A Brief History of AI

Origins of AI

- McCulloch and Pitts: Threshold logic units (1943)

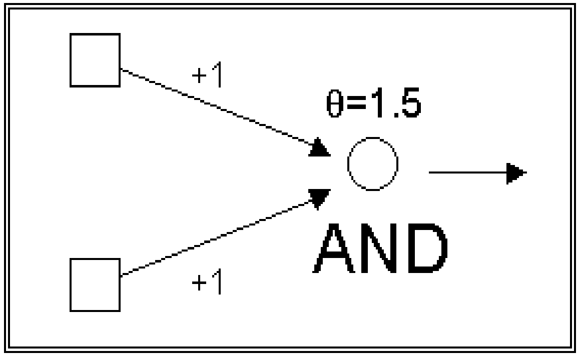

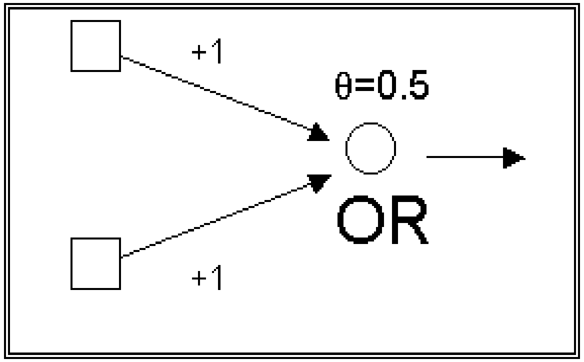

- Norbert Wiener: *Cybernetics: Or Control and Communication in the Animal and Machine* (1948)

- Donald Hebb: *The Organization of Behavior* (1949)

"When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place on one or both cells such that A's efficiency as one of the cells firing B, is increased."

Origins of AI

- Marvin Minsky: SNARC neural learning machine (1951)

  Stochastic Neural-Analog Reinforcement Calculator

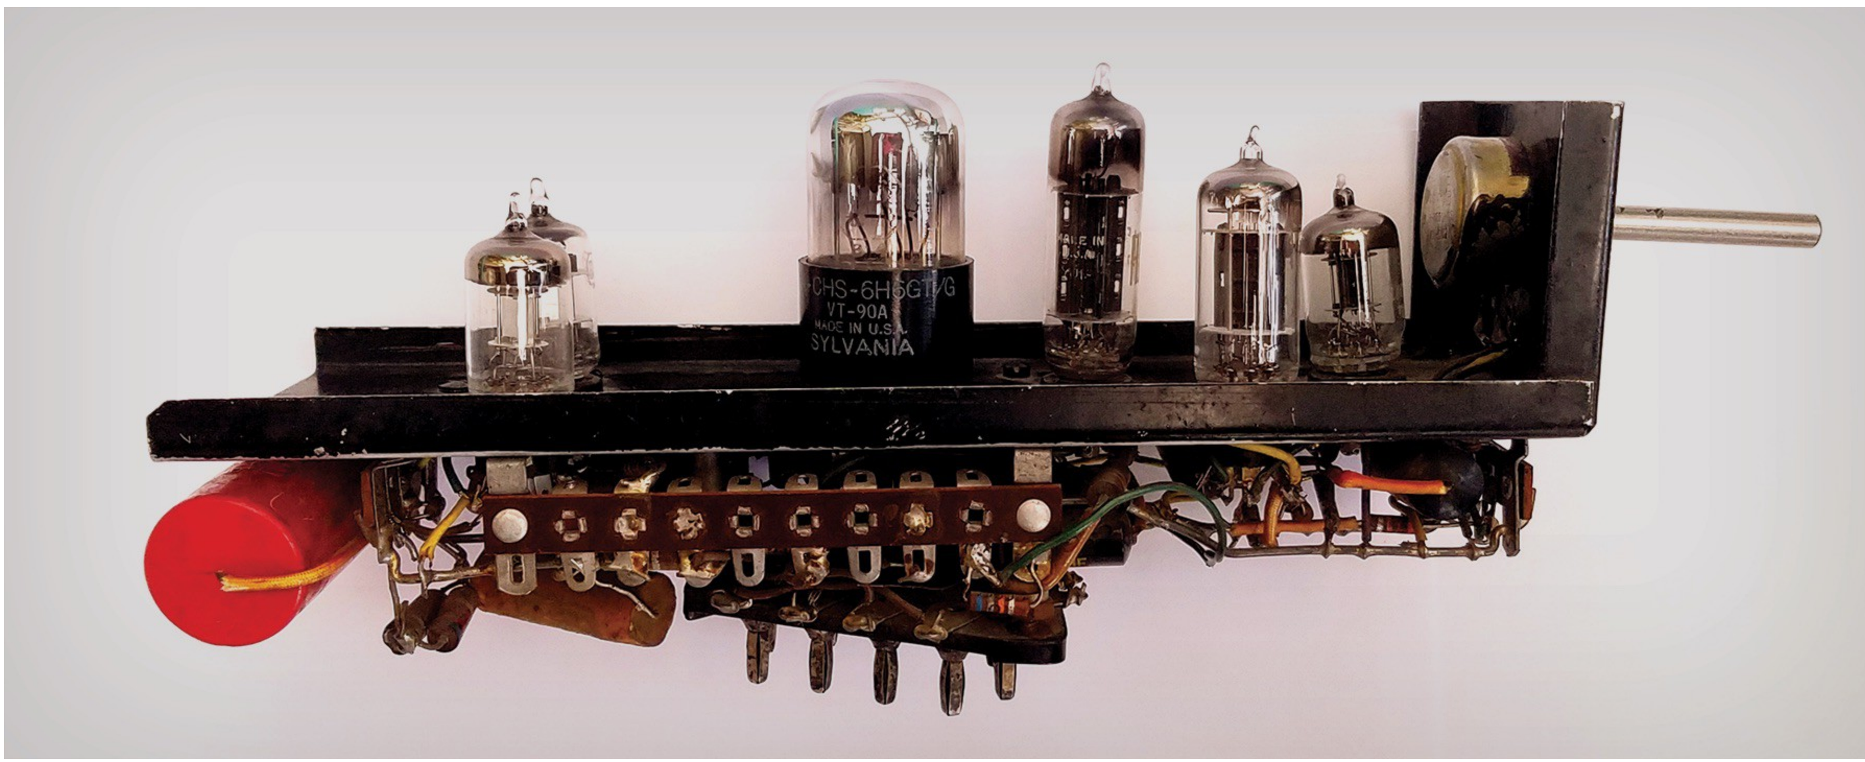

SNARC “neuron”

https://www.the-scientist.com/foundations/machine--learning--1951-65792

Origins of AI

- Marvin Minsky: SNARC neural learning machine (1951)

  Stochastic Neural-Analog Reinforcement Calculator

  "In the summer of 1951 Dean Edmonds and I went up to Harvard and built our machine. It had three hundred tubes and a lot of motors. It needed some automatic electric clutches, which we machined ourselves. The memory of the machine was stored in the positions of its control knobs — forty of them — and when the machine was learning it used the clutches to adjust its own knobs. We used a surplus gyropilot from a B-24 bomber to move the clutches. ... We sort of quit science for a while to watch the machine. We were amazed that it could have several activities going on at once in its little nervous system. Because of the random wiring, it had a sort of fail-safe characteristic. If one of the neurons wasn't working, it wouldn't make much of a difference — and, with nearly three hundred tubes and the thousands of connections we had soldered, there would usually be something wrong somewhere. ... I don't think we ever debugged our machine completely, but that didn't matter. By having this crazy random design, it was almost sure to work, no matter how you built it."

https://www.newyorker.com/magazine/1981/12/14/a-i

Origins of AI

- Alan Turing: *Computing Machinery and Intelligence* (1950)

- The Turing Test

  Language allows for in-depth probing of intelligence and consciousness in a machine

- Made groundbreaking contributions in many other fields

  - Theory of computation
  - Turing machines
  - Cryptography
  - Self-organization

Origins of AI

John von Neumann:
The Computer and the Brain

(published posthumously in 1958)

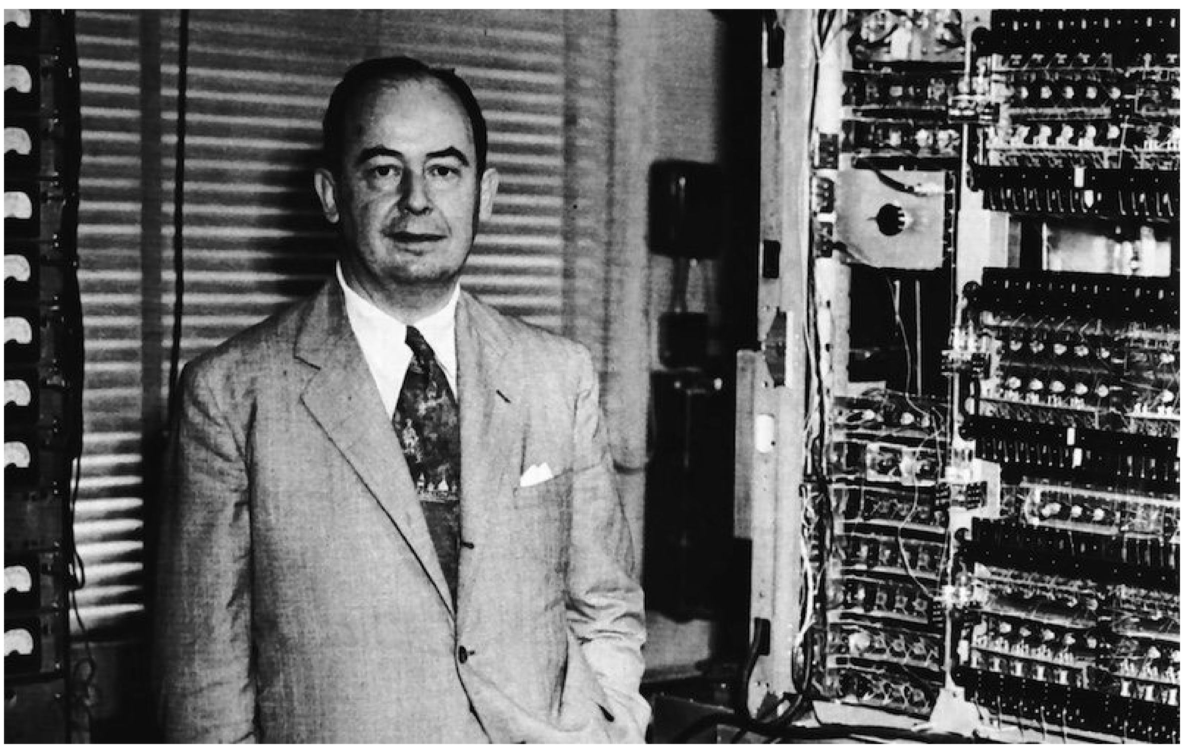

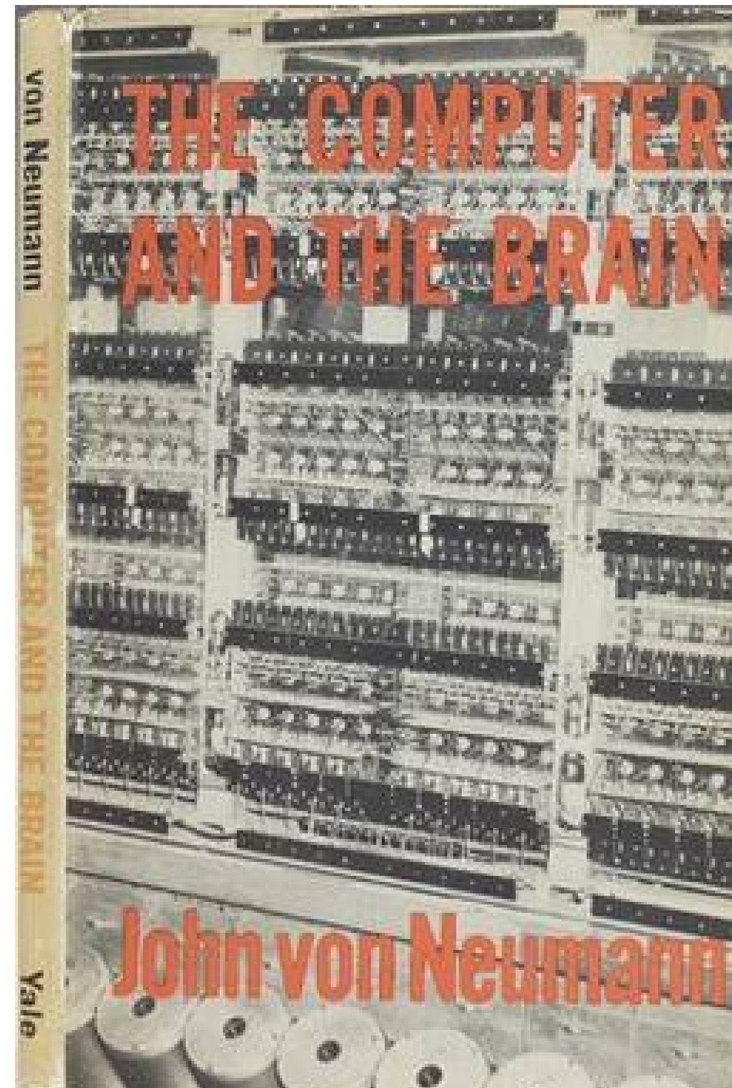

Origins of AI

1956 Summer Dartmouth Conference

- Organized by John McCarthy, who coined the term "artificial intelligence"
- Participants:

John McCarthy
Marvin Minsky
Claude Shannon
Nathaniel Rochester
Allen Newell
Herbert Simon
Oliver Selfridge
Trenchant More
Arthur Samuel
Ray Solomonoff

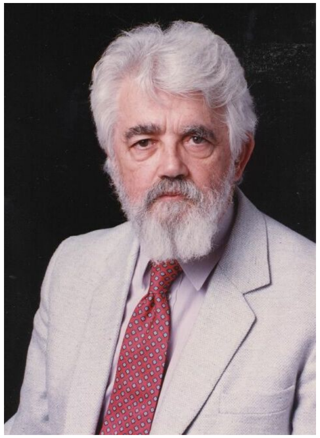

McCarthy

Minsky

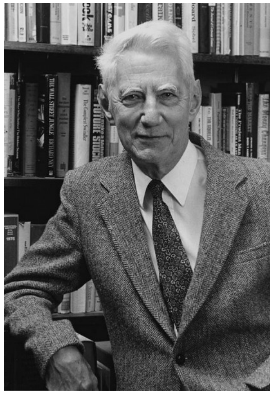

Shannon

Simon and Newell

Samuel (standing)

Early Developments

Checkers

- Arthur Samuel, IBM (1959)
  - First successful **reinforcement learning** program
  - Learned to play checkers better than Samuel himself
  - Beat 4th ranked player in the nation in 1961

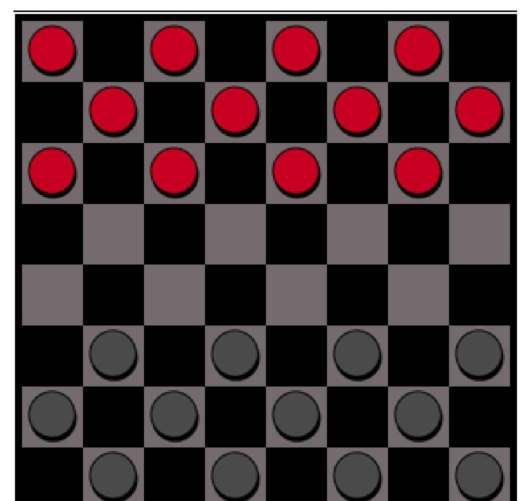

Early Developments

Perceptrons

- Frank Rosenblatt (late 1950s, early 1960s)

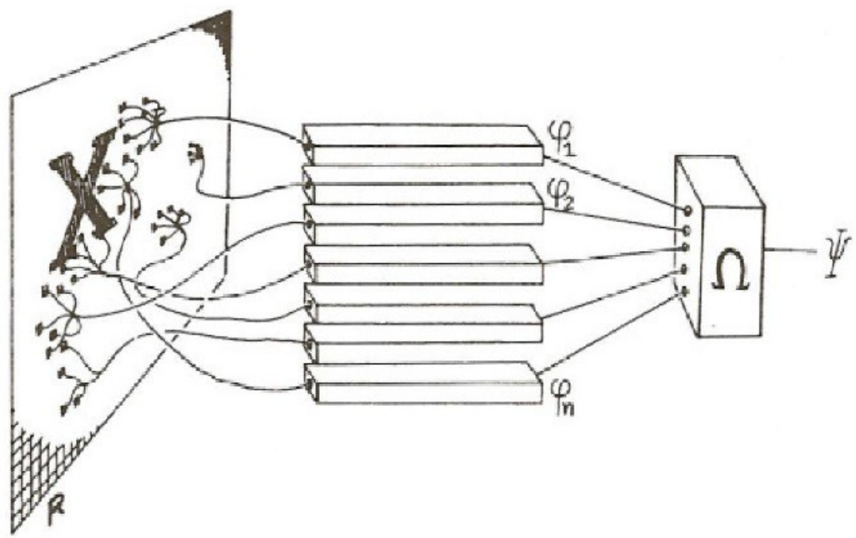

- Learned to classify “visual” patterns by adjusting weights

- Perceptron Learning Procedure

Early Developments

Perceptrons

- Single-layer perceptrons could not learn simple classification problems such as XOR

0	**0**	⇒	**0**
0	**1**	⇒	**1**
1	**0**	⇒	**1**
1	**1**	⇒	**0**

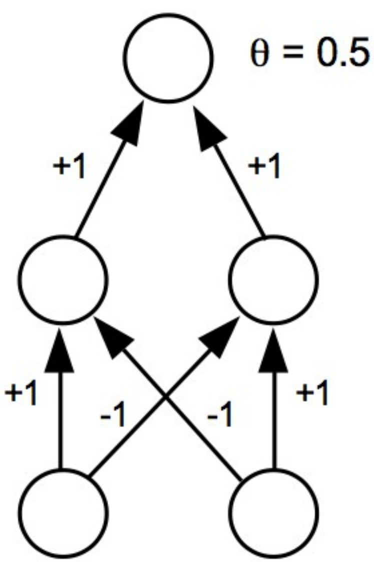

- Multi-layer perceptrons **could** solve XOR in principle, but no learning algorithm was known

Early Developments

Perceptrons

- Marvin Minsky and Seymour Papert published *Perceptrons* in 1969

[The perceptron] has many features to attract attention: its linearity; its intriguing learning theorem; its clear paradigmatic simplicity as a kind of parallel computation. ***There is no reason to suppose that any of these virtues carry over to the many-layered version. Nevertheless, we consider it to be an important research problem to elucidate (or reject) our intuitive judgment that the extension is sterile.***

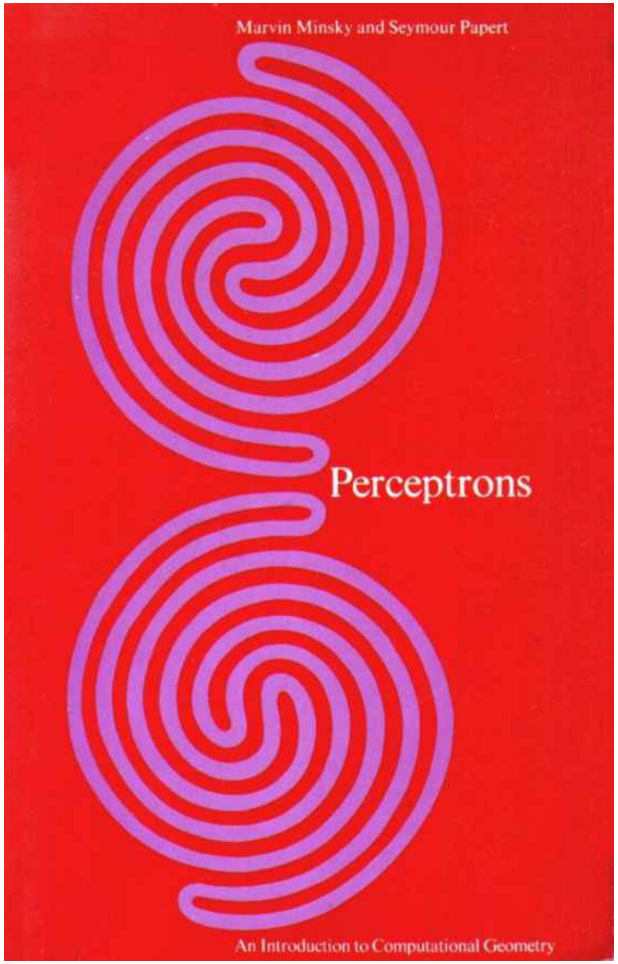

Early Developments

Heuristic search

- Allen Newell and Herbert Simon
  - Logic Theorist (1955)
  - General Problem Solver (1959)

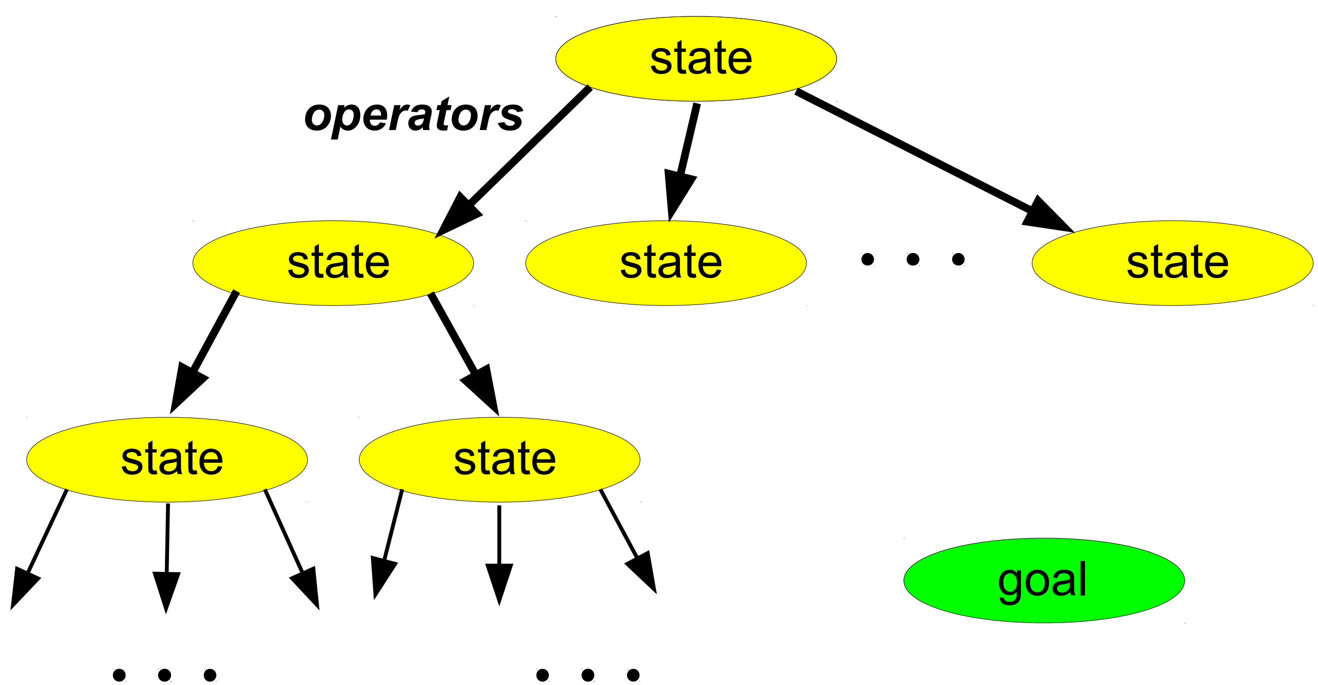

Heuristic Search: The 8-Puzzle

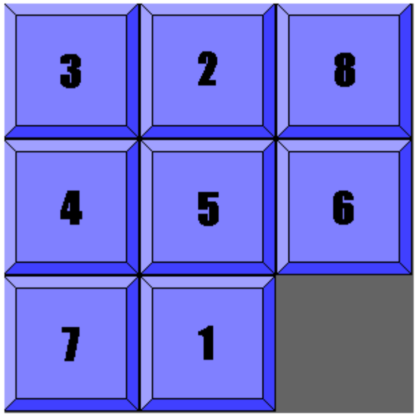

Initial state

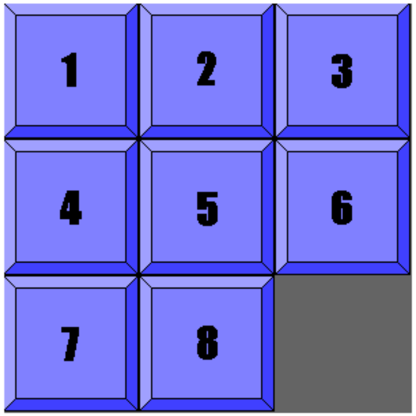

Goal state

- State representation: **(3 2 8 4 5 6 7 1 _)**

- Operators: *Up, Down, Left, Right*

  (3 2 8 4 5 6 7 1 _) $\underset{\textit{Down}}{\Rightarrow}$ **(3 2 8 4 5 _ 7 1 6)**

- 8-puzzle has 9! possible states = 362,880
- 15-puzzle has 16! possible states = 20,922,789,888,000

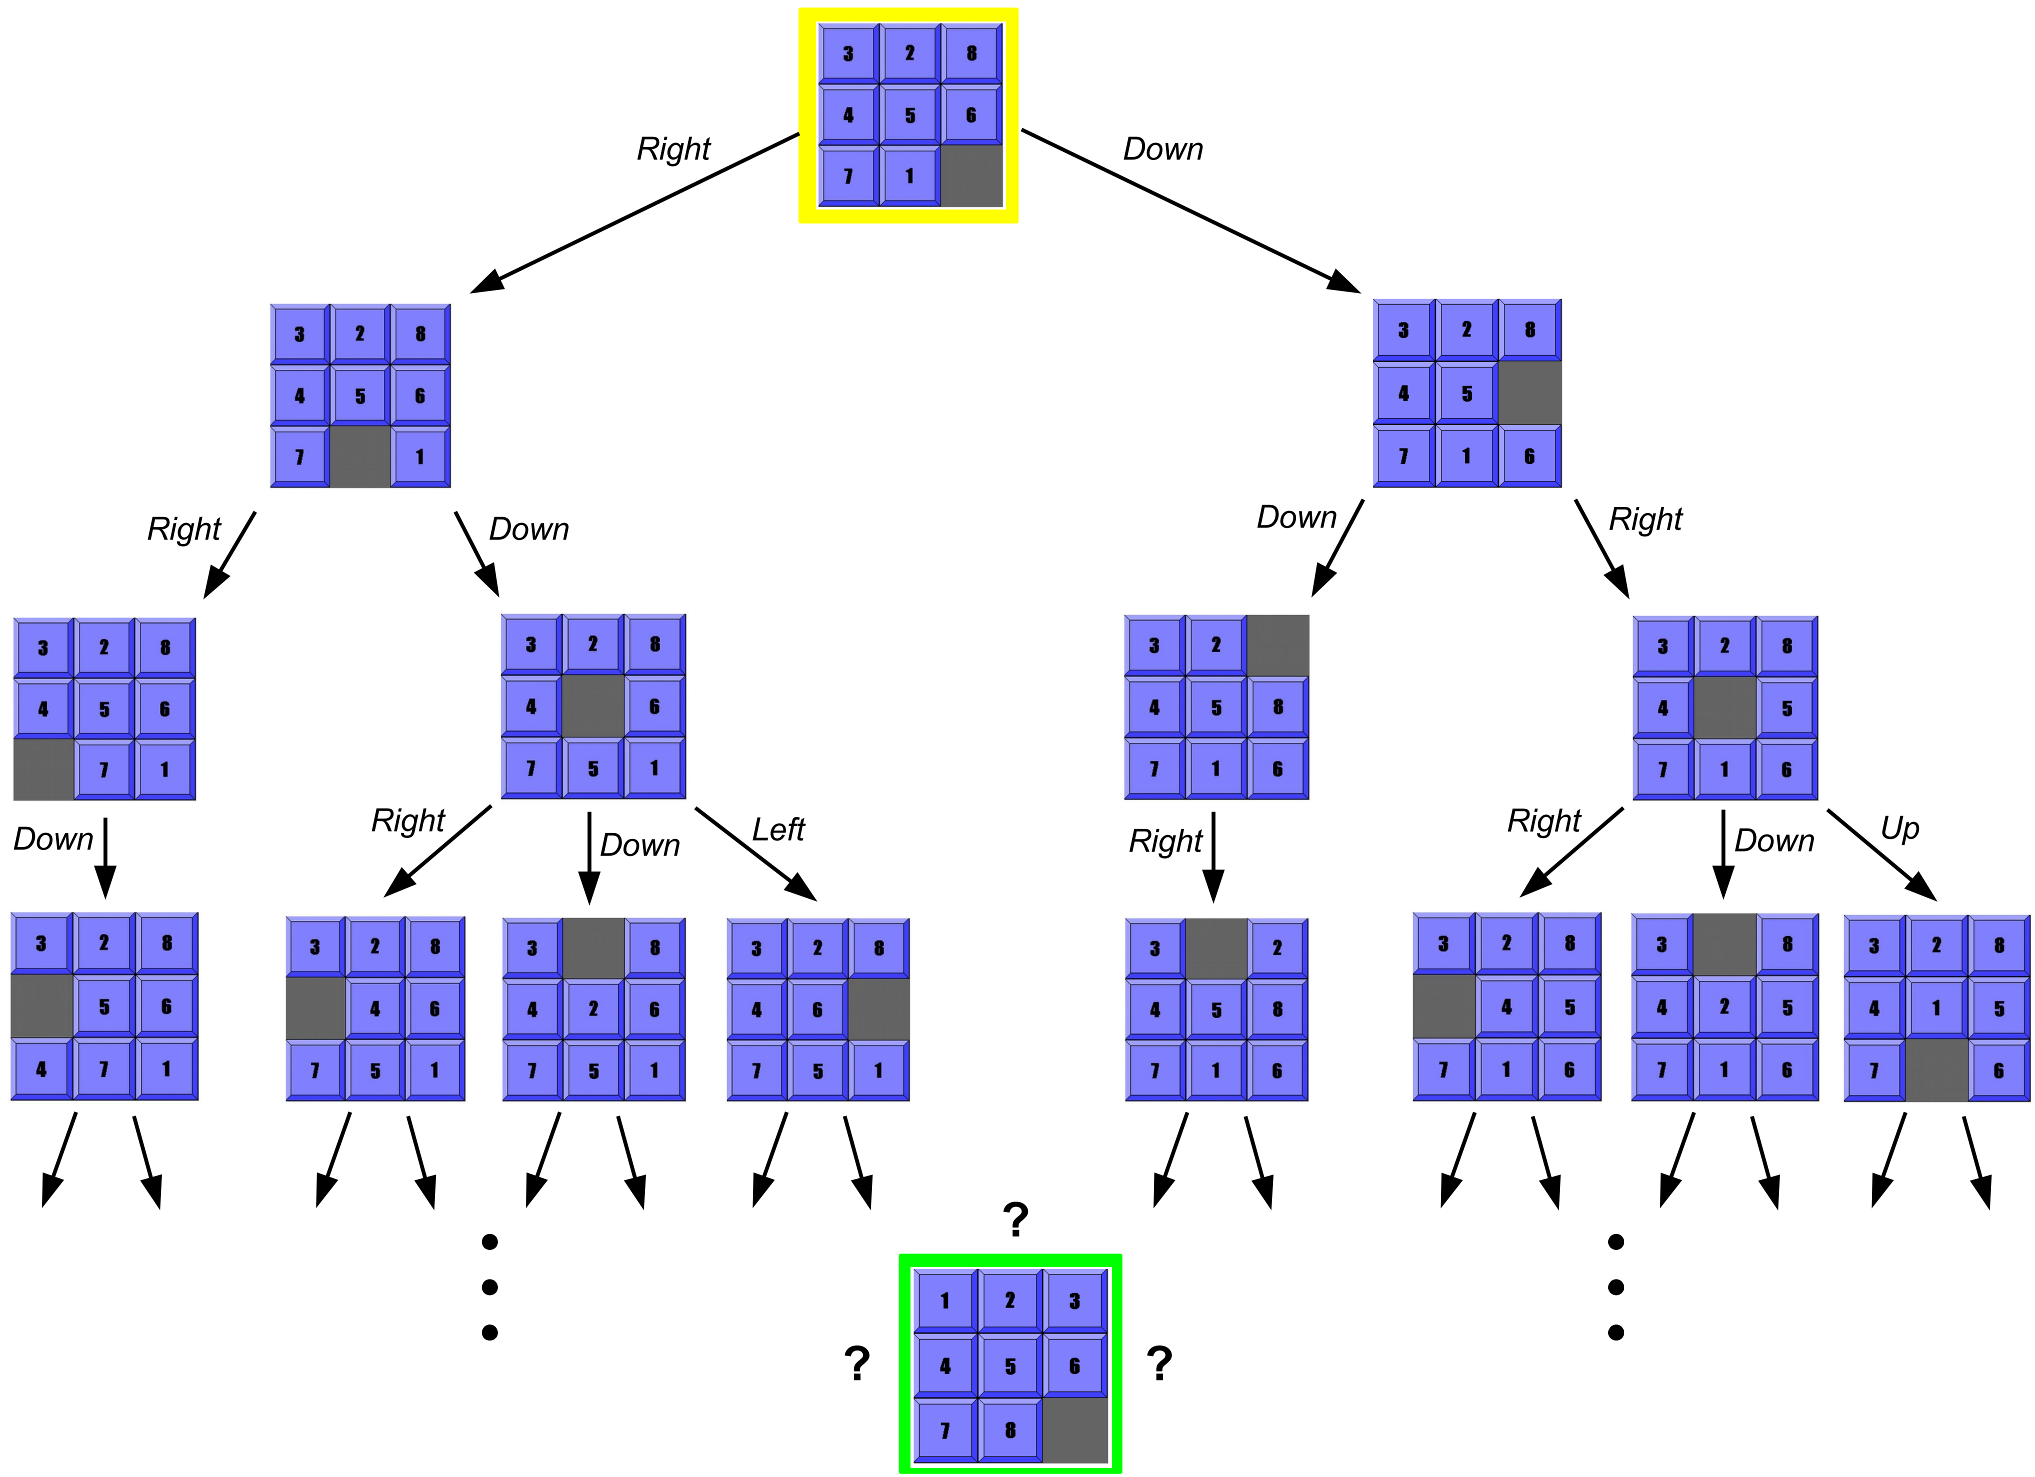
Right
Down
Right
Down
Down
Right
Down
Right
Down
Left
Right
Right
Down
Up
?
?
?

Symbolic AI Era (or GOFAI): 1970s - 1980s

- Heuristic search
- Game playing (chess, checkers, backgammon, etc.)
- Automated theorem proving (AM, Eurisko)
- Logic puzzles
- Planning
- Natural language processing (SHRDLU)
- Common-sense reasoning (CYC)
- Expert systems

  DENDRAL: inferred molecular structure from mass spectroscopy data
  MYCIN: diagnosed blood infections
  XCON: configured computer systems
  PROSPECTOR: analyzed geological data to identify oil and mineral deposits

The Return of Neural Networks: 1980s

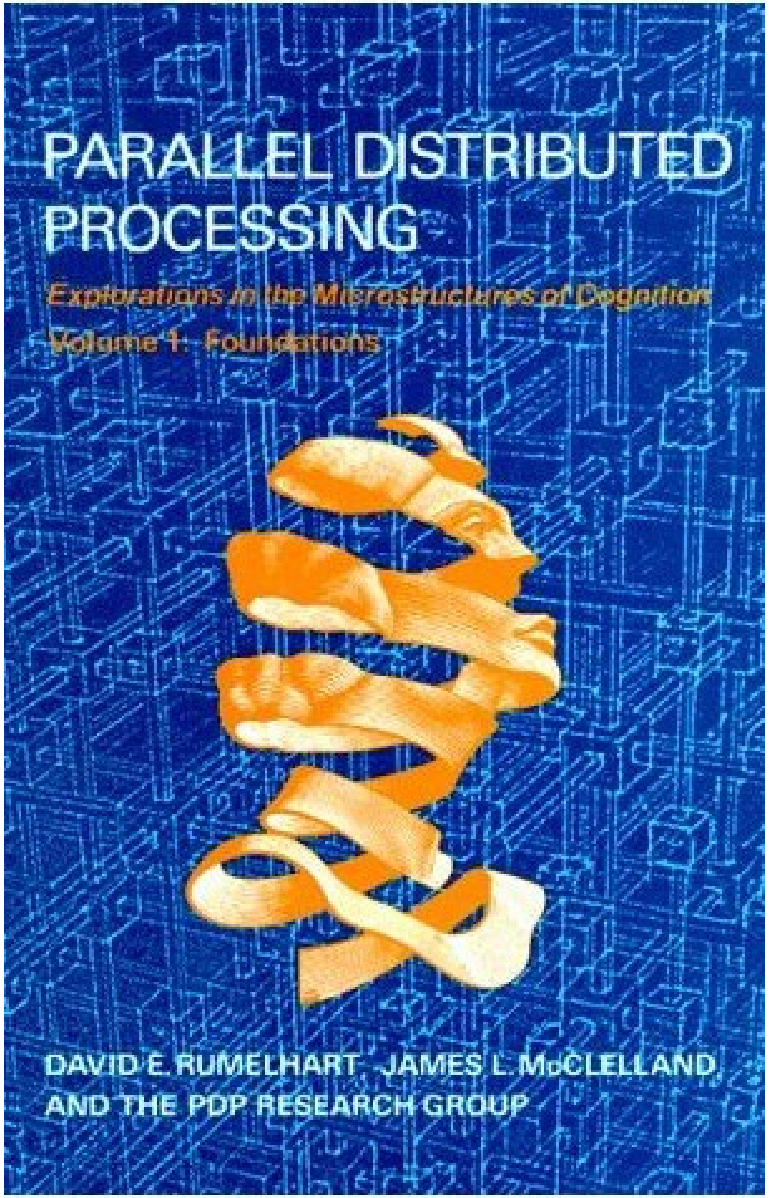

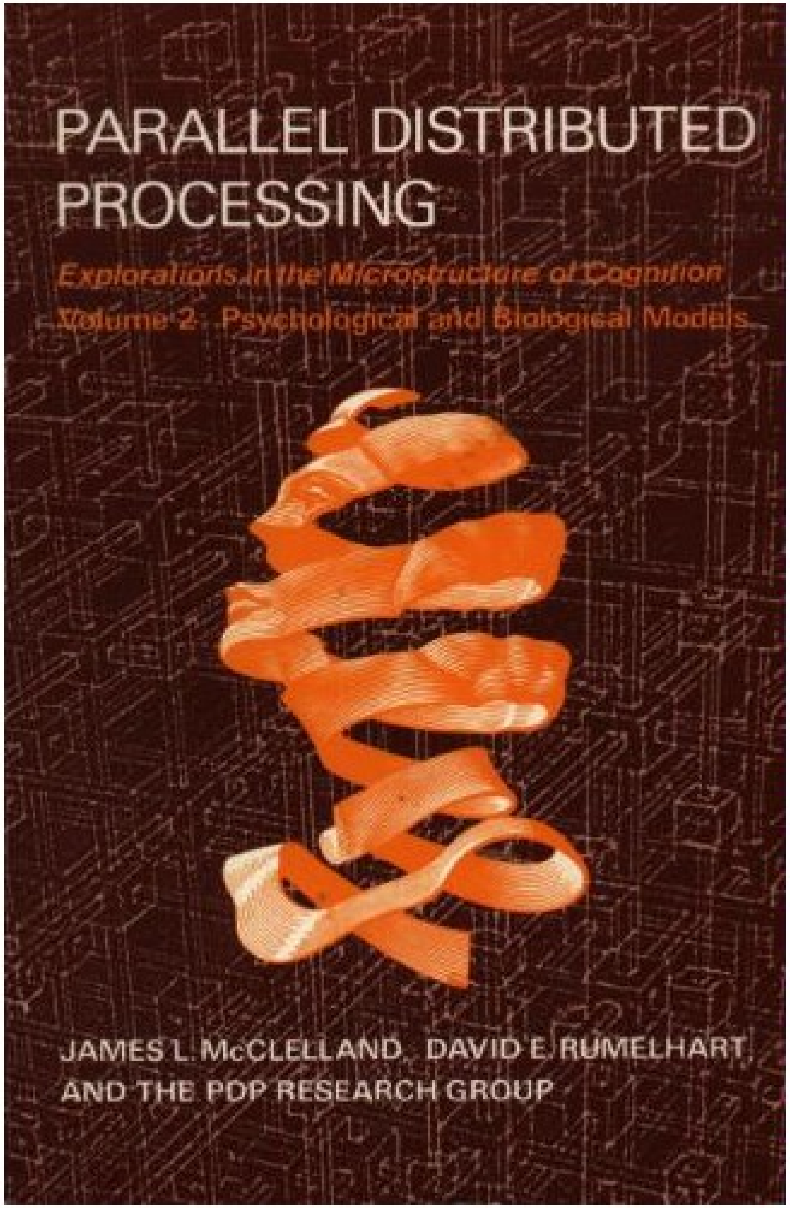

The Return of Neural Networks: 1980s

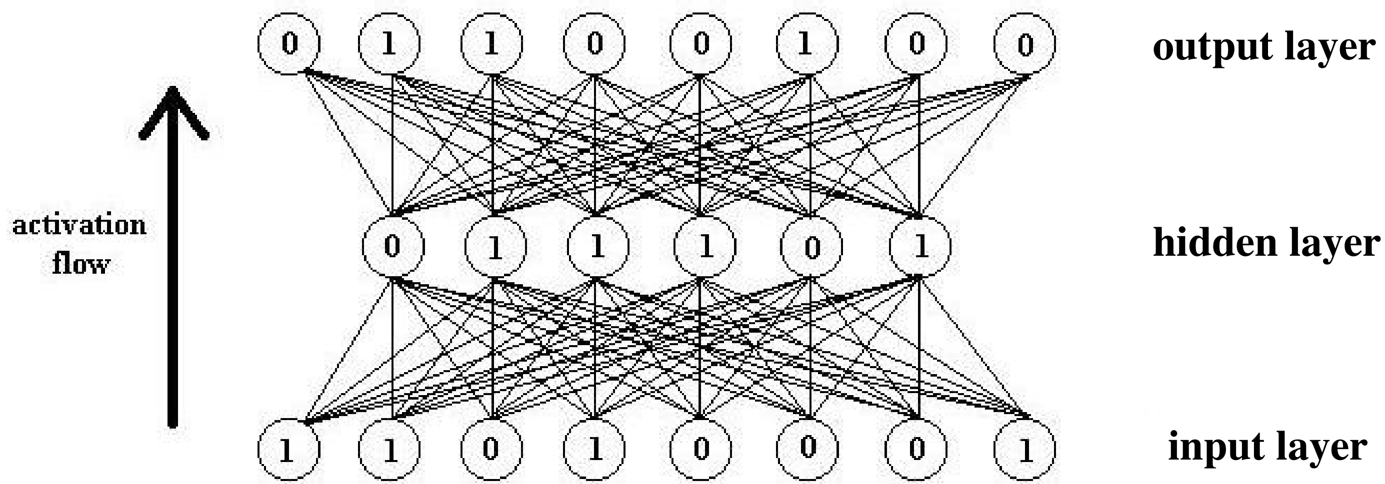

- **Connectionist** architecture
- Each neuron has an **activation** level (typically between 0 and 1)
- Each connection has a **strength** value (positive or negative)

The Return of Neural Networks: 1980s

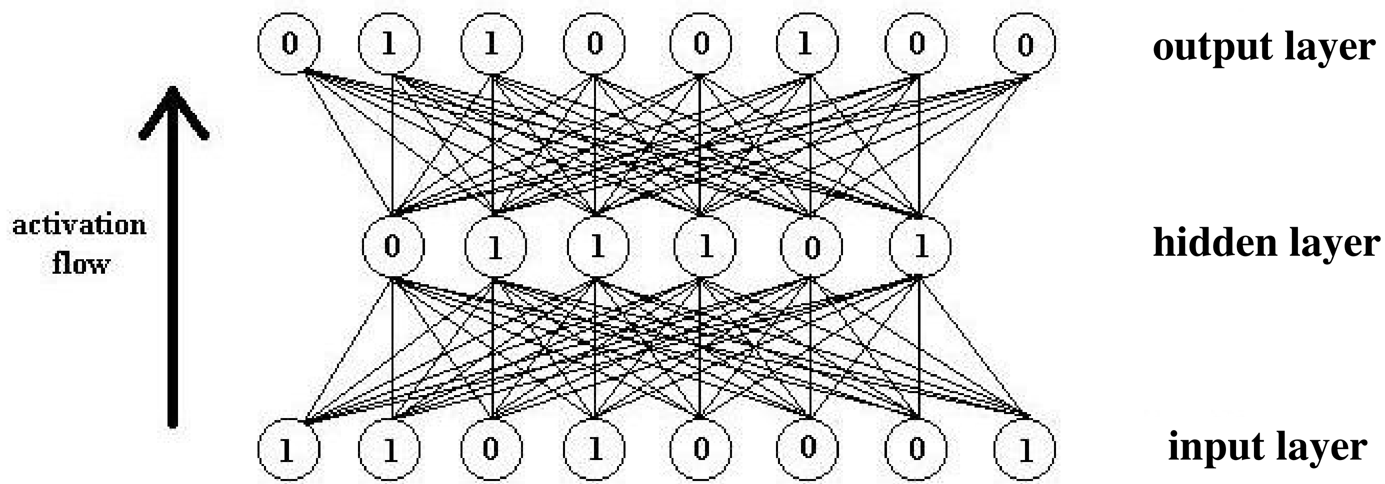

- **Backpropagation** learning algorithm
- Solved the problems pointed out by Minsky and Papert
- Many other neural net models and learning algorithms developed

The Connectionist Era: 1980s - 1990s

- Backpropagation applied mostly to small-scale problems

- Most neural networks had only **one hidden layer** and relatively small numbers of neurons and connections

- Backpropagation was **very slow** for networks with many hidden layers

- In theory, **any function** can be represented by a network with only one hidden layer, so why bother with many-layer networks?

ALVINN (1990s)

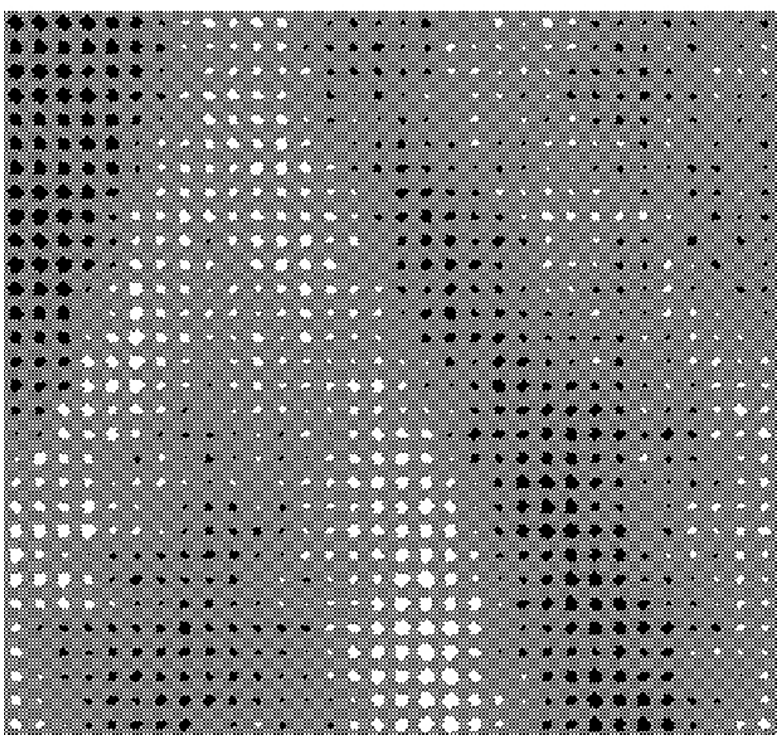

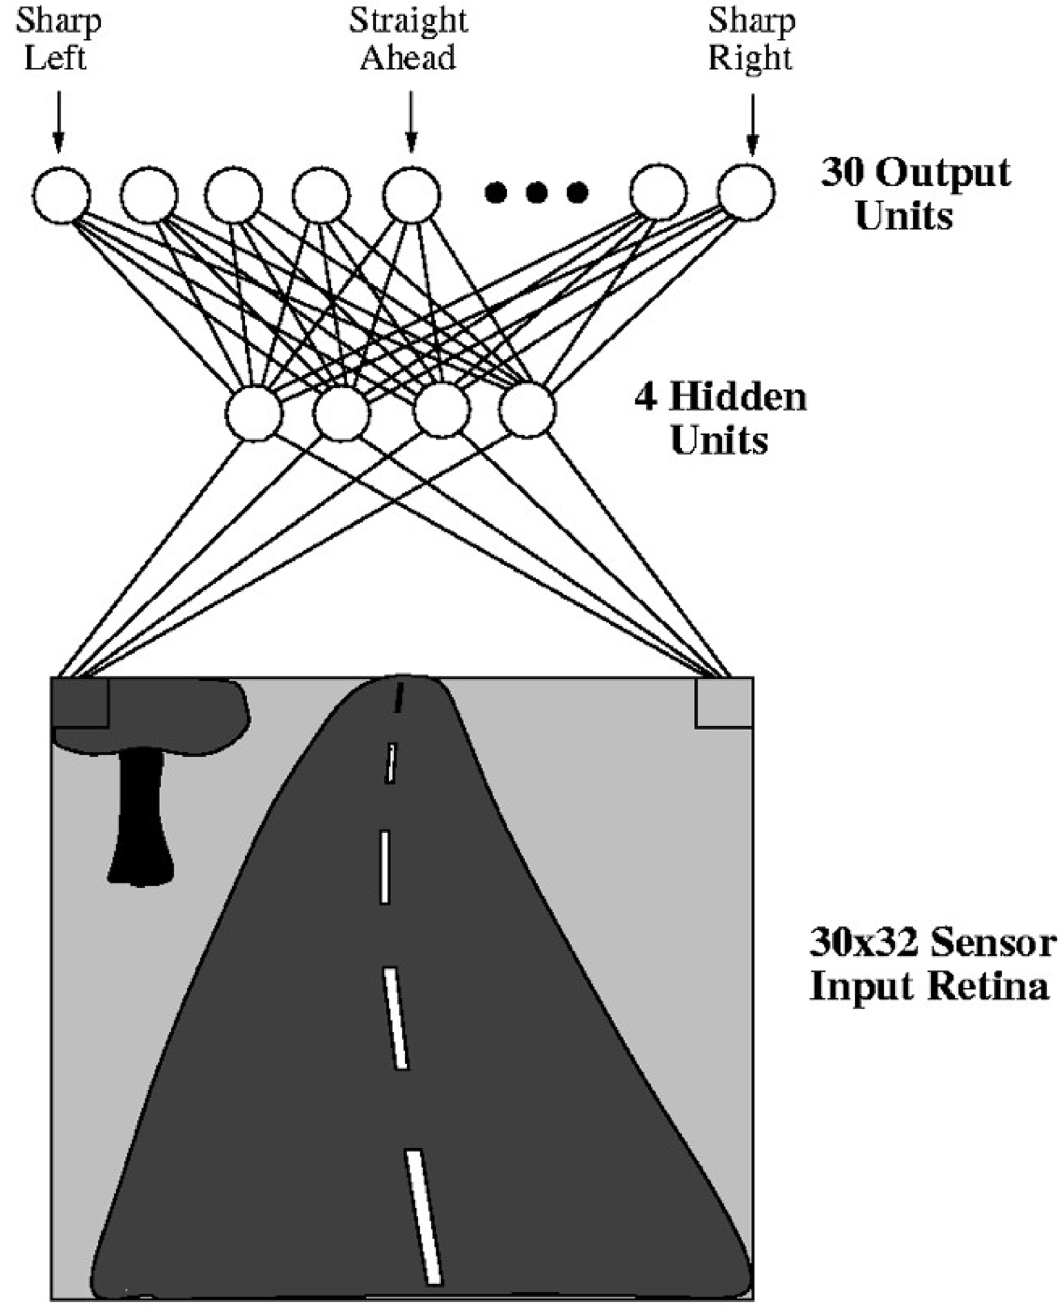

An Early Convolutional Neural Network

- **Zip code recognizer** (Yann LeCun, AT&T Bell Labs, 1980s)
  - Used a large neural network with several layers
  - Trained on handwritten zip codes from U.S. mail
  - Achieved the state of the art in digit recognition
  - Classification accuracy > 95%

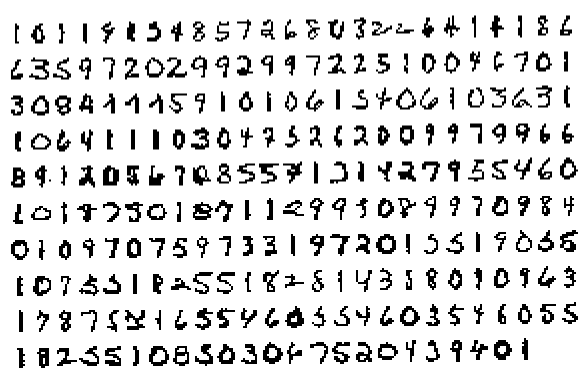

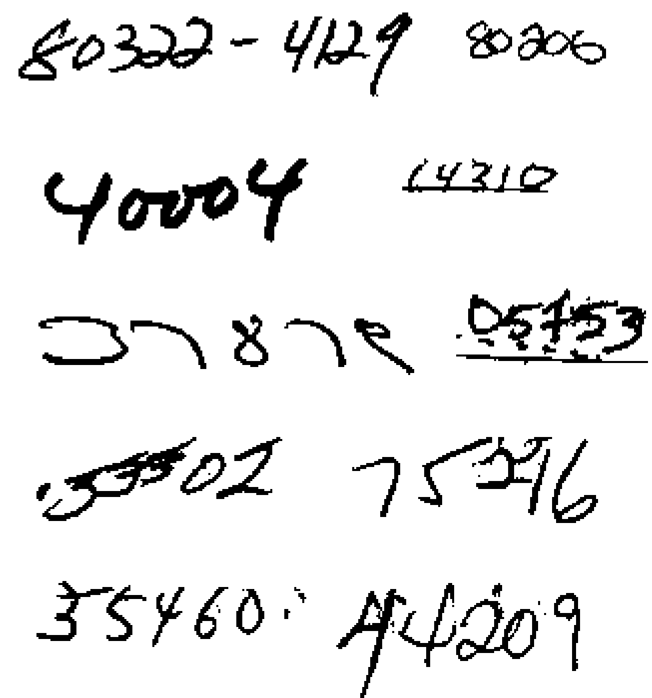

An Early Convolutional Neural Network

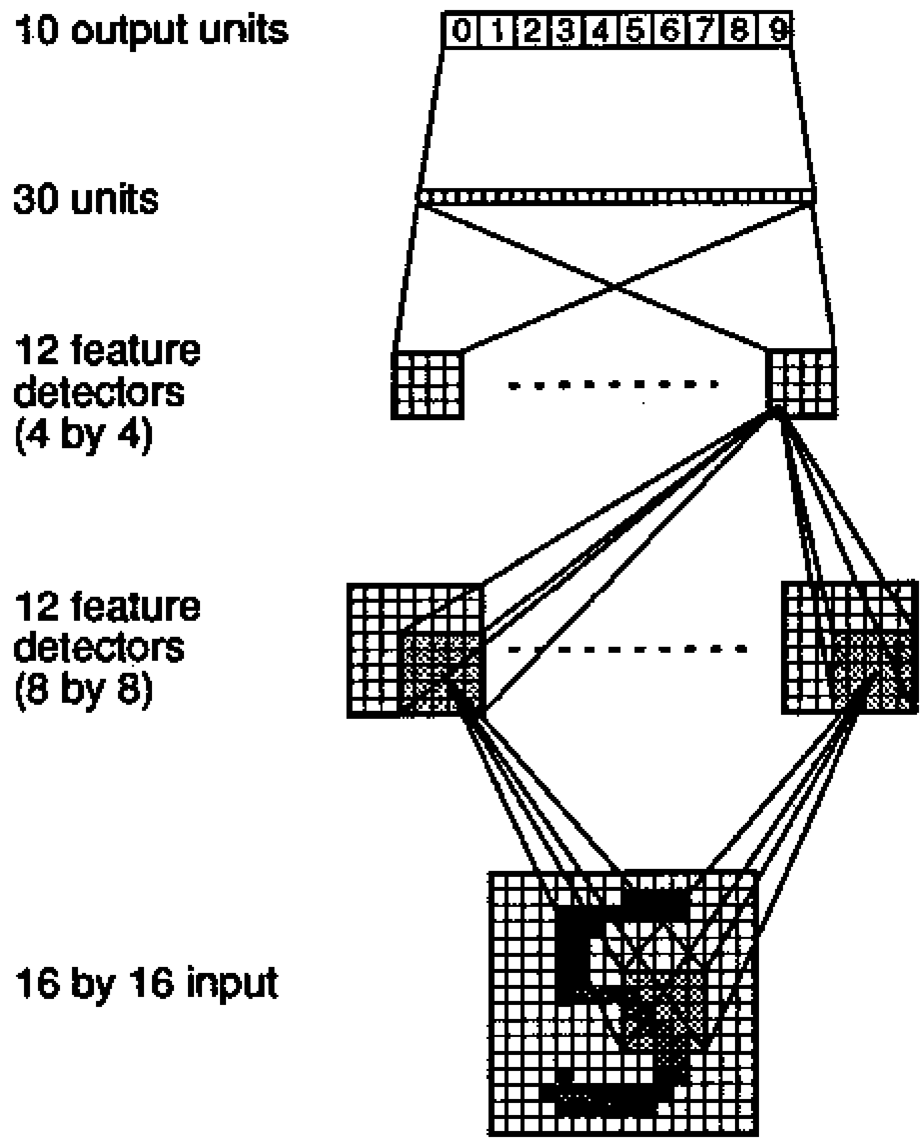

Other Developments

- IBM's **Deep Blue** became the world chess champion on May 11, 1997 by defeating Garry Kasparov (3.5 to 2.5)

- **Deep Junior** held its own against Kasparov in a February 2003 rematch (3 to 3)

Other Developments

- IBM's **Watson** answered *Jeopardy!*-style questions posed in natural language

- Watson beat human *Jeopardy!* champions Ken Jennings and Brad Rutter in 2011 for a $1 million prize

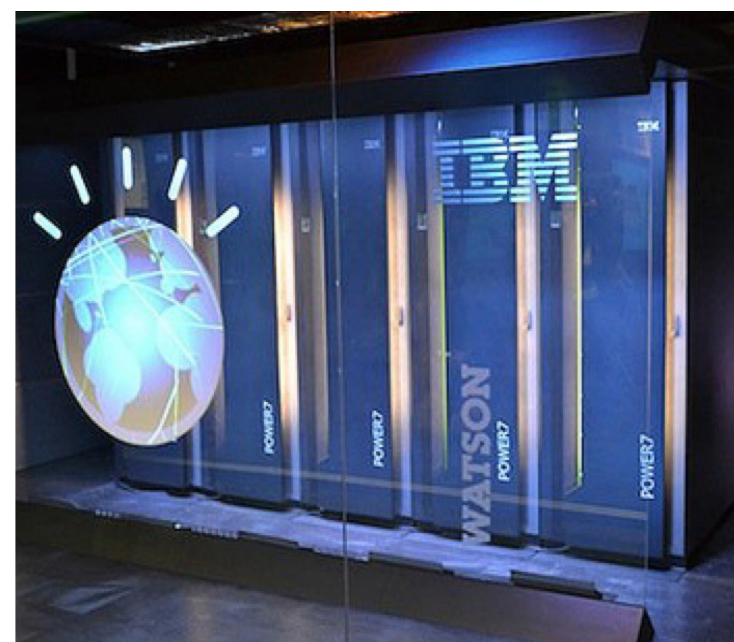

Other Developments

- **John Holland** and his students at the University of Michigan develop **genetic algorithms**
- Solutions to computational problems are **evolved** by a process analogous to natural selection
- **John Koza** develops **genetic programming** by using GAs to directly evolve computer programs
- Some GA-evolved solutions to problems are better than the best solutions designed by human engineers (e.g., the design of antennas for NASA satellites)

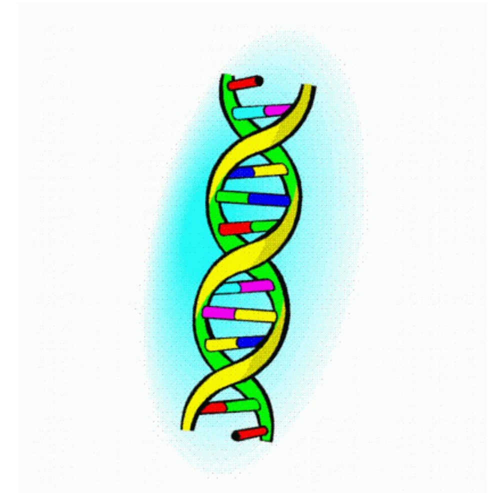

Other Developments

- **TD-Gammon** (Gerry Tesauro, IBM, 1990s)
  - Learned to play backgammon
  - Trained by playing over 1.5 million games against itself using **reinforcement learning**
  - Discovered novel board evaluation strategies
  - Achieved parity with the top 5-10 players in the world

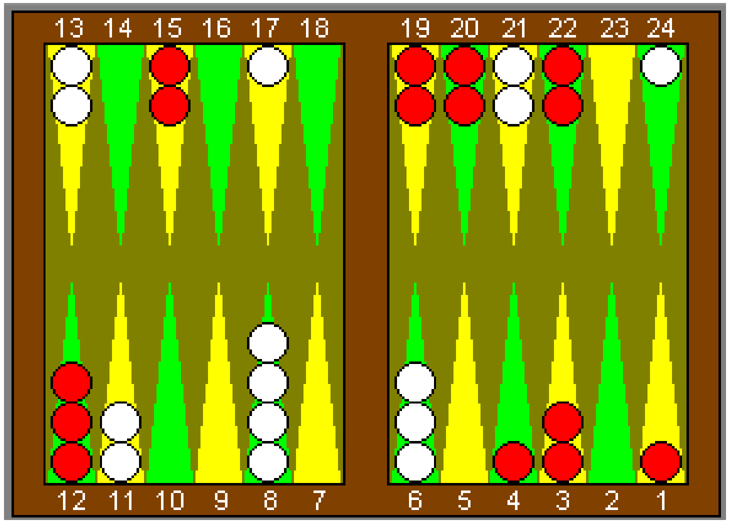

Some Early Robots

- W. Grey Walter's **Tortoises** (1950's)

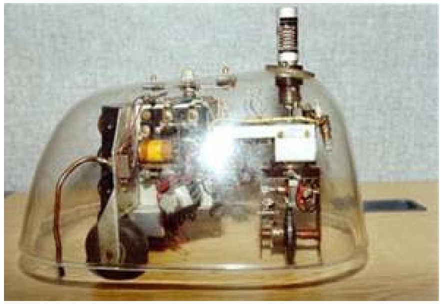

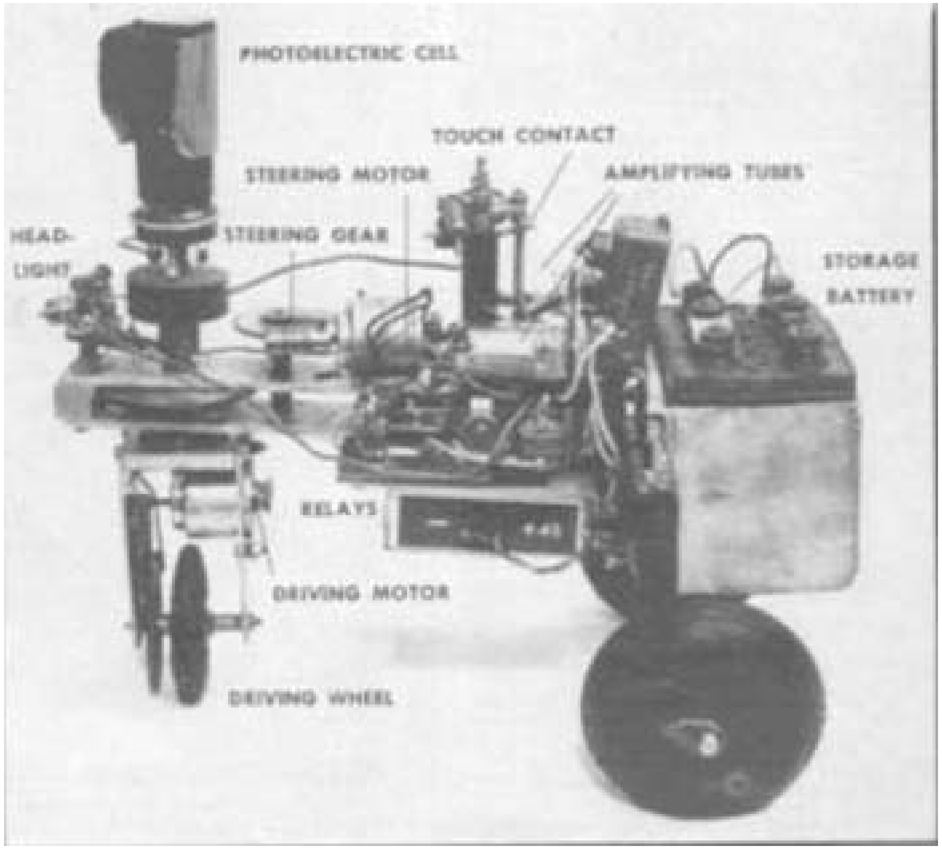

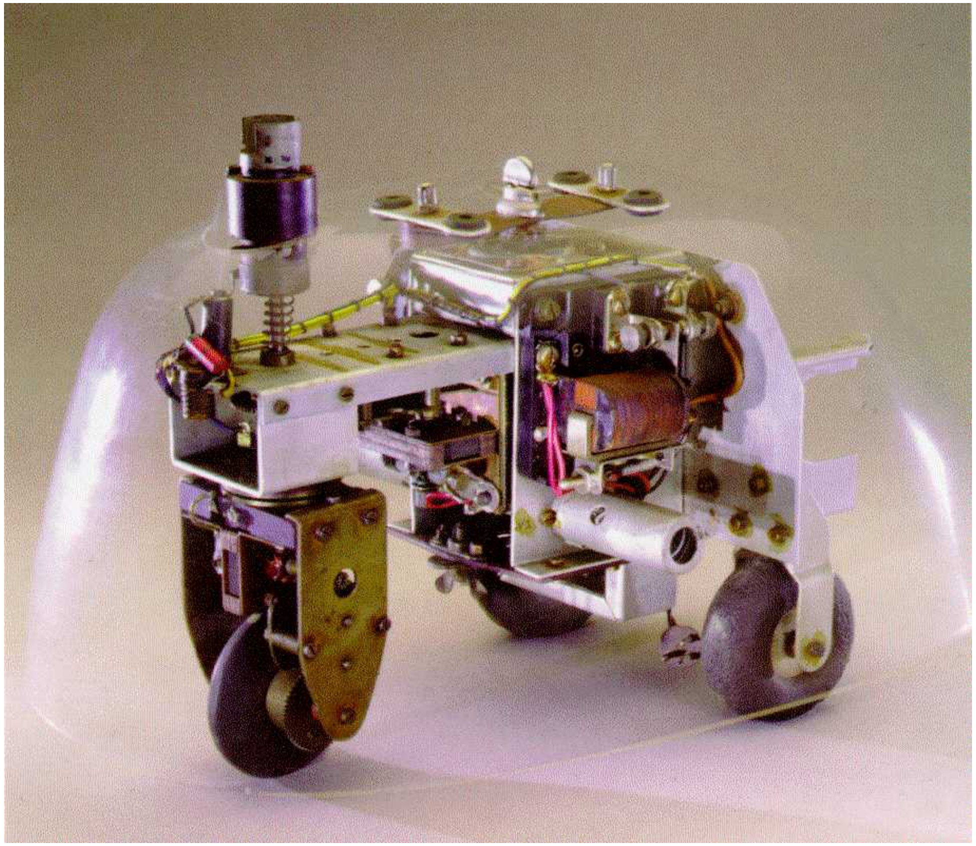

Some Early Robots

- W. Grey Walter's **Tortoises** (1950's)

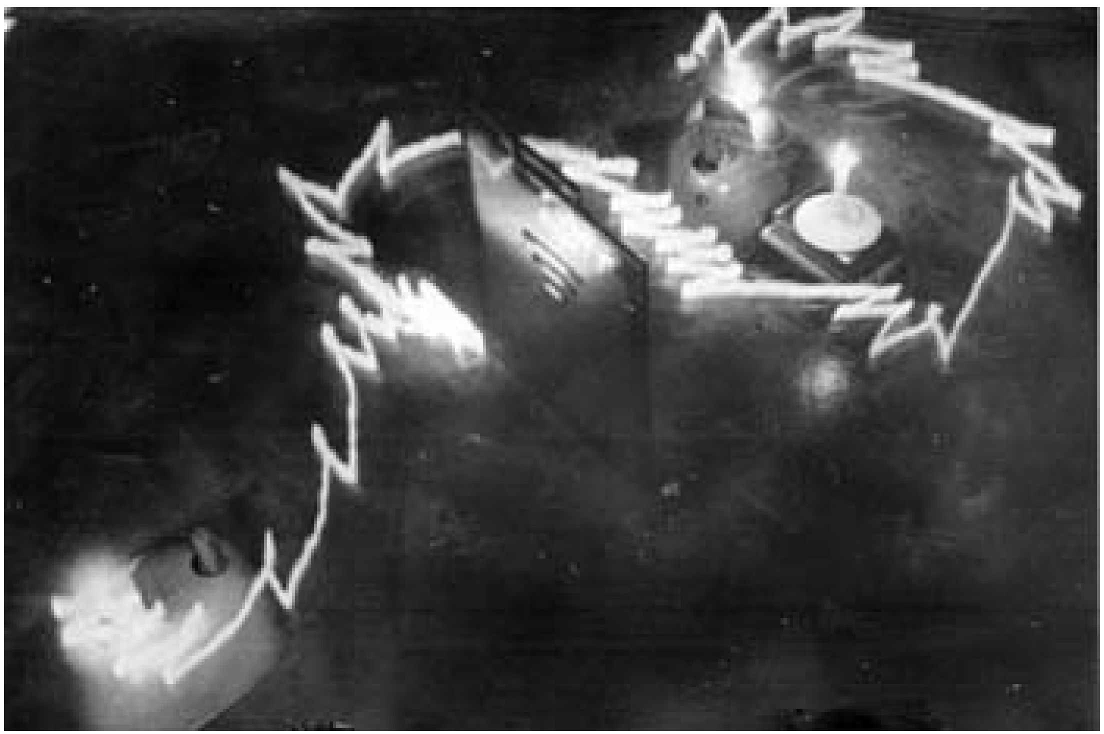

Some Early Robots

- **Shakey**
- Developed at Stanford (1969)
- Bump sensors
- Camera
- Lived in a special indoor world with a white floor and black objects (balls, pyramids, etc.)

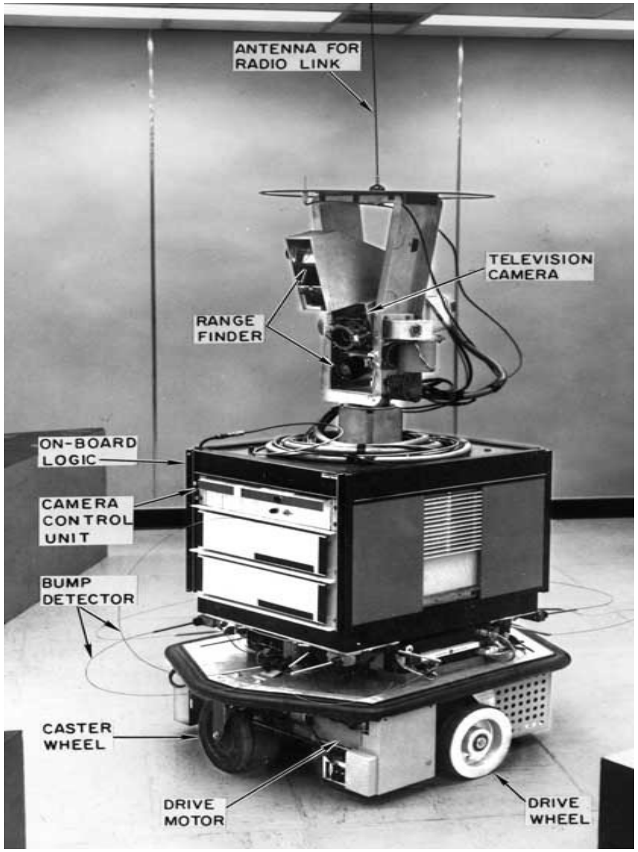

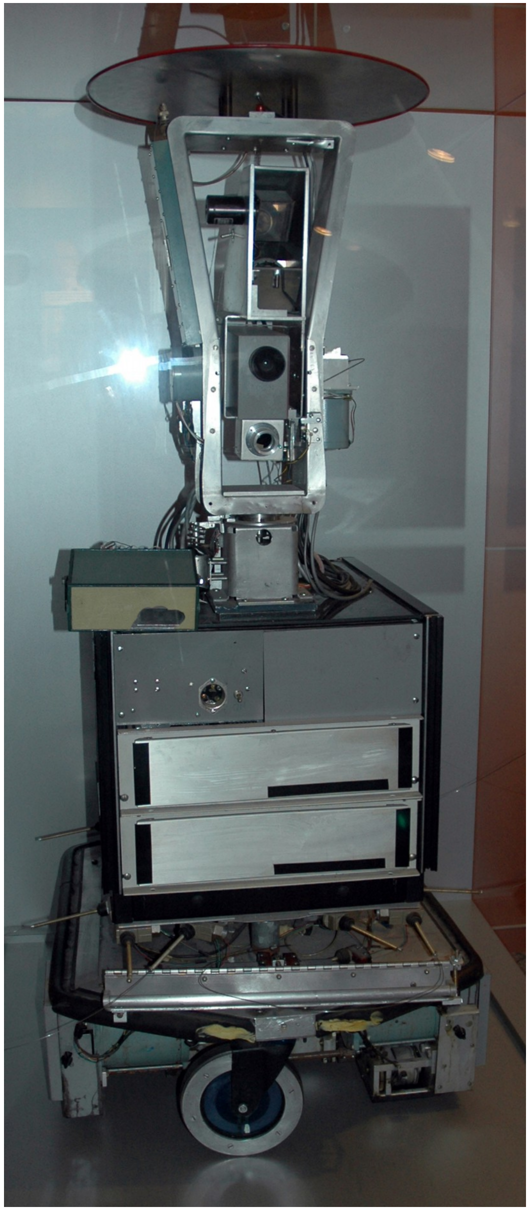

Some Early Robots

- **Stanford Cart** (1977)
- Developed by Hans Moravec
- Vision-based navigation
- Path planning
- Operated in “Cartland”

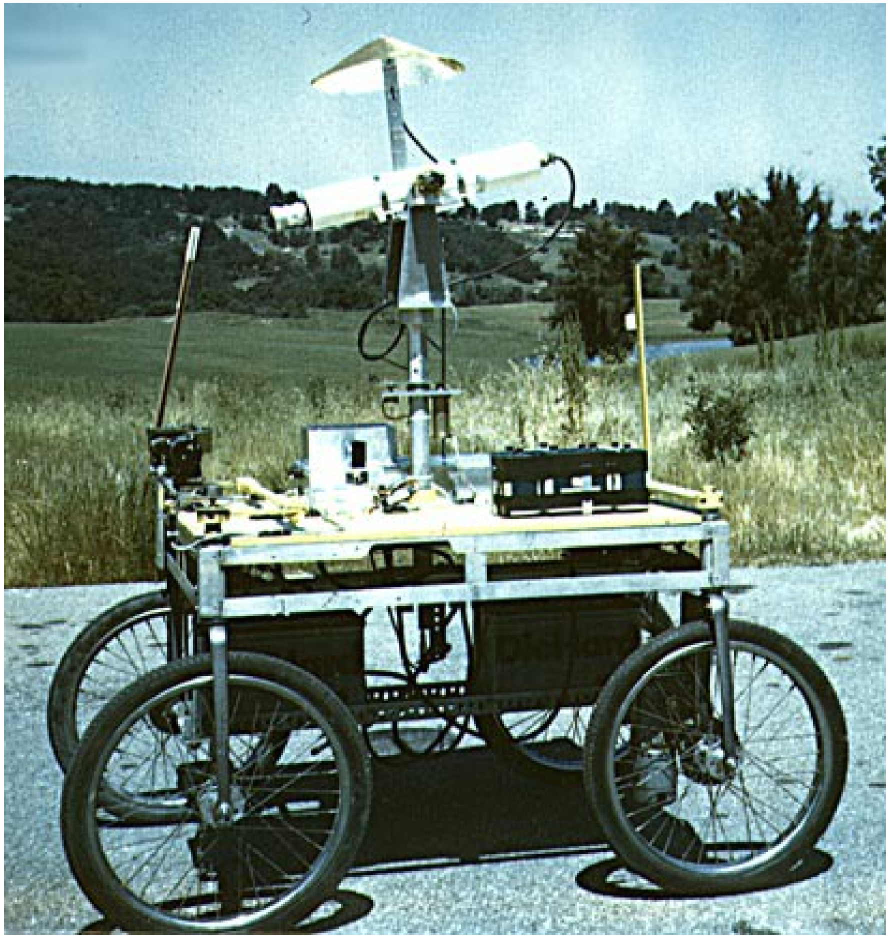

Cartland

Cartland

Behavior-Based Robotics

- Rodney Brooks and students, MIT (1980s)
- Distributed, parallel, “connectionist-style” architecture
- Emergent behaviors

Genghis

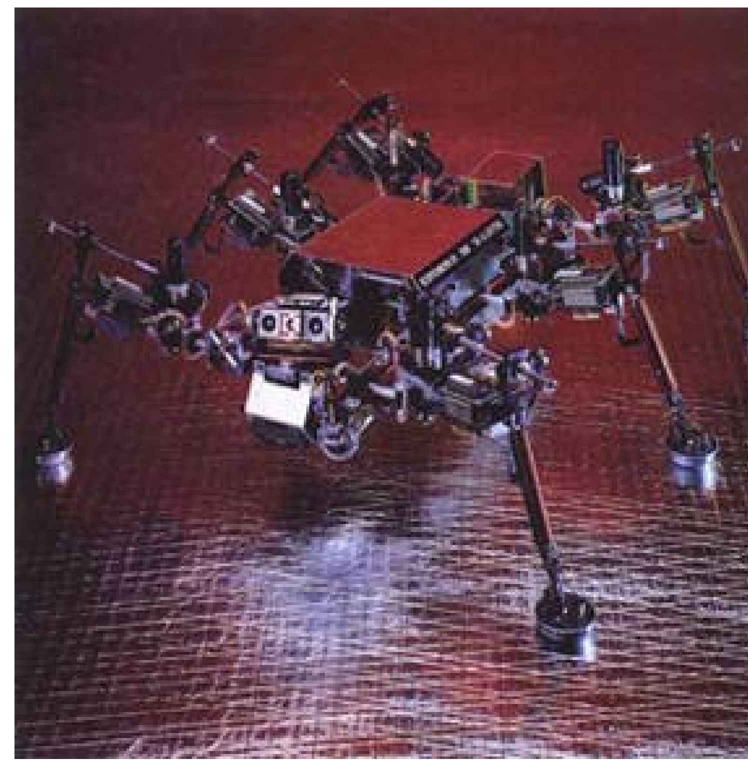

Hannibal

Behavior-Based Robotics

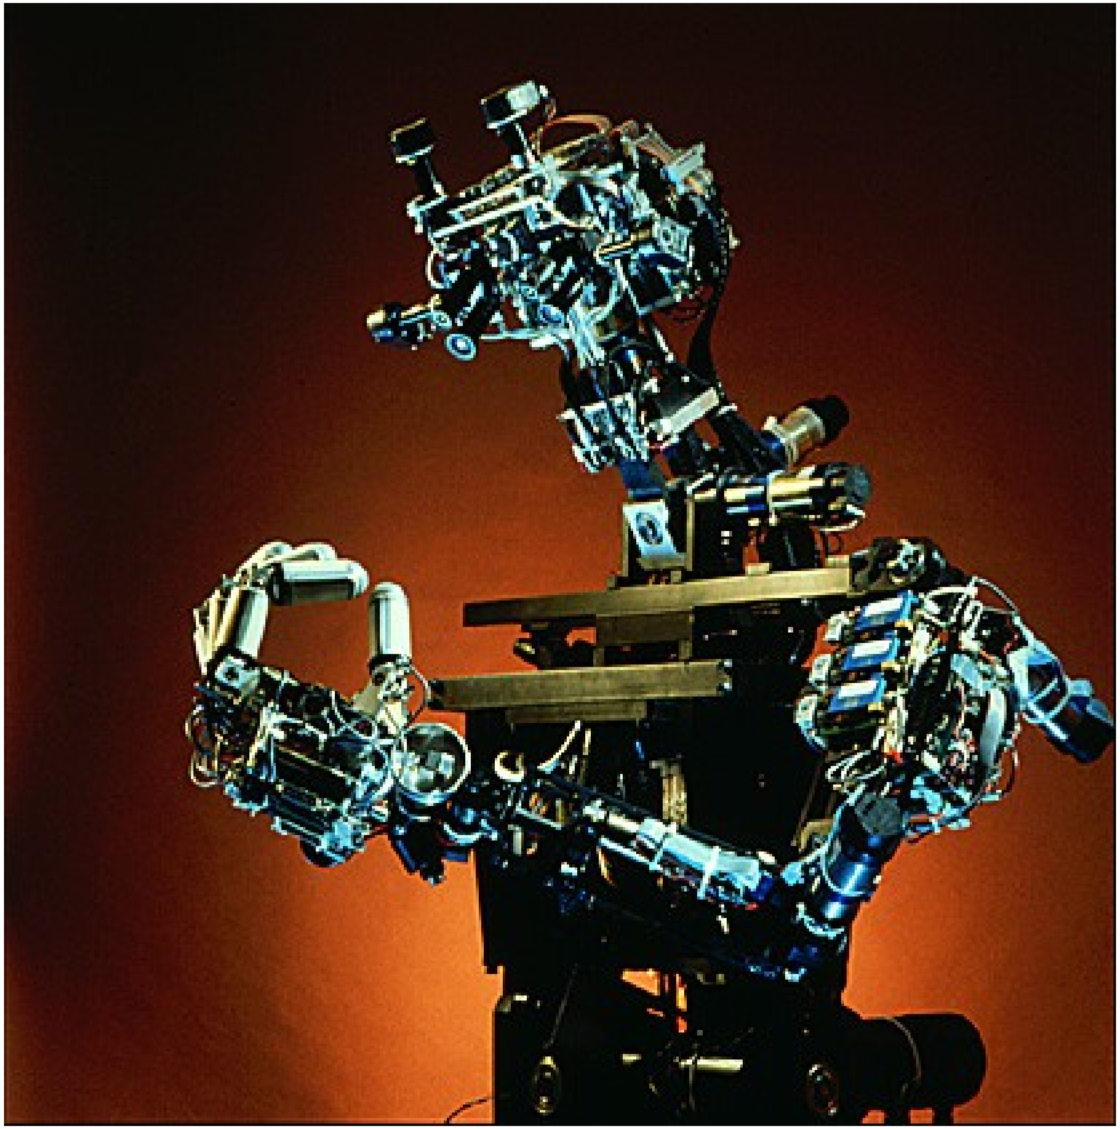

COG

Subsumption Architecture

- Simple **behaviors** connected directly to sensors and motors
- Behaviors organized into **levels** of control
- **Suppression** and **inhibition** of control signals between levels

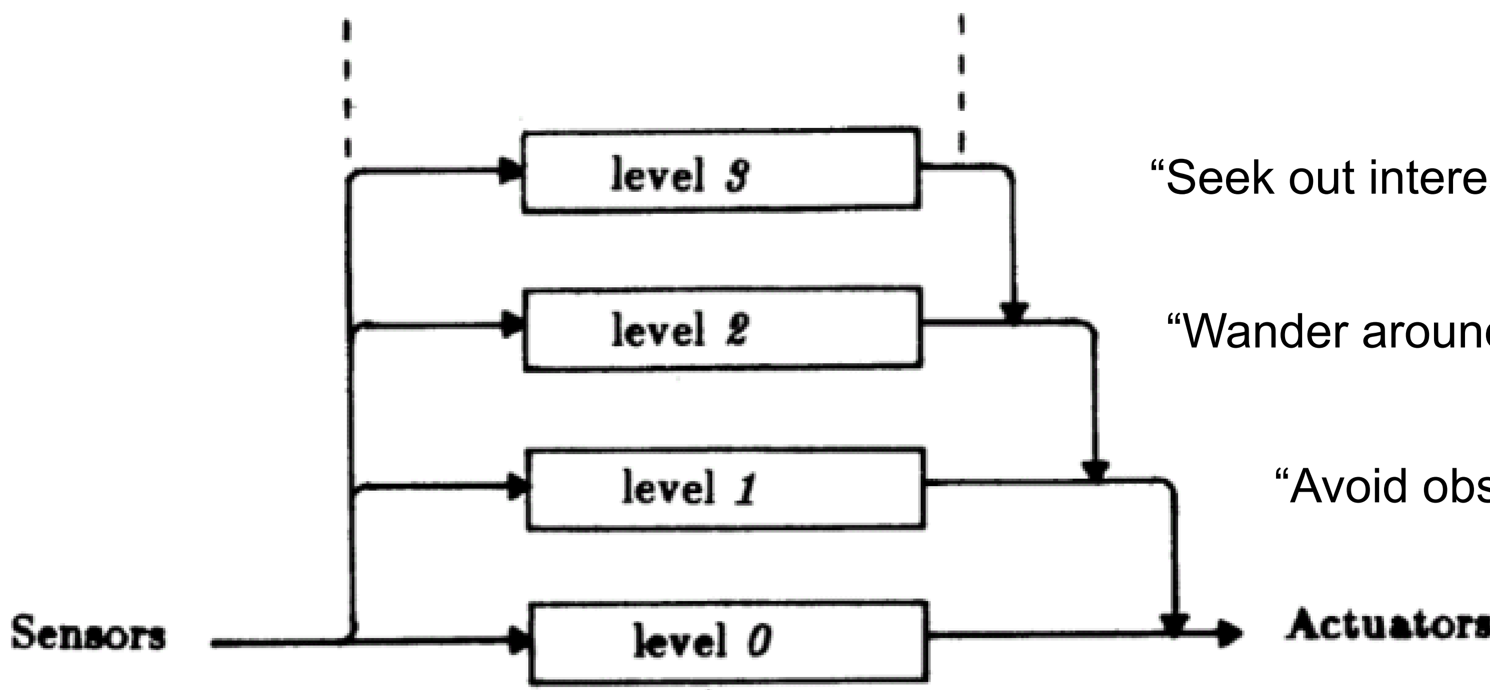

"Seek out interesting places"

"Wander around aimlessly"

"Avoid obstacles"

Subsumption Architecture

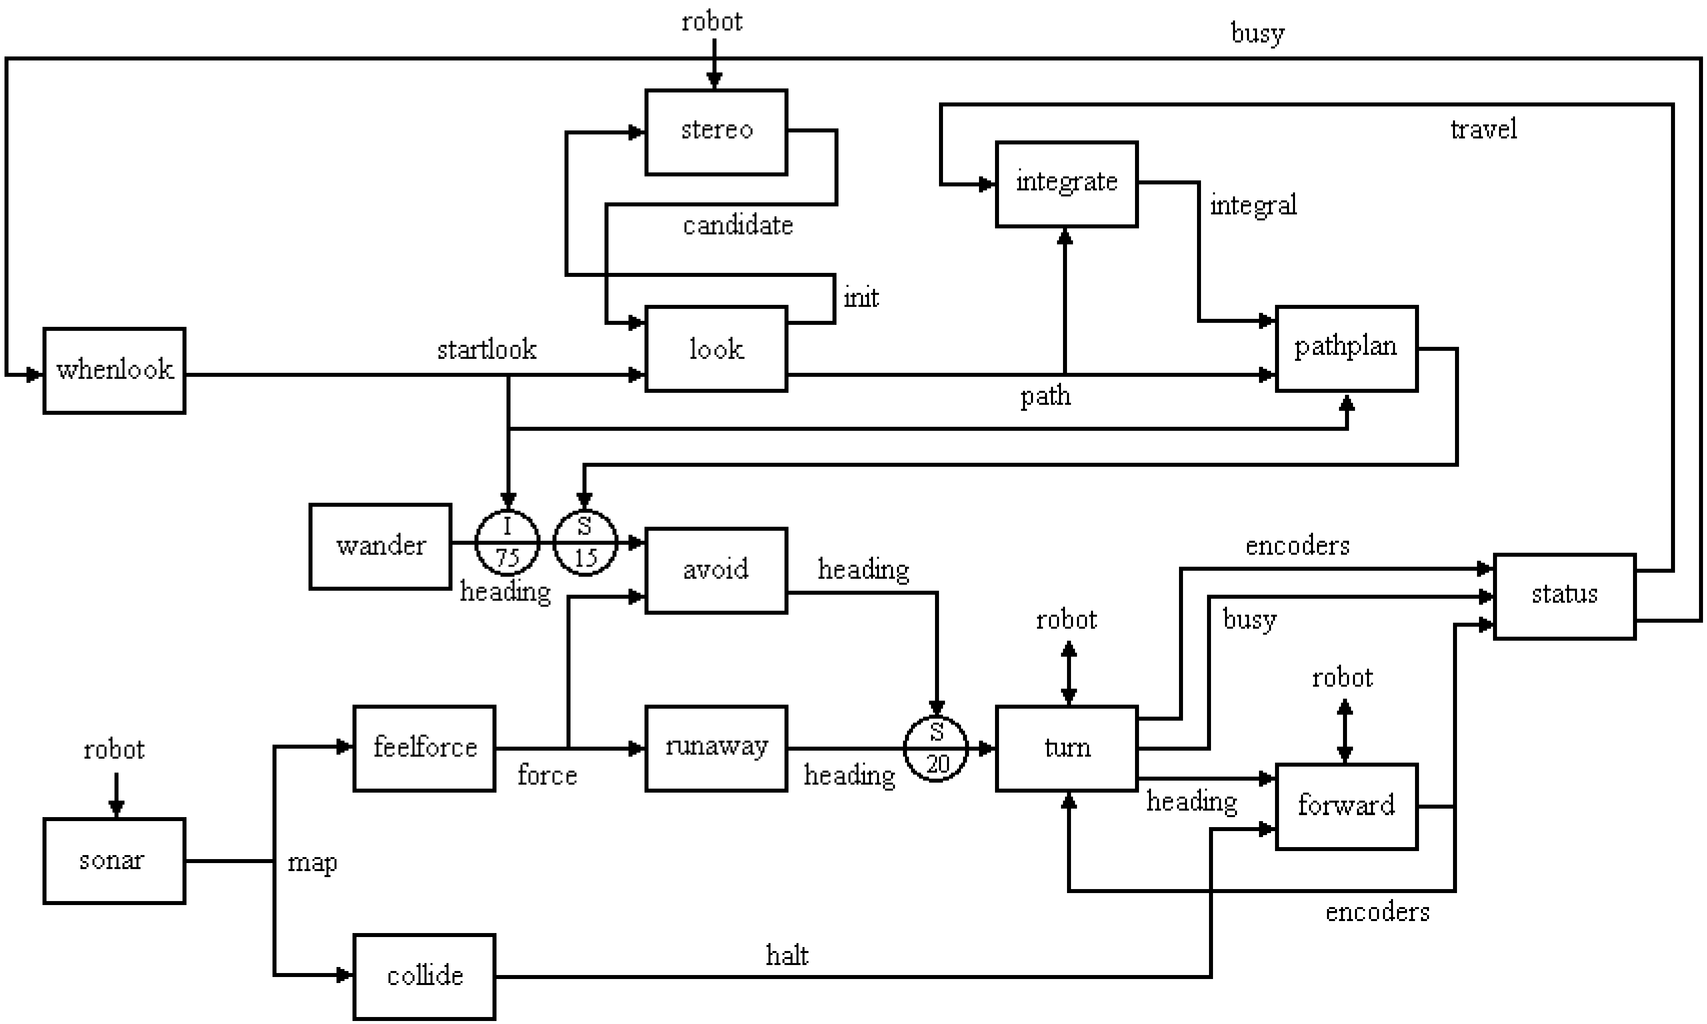

The Deep Learning Era

- **Variations on backpropagation learning** are discovered that significantly improve its performance on multi-layer networks
- **Huge amounts of training data** become widely available for training neural networks, thanks to the Internet
- **GPUs (graphics-processing units)** speed up training enormously
- A **deep convolutional neural network** called AlexNet wins the 2012 ImageNet competition

ImageNet

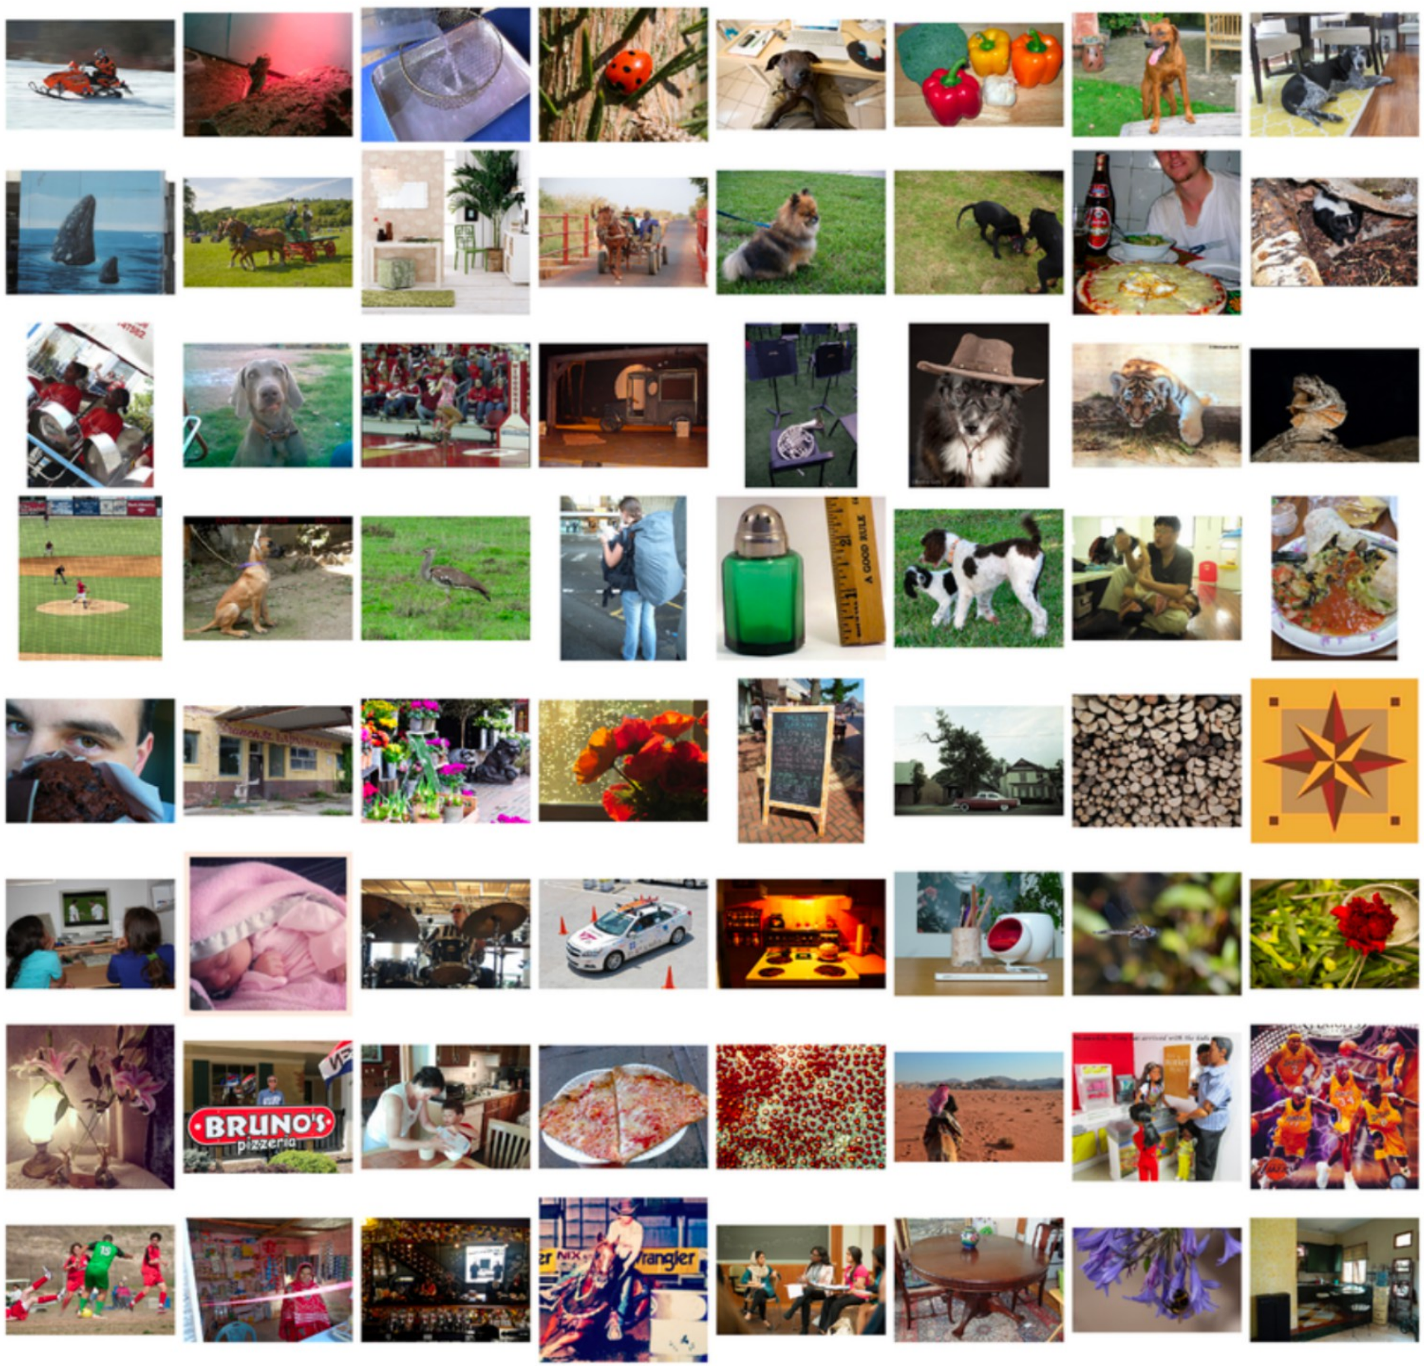

Learning to Recognize Cats

- Google/Stanford **deep neural network** project (2012)
  - 16,000 CPUs and > 1 billion connections
  - 10 million unlabeled 200$\times$200 pixel images from YouTube
  - 3 days of training
  - one particular neuron learned to recognize cats

https://nyti.ms/1EsS3ZD

Deep Reinforcement Learning

- **Google DeepMind** publishes the paper *"Playing Atari with Deep Reinforcement Learning"* in 2013

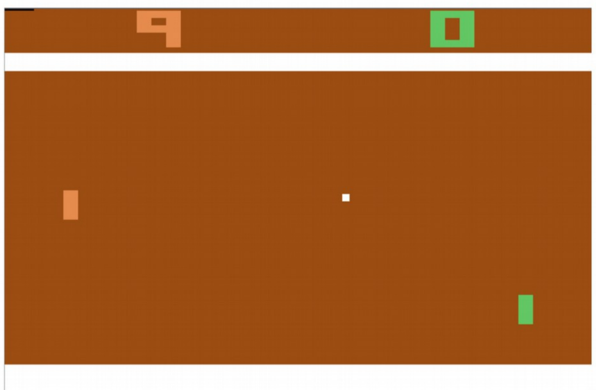

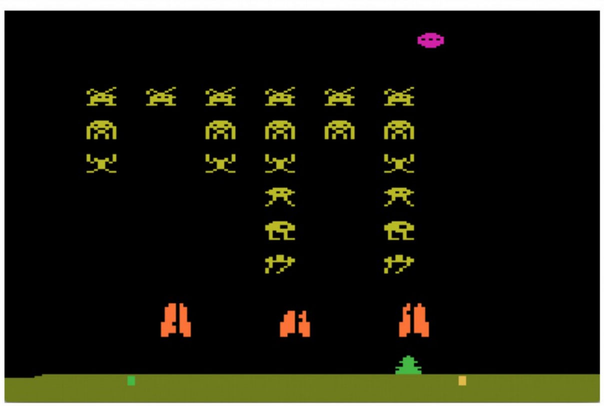


Deep Reinforcement Learning

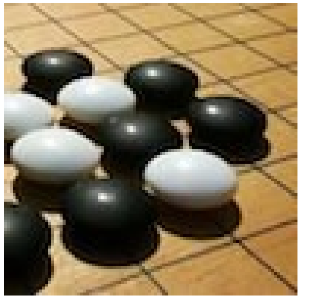

- **AlphaGo** learns to play Go using a deep neural network trained by reinforcement learning on thousands of actual games between human players, as well as games played against itself

- AlphaGo beats professional human Go player **Fan Hui** in 2015, and world champion player **Lee Sedol** in 2016

- **AlphaGo Zero** learns to play Go even better than AlphaGo – starting from random play with no human-based knowledge of the game, and playing millions of games against itself

- **AlphaZero** learns to play Go, Chess, and Shogi from scratch

Deep Reinforcement Learning

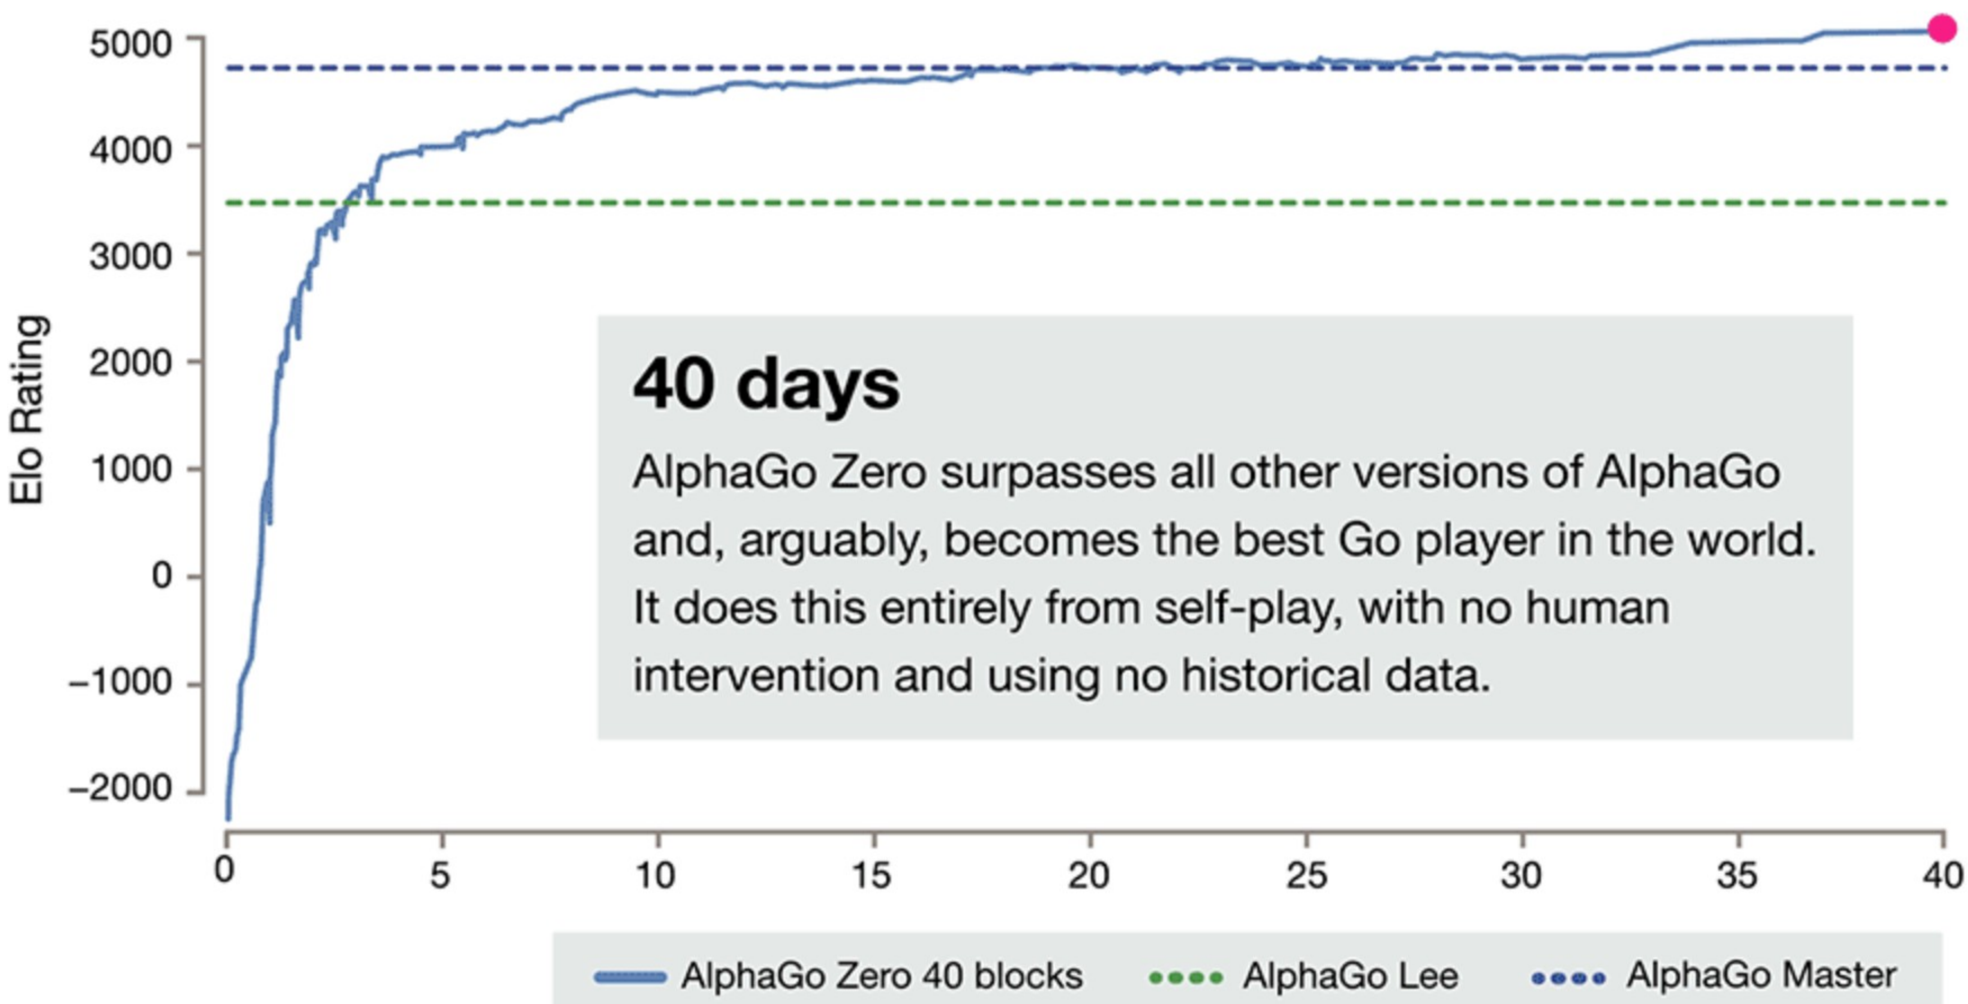

Image Captioning

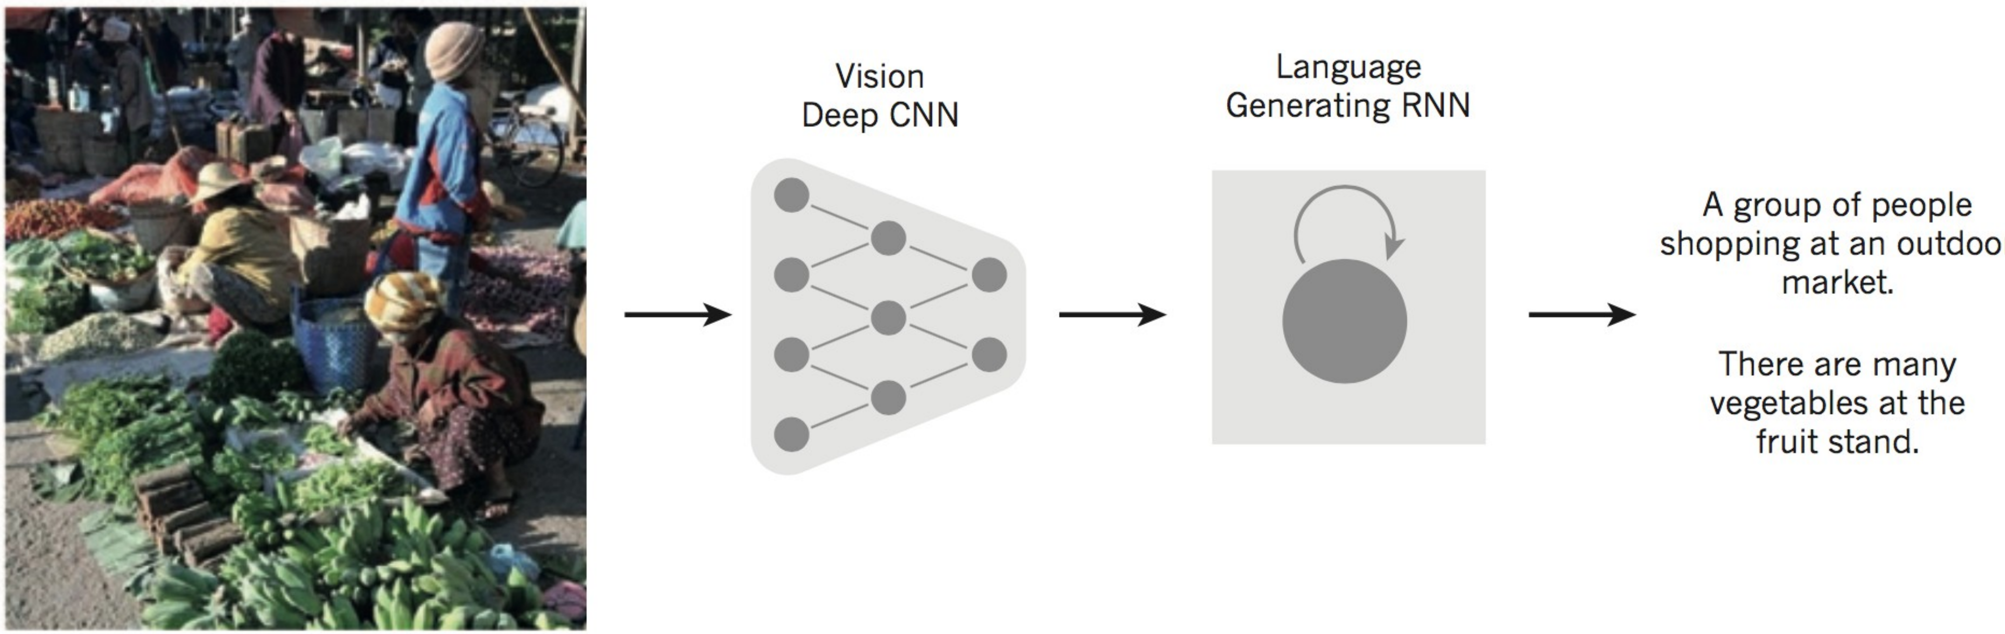

From: Deep learning, Yann LeCun, Yoshua Bengio, Geoffrey Hinton, *Nature* vol. 521, May 28, 2015

Image Captioning

A woman is throwing a **frisbee** in a park.

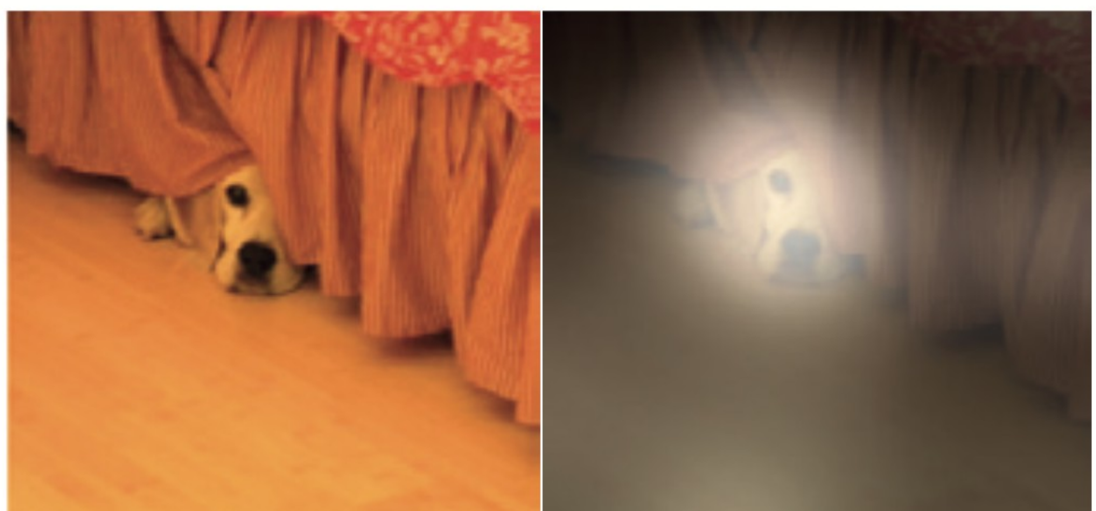

A **dog** is standing on a hardwood floor.

From: Deep learning, Yann LeCun, Yoshua Bengio, Geoffrey Hinton, *Nature* vol. 521, May 28, 2015

Image Captioning

A **stop** sign is on a road with a mountain in the background

A little **girl** sitting on a bed with a teddy bear.

From: Deep learning, Yann LeCun, Yoshua Bengio, Geoffrey Hinton, *Nature* vol. 521, May 28, 2015

Image Captioning

A group of **people** sitting on a boat in the water.

A giraffe standing in a forest with
trees in the background.

From: Deep learning, Yann LeCun, Yoshua Bengio, Geoffrey Hinton, *Nature* vol. 521, May 28, 2015

Generative AI

Large Language Models

- GPT: Generative Pre-trained Transformer (OpenAI)
  GPT-1 (2018) 117 million parameters, 12 layers
  GPT-2 (2019) 1.5 billion parameters, 48 layers
  GPT-3 (2020) 175 billion parameters, 96 layers

- BERT (Google, 2018)

- LaMDA (Google, 2021)

- Codex (OpenAI, 2021)

- ChatGPT (OpenAI, 2022)

- GPT-4 (OpenAI, 2023?)

DALL-E 2

“a teapot in the shape of an avocado”

DALL-E 2

“cats playing chess”

DALL-E 2

“a living room filled with sand, sand on the floor, piano in the room”

The Future

?

?

?