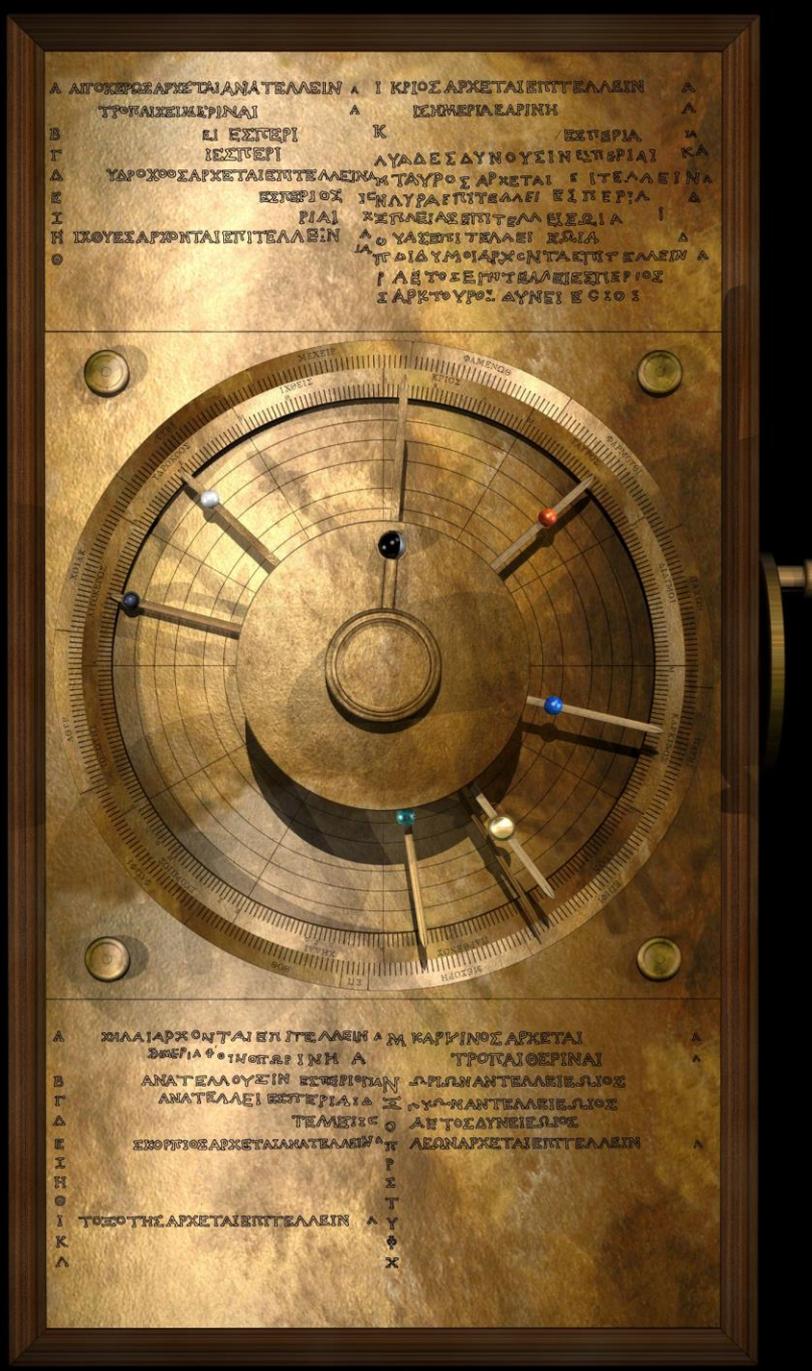
- “Tally sticks” help in basic computations

- Ancient Greece

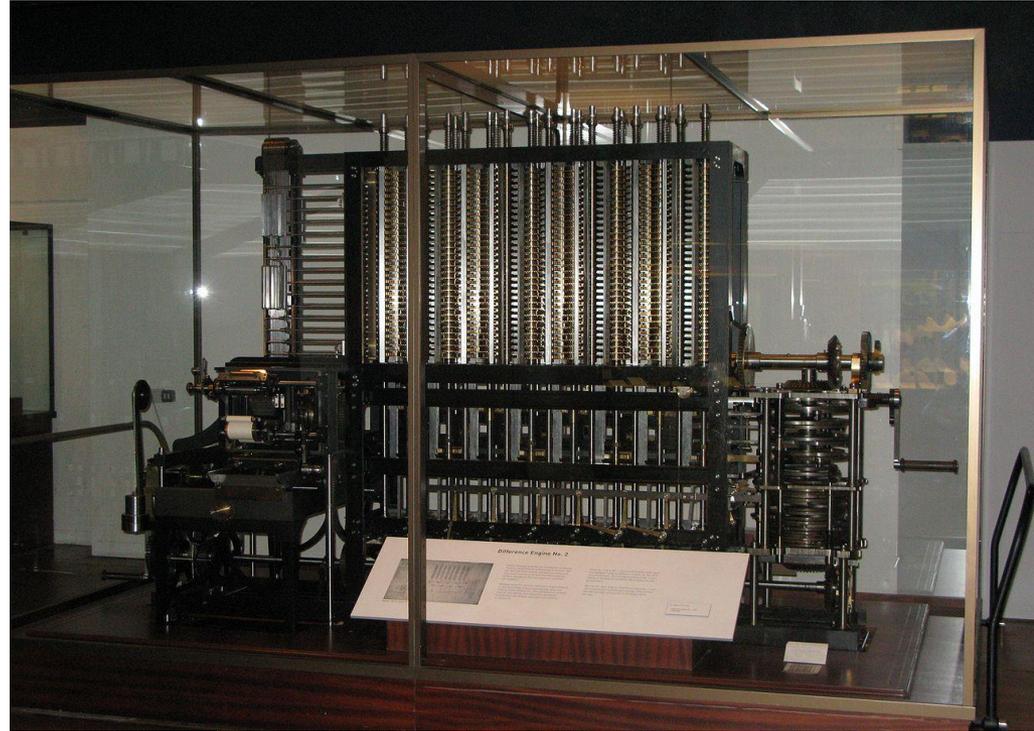
- Antikythera Mechanism
 - Predicts astronomical positions
 - Computes Olympiads

- First mathematical tables
 - Document numerical values of common, important functions
 - Can be found in older textbooks

- Early algorithms
 - Sieve of Erasthones
 - Computes prime numbers
 - Name comes from Al-Khwarizmi
 - 9th century Persian mathematician



- 1613 – First written reference to a “computer”
 - This is a person whose job it is to compute things.
 - People worked in this job until the 1970s



- 1642 – Mechanical Calculators
 - First one by Blaise Pascal

- 1822 – Difference Engine

- Charles Babbage develops the hardware to compute polynomials
- Ada Lovelace
 - Has the idea that this machine could be programmable
 - Develops the first algorithm that can run on a machine
- Two were built for museums in 2008 – These were the first to be built

- 1907 – Vacuum tubes
- 1909 – Crystal oscillators
- 1925 – Field-effect transistors

- 1936 – Universal Turing Machines
 - A practical, programmable computer
 - Based on physical components



- 1939 – WWII starts
- 1941 – Bombes in Bletchley Park
 - Faster method for breaking Enigma
 - Key in winning WWII
 - Not at all like in the movie
 - Improvement of an earlier Polish machine.
 - He didn't have to fight the other staff on this. Turing was hired to build this machine!



- 1956 – TX-O
 - First transistor computer
- 1960 – PDP-1
 - First “minicomputer”
- 1971 – Intel 4040
 - First microchip
- 1981 – IBM PC

- The rest should be familiar to you



Developmental Robotics

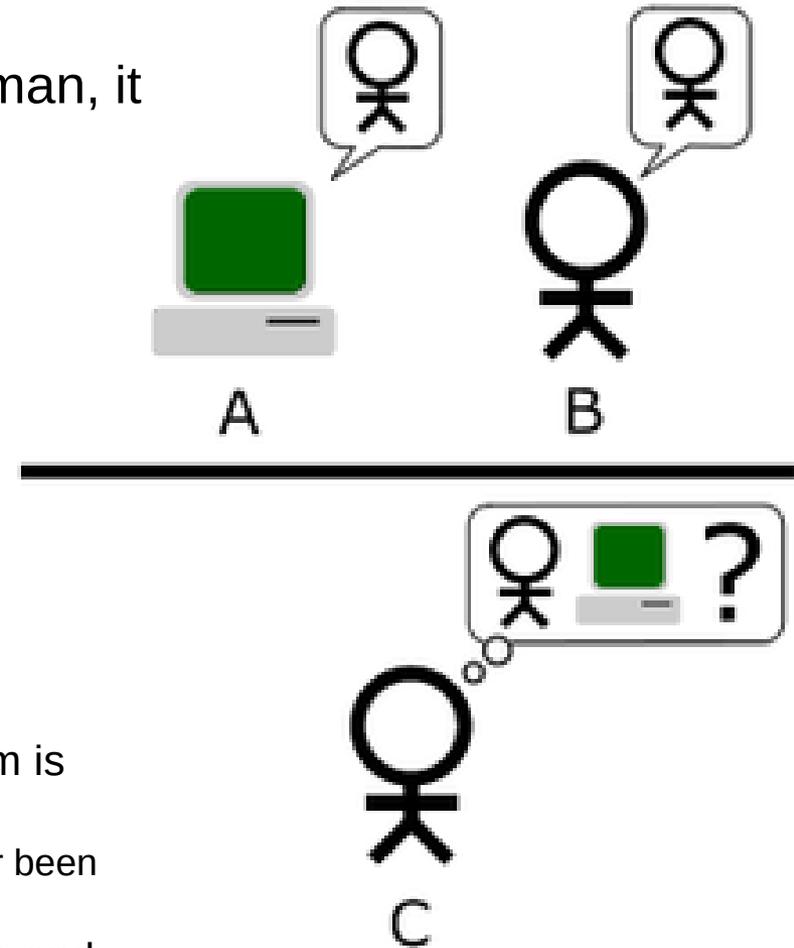
- 1950 – Alan Turing writes “Computing Machinery and Intelligence.”
 - What it means to “think” is a tough question
 - Let's “replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.”
- “Imitation Game”
 - Players
 - Player A – Man
 - Player B – Woman
 - Player C – Interrogator
 - Can the interrogator determine the sex of the players by asking questions?
 - Both players claim to be a woman
 - What happens if a machine replaces Player A?
 - If the interrogator cannot identify the machine, the machine passes

The Turing Test

- The modern test
 - Jury of people & computers
 - If participants believe the machine is human, it passes
- The Loebner Prize
 - Ranks chatbots as most convincing
 - Generally scorned by AI experts
 - Very old chat programs can now do fairly well
 - Regarded as a publicity stunt
 - Cash prize
 - \$3000 – Best program
 - \$25,000 – Convinces the jury that the program is human
 - This will only be awarded once, and has never been awarded
 - \$100,000 – Adds understanding text, auditory, and visual input
 - Once this happens, the contest ends
 - Might have the CEO of the company that made of Mitsuku come visit in the spring
 - 5x winner of Loebner prize!

The Turing Test

- The modern test
 - Jury of people & computers
 - If participants believe the machine is human, it passes
- The Loebner Prize
 - Ranks chatbots as most convincing
 - Generally scorned by AI experts
 - Very old chat programs can now do fairly well
 - Regarded as a publicity stunt
 - Cash prize
 - \$3000 – Best program
 - \$25,000 – Convinces the jury that the program is human
 - This will only be awarded once, and has never been awarded
 - \$100,000 – Adds understanding text, auditory, and visual input
 - Once this happens, the contest ends



AI gets its start

- 1956 – The Dartmouth Conference
 - 6 weeks at Dartmouth College
 - Clarify and develop the ideas of researchers working on intelligent machines
 - Considered to be the meeting that started the field of AI
- Around this time, both AI and computer research experienced rapid growth and achievement

Early Optimism

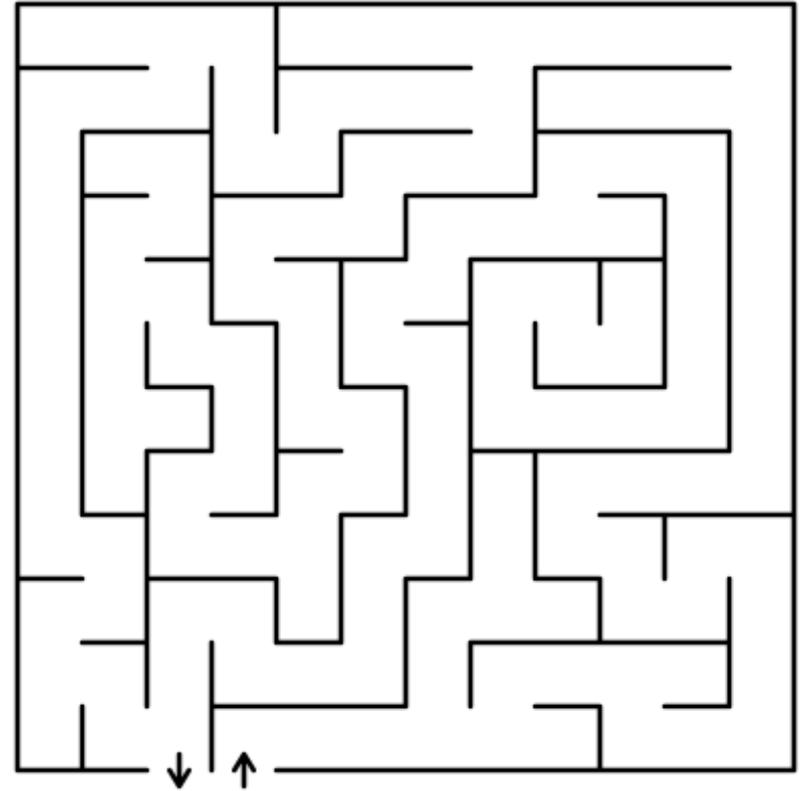
- AI and computing advanced much more quickly than expected at the start of the 20th century
 - Leading to unreasonable levels of optimism
- 1958 – Newell & Simon
 - “Within ten years, a digital computer will be the world’s chess champion”
 - 1997 – Deep Blue vs Garry Kasparov
 - “Within ten years, a digital computer will prove a new mathematical theorem”
 - So far, only computer-assisted proofs have been generated
- 1965 – Simon
 - “Machines will be capable, within 20 years, of doing any work a man can do.”
- Today – Kurzweil
 - 2019 – A computer has as much computer power as the human brain
 - 2045 – The Singularity – The first ultra-intelligent machine

AI spins into multiple areas of research

- Planning & Scheduling
 - Problem Solving
 - Knowledge Representation & Reasoning
- Machine Learning
 - Classification
 - Regression
 - Clustering
- Natural Language Processing
- Computer Vision
 - Perception
- Robotics
- ..and of course others

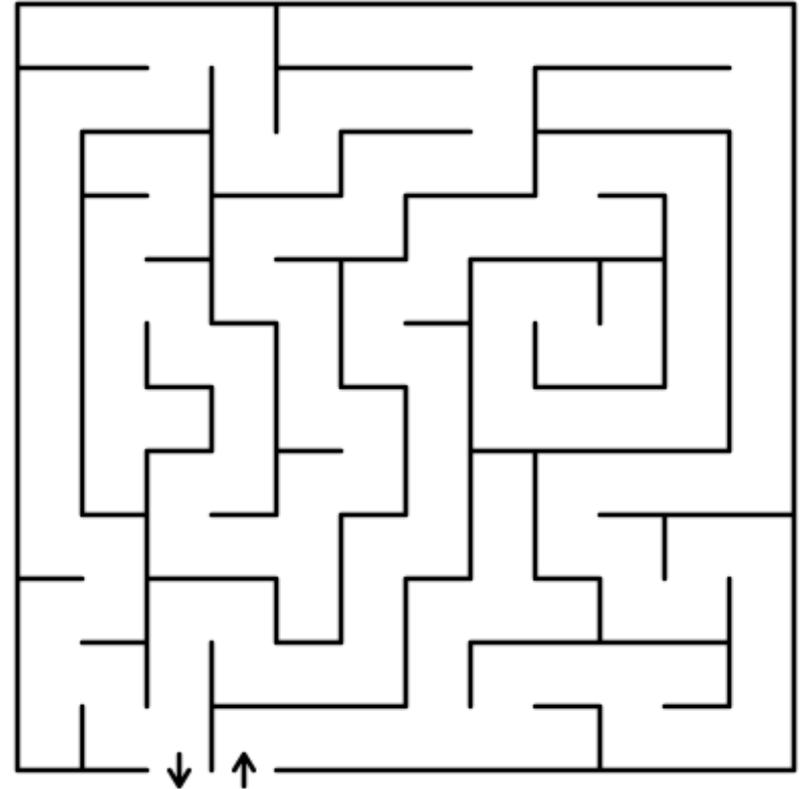
Planning & Scheduling

- Example: Solving a maze



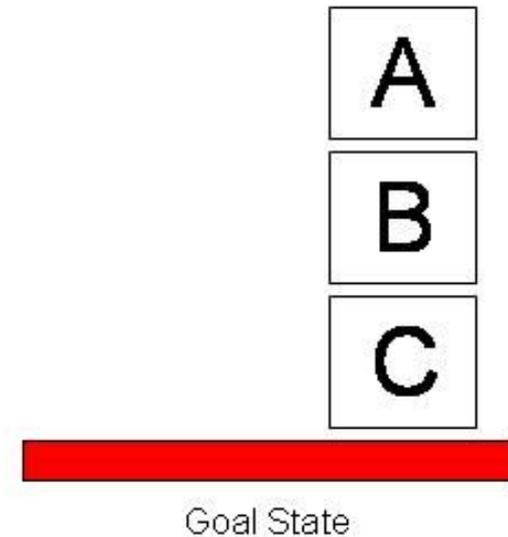
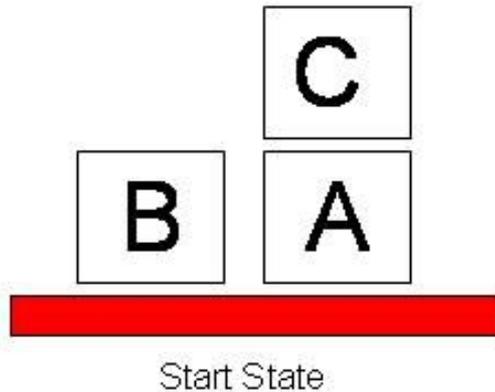
Planning & Scheduling

- Picture a robot in this maze
 - It runs a “search” algorithm
 - Up: Doesn't work
 - Left: Doesn't work
 - Right: Works!
 - Down: Doesn't work
- “Search”
 - Repeat until solved.
 - Store each position reached.
 - Try each move from each position.



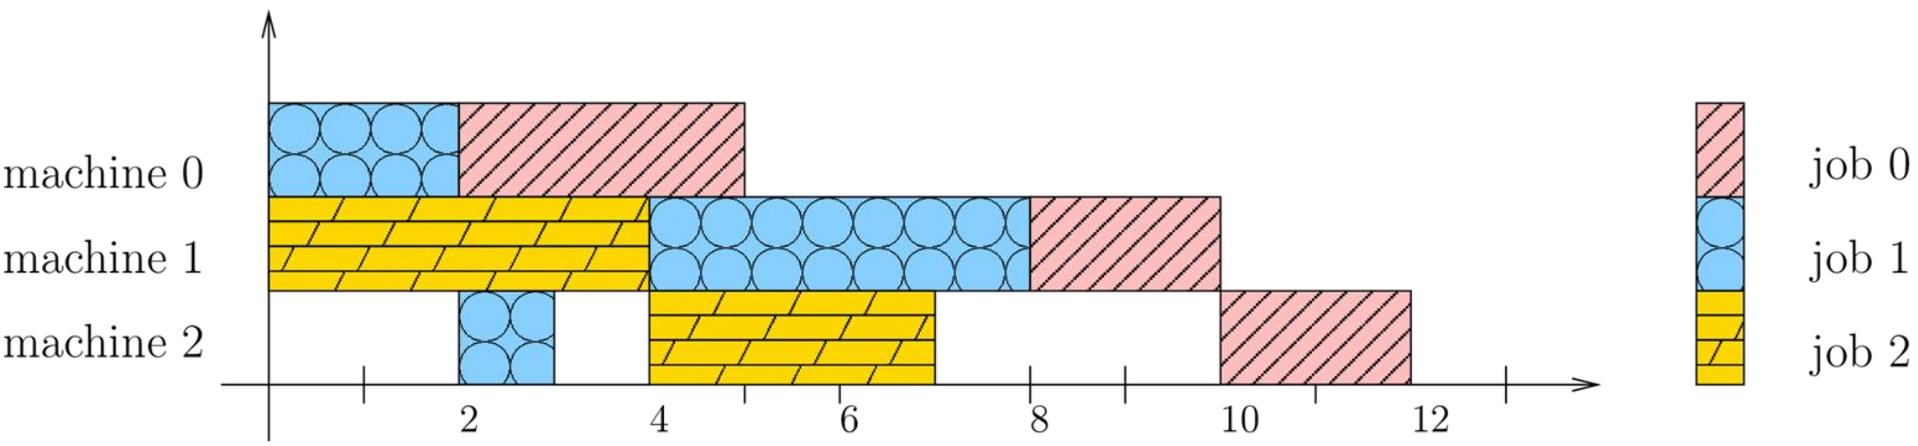
Planning & Scheduling

- Blocks World
 - Different problem
 - Same algorithm
- Agent can
 - Pick up blocks
 - Put down blocks
 - Stack blocks on each other
 - Stack blocks on the table
- Try every action you can do
- Remember the “state” after each action
- Try every action in each generated state



Planning & Scheduling

- Scheduling Problems
 - Classic: Job Shop Scheduling
 - N jobs
 - M machines
 - Find the fastest schedule to complete the job
 - Related (and current) problem
 - How do you sell Superbowl ads in order to maximize your profit?
 - Certain advertisers will offer more money
 - Certain slots are worth more money

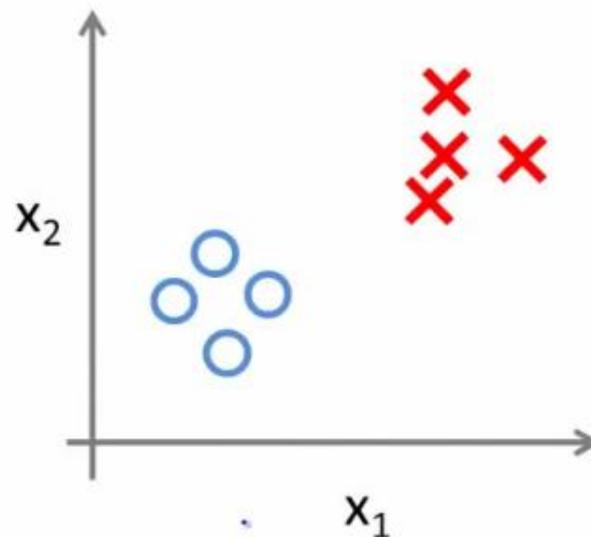


Machine Learning

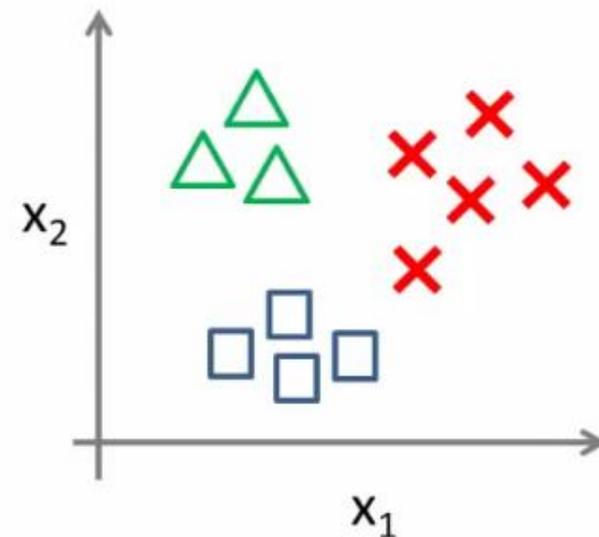
- Classification
 - Identifying a class that a datum fits into
- Binary classification
 - Two classes
 - Often “it is or it isn’t something”
 - Medical diagnoses

- Multi-class
 - Image classification
 - Dog
 - Cat
 - Soda can

Binary classification:

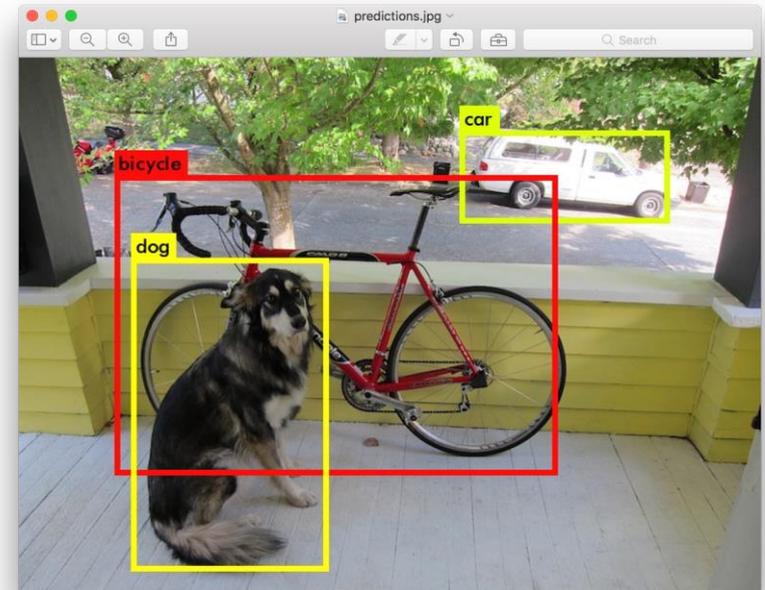


Multi-class classification:



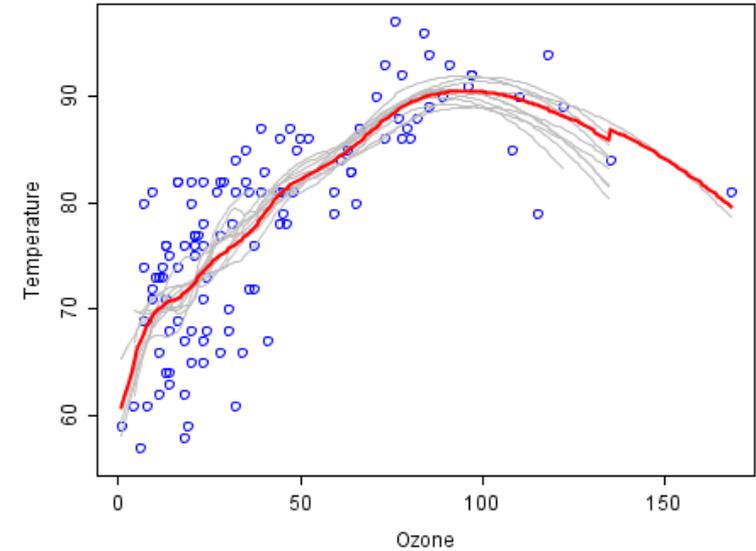
Machine Learning

- YOLO
 - Currently a common object recognition system
 - You Only Look Once

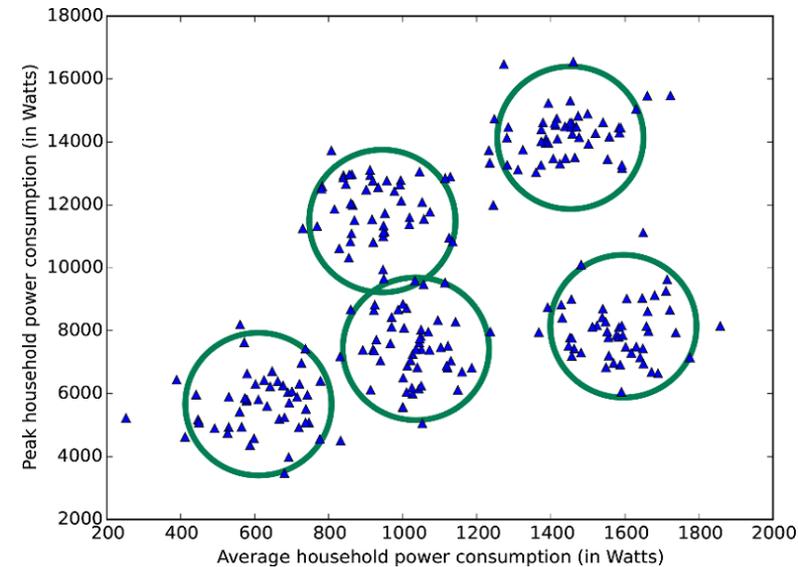


Machine Learning

- Regression
 - Given these values
 - What is the numerical value of <blank>
 - I want my car to go at this speed
 - I want to predict a stock's value



- Clustering
 - These data are similar in some way



Natural Language Processing

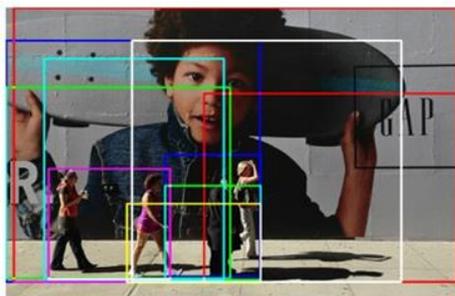
- Parsing
 - Syntactic
 - The dog is in the yard.
 - The/DT dog/NN is/VBZ in/IN the/DT yard/NN
 - Semantic
 - in(yard,dog)
- Perceptual Grounding
 - Pairing percepts to semantics
 - For instance, teaching a robot what a can looks like, or the color red, or the word “heavy”

Natural Language Processing

- Sentiment analysis
 - Is this a positive or a negative product review?

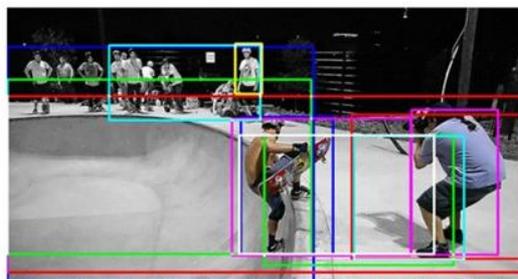
- Text summarization
 - Take a newspaper article, condense into 1-10 lines

- Image captioning
 - Look at a picture. What is in the picture?



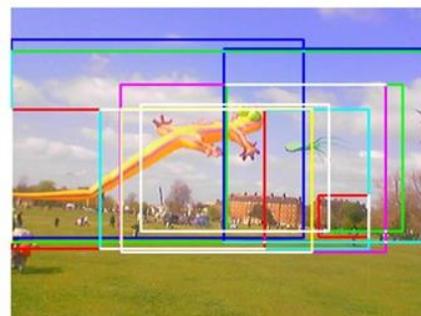
[men (0.59)] [group (0.66)] [woman (0.64)]
 [people (0.89)] [holding (0.60)] [playing (0.61)] [tennis (0.69)]
 [court (0.51)] [standing (0.59)] [skis (0.58)] [street (0.52)]
 [man (0.77)] [skateboard (0.67)]

a group of people standing next to each other
 people stand outside a large ad for gap featuring a young boy



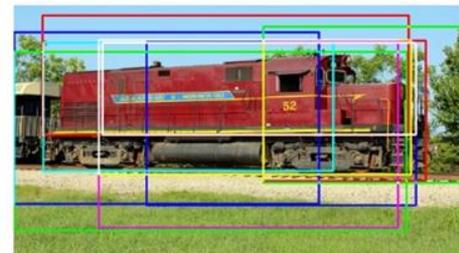
[person (0.55)] [street (0.53)] [holding (0.55)] [group (0.63)] [slope (0.51)]
 [standing (0.62)] [snow (0.91)] [skis (0.74)] [player (0.54)]
 [people (0.85)] [men (0.57)] [skiing (0.51)]
 [skateboard (0.89)] [riding (0.75)] [tennis (0.74)] [trick (0.53)] [skate (0.52)]
 [woman (0.52)] [man (0.86)] [down (0.61)]

a group of people riding skis down a snow covered slope
 a guy on a skateboard on the side of a ramp



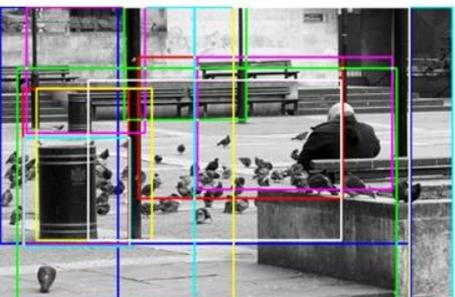
[airplane (0.57)] [plane (0.58)] [kites (0.93)] [people (0.80)]
 [flying (0.93)] [man (0.57)] [beach (0.64)] [wave (0.61)]
 [sky (0.61)] [kite (0.74)] [field (0.75)]

a couple of people flying kites in a field
 people in a field flying different styles of kites



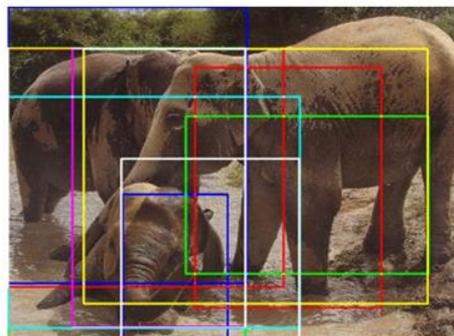
[parked (0.72)] [bench (0.63)] [truck (0.70)] [red (0.88)]
 [train (1.00)] [sitting (0.73)] [cars (0.58)] [traveling (0.52)]
 [grass (0.65)] [track (0.69)] [car (0.59)] [yellow (0.57)]
 [field (0.80)] [engine (0.56)] [down (0.54)] [tracks (0.94)]

a train traveling down train tracks near a field
 a red train is coming down the tracks



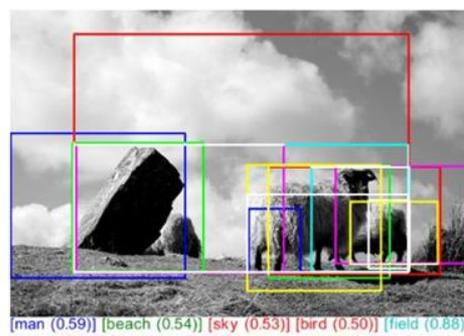
[umbrella (0.59)] [woman (0.52)]
 [fire (0.96)] [hydrant (0.96)] [street (0.79)] [old (0.50)]
 [bench (0.81)] [building (0.75)] [standing (0.57)] [baseball (0.55)]
 [white (0.82)] [sitting (0.65)] [people (0.79)] [photo (0.53)]
 [black (0.84)] [kitchen (0.54)] [man (0.72)] [water (0.56)]

a black and white photo of a fire hydrant
 a courtyard full of poles pigeons and garbage cans also has benches on
 either side of it one of which shows the back of a large person facin
 g in the direction of the pigeons



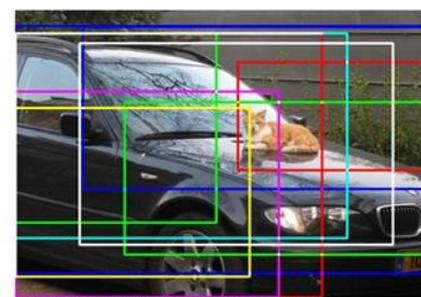
[horse (0.53)] [bear (0.71)] [elephant (0.99)] [elephants (0.95)]
 [brown (0.68)] [baby (0.82)] [walking (0.57)] [playing (0.61)]
 [man (0.57)] [standing (0.79)] [field (0.65)]
 [water (0.83)] [large (0.71)] [dirt (0.65)] [river (0.58)]

a baby elephant standing next to each other on a field
 elephants are playing together in a shallow watering hole



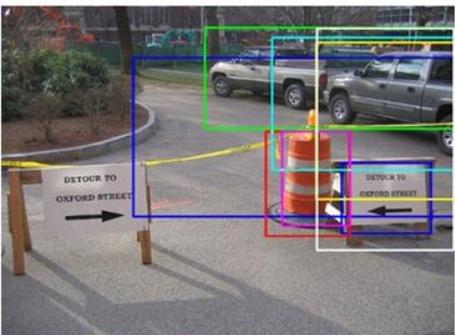
[man (0.59)] [beach (0.54)] [sky (0.53)] [bird (0.50)] [field (0.88)]
 [snow (0.86)] [mountain (0.59)] [standing (0.81)] [white (0.64)]
 [people (0.51)] [dog (0.60)] [cows (0.55)]
 [sheep (0.97)] [black (0.84)] [grass (0.64)] [horse (0.60)]
 [elephants (0.57)] [bear (0.81)]

a black bear standing on top of a grass covered field
 a couple of sheep standing up on a small hill



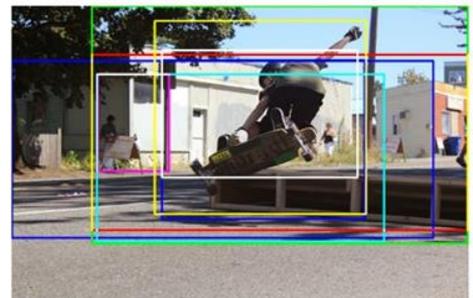
[bus (0.56)] [car (0.79)] [black (0.57)] [truck (0.86)]
 [street (0.57)] [bed (0.51)] [parked (0.55)] [dog (0.65)]
 [sitting (0.55)] [man (0.53)] [cat (0.72)]

a dog sitting on top of a car
 a cat is lying on the hood of a black car



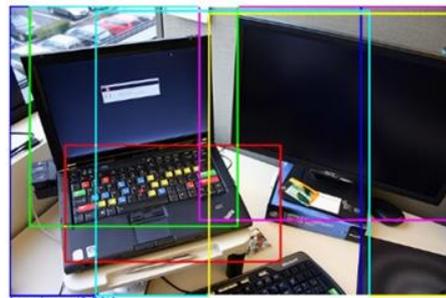
[street (0.89)] [truck (0.76)] [road (0.58)]
 [fire (0.95)] [hydrant (0.91)] [sitting (0.53)] [black (0.51)]
 [red (0.53)] [parking (0.69)] [parked (0.82)] [sign (0.78)]

a fire hydrant on the side of a road
 two signs with arrows pointing to each other for detour



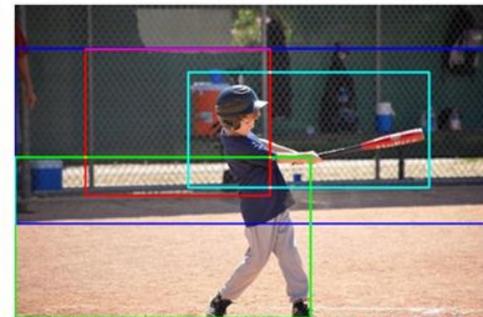
[steps (0.68)] [sitting (0.53)]
 [skateboard (0.99)] [skateboarder (0.76)] [doing (0.65)] [skate (0.64)] [ramp (0.54)] [board (0.51)]
 [street (0.79)] [riding (0.73)] [motorcycle (0.56)] [person (0.54)] [people (0.57)]
 [man (0.91)] [truck (0.76)] [parked (0.51)] [horse (0.53)] [truck (0.55)]

a man doing a trick on a skateboard
 a skateboarder is mid air performing a stunt



[monitors (0.56)]
 [laptop (0.97)] [table (0.74)] [open (0.71)] [sitting (0.61)]
 [station (0.52)]
 [desk (0.97)] [computer (0.94)] [keyboard (0.68)] [computers (0.65)]
 [tv (0.54)] [television (0.50)] [monitor (0.69)]

an open laptop computer sitting on top of a desk
 two computers are shown together on a desk



[tennis (0.65)] [holding (0.53)] [field (0.56)] [ball (0.79)] [court (0.52)] [boy (0.51)]
 [baseball (0.97)] [player (0.83)] [bat (0.82)] [man (0.80)] [playing (0.65)] [game (0.60)]

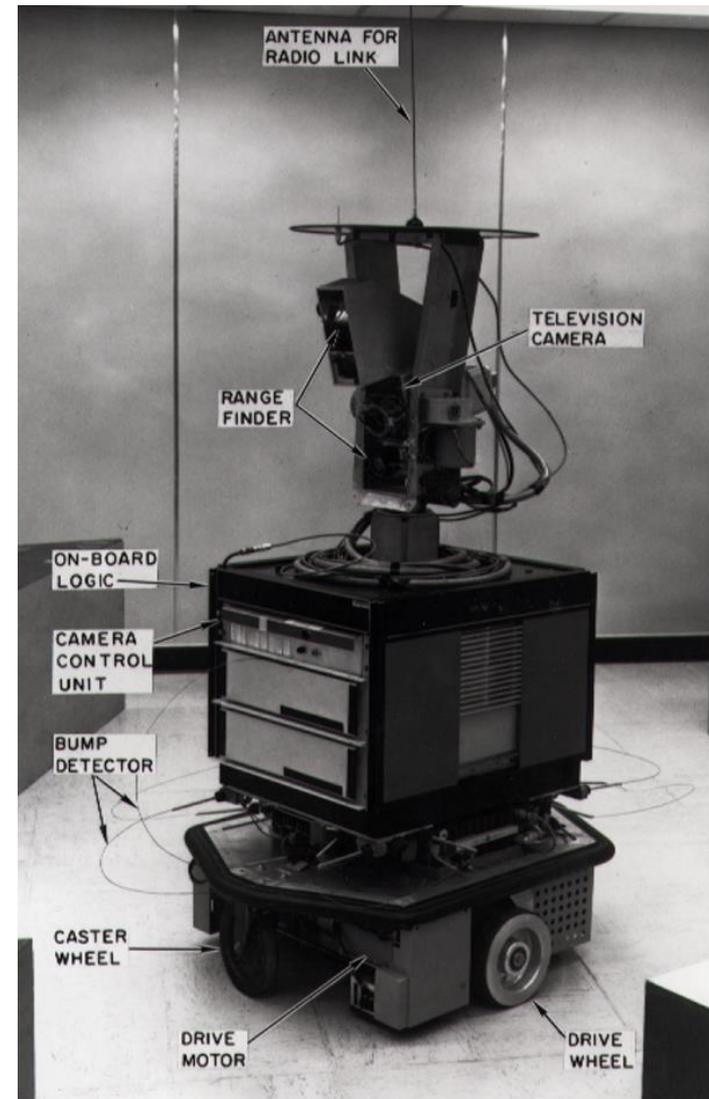
a baseball player swinging a bat at a ball
 a boy is playing with a baseball bat

Computer Vision

- Image Recognition
 - Identify image contents
 - YOLO
- Stereo reconstruction
 - Given 2 images, reconstruct the 3D scene
- Segmentation
 - Pick apart the pieces of the image

Robotics

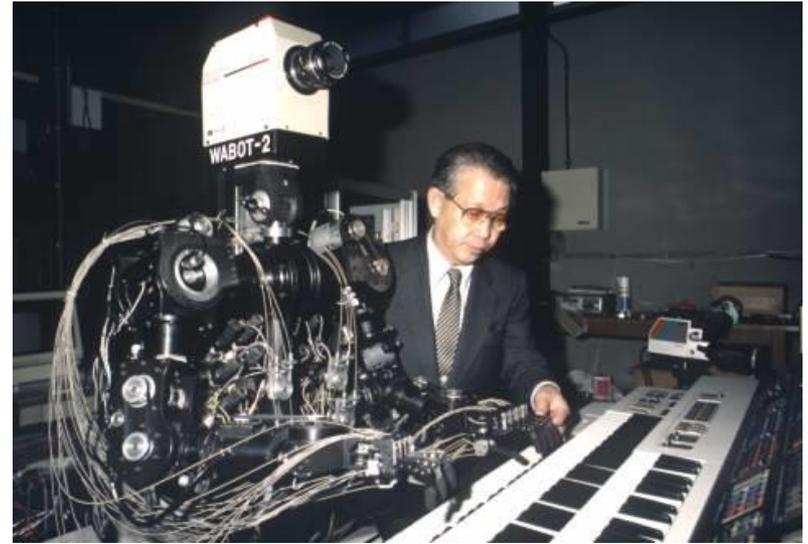
- 1966 - Shakey the robot
 - Stanford Research Institute (Now SRI international)
 - Simple computer vision
 - Navigation in multiple rooms
 - Blocks
 - Planning in STRIPS
 - Stanford Research Institute Planning System
 - Planning language & solver



Robotics

- 1967 – Waseda WABOT
 - First full-scale humanoid robot

- 1970s – Kuka Robotics
 - Used in automobile production



Robotics

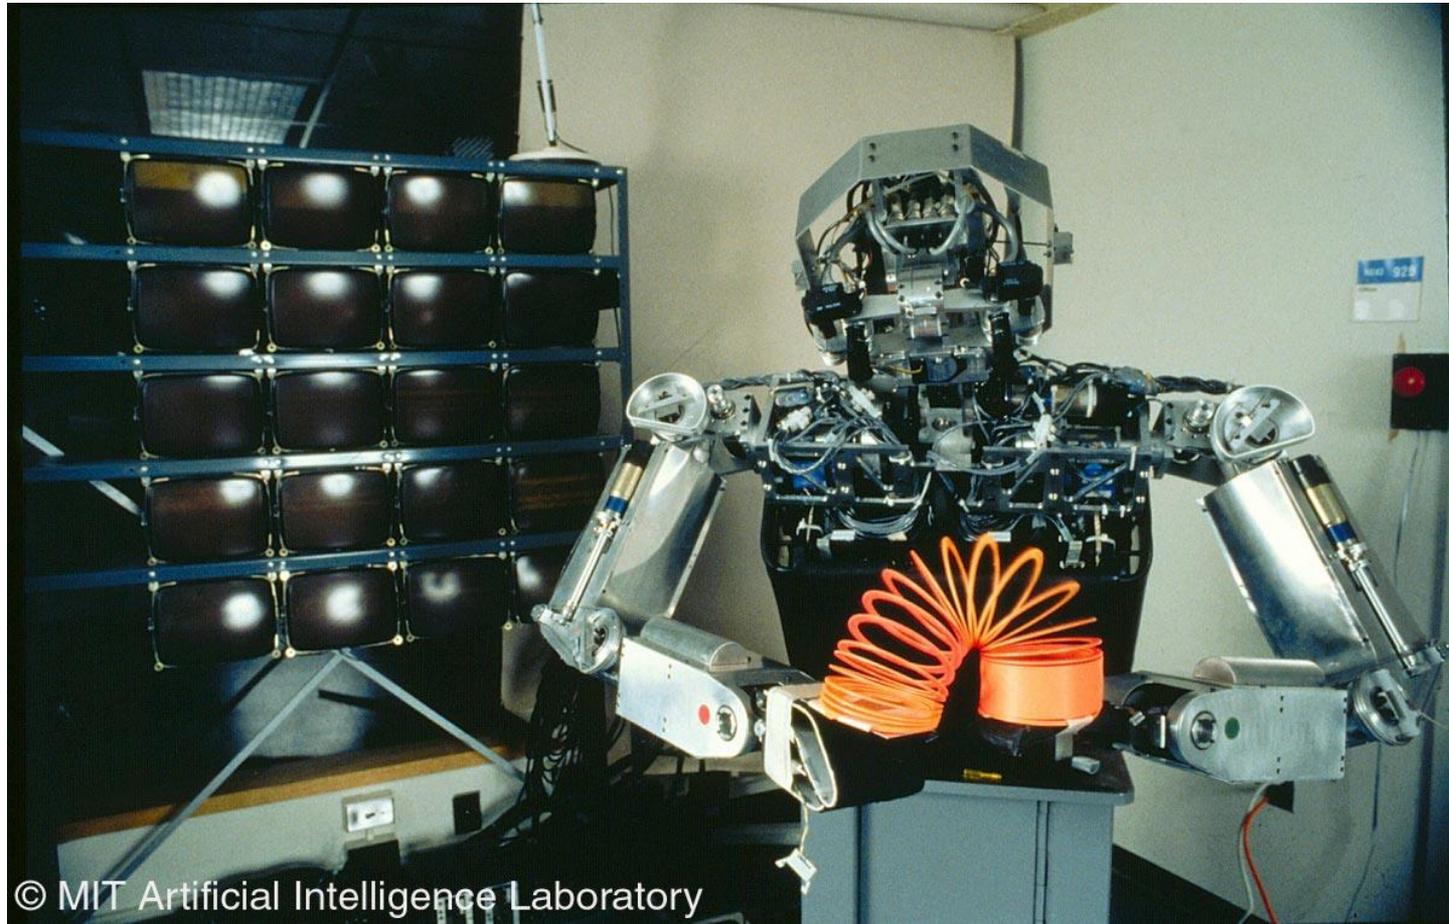
- 1989 – Ghengis
 - Inexpensive
 - Tested gait patterns

- 1995 – No Hands Across America
 - Mostly autonomous drive
 - CMU NavLab



Robotics

- Late 1990s – Cog
 - Attempt to emulate human-like intelligence & development



Robotics

- DARPA Grand Challenge
 - Autonomous vehicle race across Mojave Desert
 - Kicked off commercial autonomous vehicles



Robotics

- Androids
 - Geminoid
 - Erica







Robotics

- Asimo
- Toyota Human Support Robot (HSR)



Robotics

- Building-Wide Intelligence

