

# Foundations of Artificial Intelligence

## A3. Introduction: AI Past and Present

Malte Helmert

University of Basel

February 28, 2024

# Introduction: Overview

## Chapter overview: introduction

- A1. Organizational Matters
- A2. What is Artificial Intelligence?
- A3. AI Past and Present
- A4. Rational Agents
- A5. Environments and Problem Solving Methods

# A Short History of AI

## Precursors (Until ca. 1943)

1950

1960

1970

1980

1990

2000

...

Philosophy and mathematics ask similar questions that influence AI.

- Aristotle (384–322 BC)
- Leibniz (1646–1716)
- Hilbert program (1920s)



## Gestation (1943–1956)

1950

1960

1970

1980

1990


2000

...

Invention of electrical computers raised question:  
Can computers mimic the human mind?

# Gestation (1943–1956)

## Artificial Neurons



1950

1960

1970

1980

1990

2000


...

W. McCulloch & W. Pitts (1943)

- first computational model of **artificial neuron**
- **network of neurons** can compute any computable function
- basis of **deep learning**

## Gestation (1943–1956)

Artificial  
Neurons



1950

1960

1970

1980

1990

2000

...

Turing Test

Computing Machinery and Intelligence (A. Turing, 1950)

- famous for introducing **Turing test**
- (still) relevant discussion of **AI potential** and **requirements**
- suggests core AI aspects: **knowledge representation, reasoning, language understanding, learning**

# Gestation (1943–1956)

Artificial  
Neurons

Dartmouth

1950

1960

1970

1980

1990

2000

...

Turing Test



John McCarthy



Marvin Minsky



Claude Shannon



Ray Solomonoff



Alan Newell



Herbert Simon



Arthur Samuel



Oliver Selfridge



Nathaniel Rochester




Trenchard More

## Dartmouth workshop (1956)

- ambitious proposal: “An attempt will be made to find how to make machines use language, [...] solve kinds of problems now reserved for humans, and improve themselves.”
- no important breakthrough
- J. McCarthy coins term **artificial intelligence**

# Early Enthusiasm (1952–1969)

Artificial  
Neurons



Dartmouth

1950

1960

1970

1980

1990

2000

...

Turing Test

early enthusiasm (H. Simon, 1957):

“[...] there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until – in the visible future – the range of problems they can handle will be coextensive with the range to which the human mind has been applied.”

# Early Enthusiasm (1952–1969)

Artificial  
Neurons

Dartmouth

1950

1960

1970

1980


1990

2000

...

GPS

Turing Test



General Problem Solver (H. Simon & A. Newell, 1957)

- universal problem solving machine
- imitates human problem solving strategies
- in principle able to solve every formalized symbolic problem
- in practice, GPS solves simple tasks like towers of Hanoi

# Early Enthusiasm (1952–1969)

Artificial  
Neurons

RL for  
Checkers

Dartmouth



1950

1960

1970

1980

1990

2000

...

Turing Test

GPS

Checkers AI (A. Samuel, 1959)



- popularized term **machine learning**
- learned to play at strong amateur level
- uses ideas of **reinforcement learning**

# Early Enthusiasm (1952–1969)

Artificial  
Neurons

RL for  
Checkers

Dartmouth



1950

1960

1970

1980

1990

2000

...

GPS

Microworlds

Turing Test


intelligence in **microworlds**, e.g. **SHRDLU** (T. Winograd, 1968)

- understands natural language
- communicates with user via teletype on **blocks world**
- graphical representation

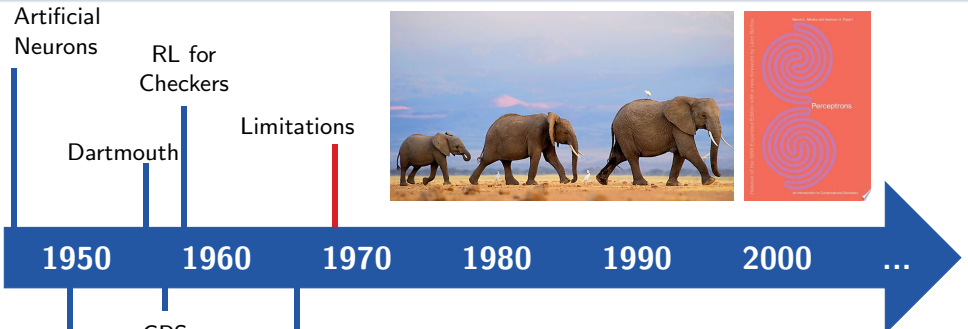
↪ <https://hci.stanford.edu/winograd/shrdlu/>



# Early Enthusiasm (1952–1969)



# A Dose of Reality (1966–1973)



- realization that unlimited computational power is illusion (birth of complexity theory, NP-completeness)
- AI systems (e.g., GPS, systems for micro worlds) *fail to scale*
- fundamental **limitations on basic structures** e.g., XOR problem of perceptrons

# Expert Systems (1969–1986)

Artificial  
Neurons

RL for  
Checkers

Dartmouth

Limitations

DISTRIBUTE-MB-DEVICES-3

```
IF:  the most current active context is distributing massbus devices
&   there is a single port disk drive that has not been assigned to a massbus
&   there are no unassigned dual port disk drives
&   the number of devices that each massbus should support is known
&   there is a massbus that has been assigned at least one disk drive and that should support additional
    disk drives
&   the type of cable needed to connect the disk drive to the previous device on the disk drive is known
THEN: assign the disk drive to the massbus
```

1950

1960

1970

1980

1990

2000

...

GPS

Microworlds

Expert  
Systems

Turing Test

1980s: AI gold rush

- rule-based expert systems commercially successful
- (human) expert knowledge as input
- allows automatic reasoning on larger problems in narrower applications
- also: second heyday of neural networks

# Expert Systems (1969–1986)

Artificial  
Neurons

RL for  
Checkers

Dartmouth

Limitations

DISTRIBUTE-MB-DEVICES-3

```
IF:  the most current active context is distributing massbus devices
&   there is a single port disk drive that has not been assigned to a massbus
&   there are no unassigned dual port disk drives
&   the number of devices that each massbus should support is known
&   there is a massbus that has been assigned at least one disk drive and that should support additional
    disk drives
&   the type of cable needed to connect the disk drive to the previous device on the disk drive is known
THEN: assign the disk drive to the massbus
```

1950

1960

1970

1980

1990

2000

...

Turing Test

GPS

Microworlds

Expert  
Systems

example: R1/XCON (J. McDermott, 1978)

- **input:** desired properties of a VAX computer system according to customer specifications
- **output:** specification of the computer system
- **inference engine:** simple forward chaining of rules

# Expert Systems (1969–1986)

Artificial  
Neurons

RL for  
Checkers

Dartmouth

Limitations

DISTRIBUTE-MB-DEVICES-3

```
IF:  the most current active context is distributing massbus devices
&   there is a single port disk drive that has not been assigned to a massbus
&   there are no unassigned dual port disk drives
&   the number of devices that each massbus should support is known
&   there is a massbus that has been assigned at least one disk drive and that should support additional
    disk drives
&   the type of cable needed to connect the disk drive to the previous device on the disk drive is known
THEN: assign the disk drive to the massbus
```

1950

1960

1970

1980

1990

2000

...

GPS

Microworlds

Expert  
Systems

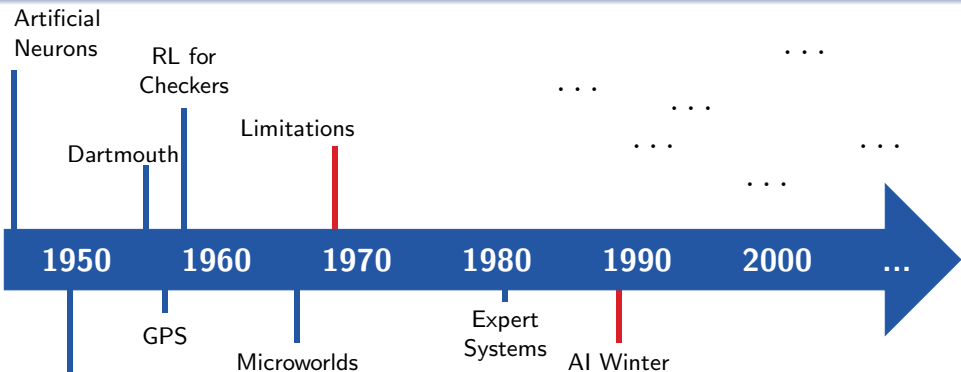
AI Winter

Turing Test

end of 1980s: AI Winter


- companies failed to deliver promises
- expert systems difficult to maintain
- expert systems susceptible to uncertainty

# Coming of Age (1990s and 2000s)



- advent of **probabilistic methods**
- **formalization** of AI techniques
- better understanding of **theoretical complexity**
- increased use of **mathematical methods**
- exploitation of large data sets (**big data**)

# Broad Visibility in Society (Since 2010s)



- well known systems and famous breakthroughs, e.g.,
- broadly used systems (e.g., virtual assistants)
  - AI systems act in real-world (e.g., self-driving cars)
  - systems outperform humans in hard tasks (e.g., AlphaGo)
  - AI and human-written text hard to distinguish (ChatGPT)

# Where are We Today?



# AI Approaching Maturity

## Russell & Norvig (1995)

Gentle revolutions have occurred in robotics, computer vision, machine learning, and knowledge representation.

A better understanding of the problems and their complexity properties, combined with increased mathematical sophistication, has led to workable research agendas and robust methods.

# Where are We Today?



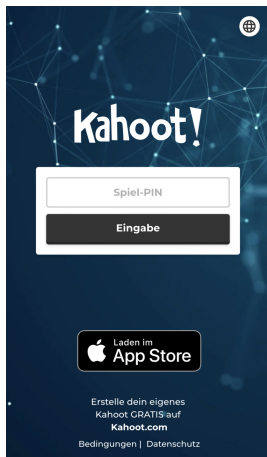
- many coexisting paradigms
  - reactive vs. deliberative
  - data-driven vs. model-driven
  - often hybrid approaches
- many methods, often borrowing from other research areas
  - logic, decision theory, statistics, ...
- different approaches
  - theoretical
  - algorithmic/experimental
  - application-oriented

# Focus on Algorithms and Experiments

Many AI problems are inherently difficult (NP-hard), but strong search techniques and heuristics often solve large problem instances regardless:







- **satisfiability in propositional logic**
  - 10,000 propositional variables or more via **conflict-directed clause learning**
- **constraint solvers**
  - good scalability via **constraint propagation** and automatic exploitation of **problem structure**
- **action planning**
  - $10^{100}$  search states and more by search using **automatically inferred heuristics**

# What Can AI Do Today?



<https://kahoot.it/>

# What Can AI Do Today? – Videos, Articles and AIs



# What Can AI Do Today?

results of our classroom poll:

- ✓ successfully complete an off-road car race
- ✓ beat a world champion table tennis player
- ✓ play guitar in a robot band
- ✓ do and fold the laundry
- ✗ drive safely in downtown Basel
- ✗ win a football match against a human team
- ✓ dance synchronously in a group of robots
- ✓ write code on the level of a CS student
- ✓ beat a world champion Chess, Go or Poker player
- ✓ create inspiring quotes
- ✓ compose music
- ✓ engage in a scientific conversation

# Summary

# Summary

- 1950s/1960s: beginnings of AI; early enthusiasm
- 1970s: micro worlds and knowledge-based systems
- 1980s: gold rush of expert systems followed by “AI winter”
- 1990s/2000s: AI comes of age; research becomes more rigorous and mathematical; mature methods
- 2010s: AI systems enter mainstream