

Foundations of Artificial Intelligence

2. Introduction: AI Past and Present

Malte Helmert

Universität Basel

February 26, 2016

Introduction: Overview

Chapter overview: introduction

- 1. What is Artificial Intelligence?
- 2. AI Past and Present
- 3. Rational Agents
- 4. Environments and Problem Solving Methods

A Short History of AI

The Origins of AI

Before AI, **philosophy**, **mathematics**, **psychology** and **linguistics** asked similar questions and influenced AI.

Gestation of AI (~1943–1956)

With the advent of electrical computers, many asked:

Can computers mimic the human mind?

↪ Turing test

60 Years of AI: 1950s

Dartmouth workshop (1956):

John McCarthy coins the term **artificial intelligence**

↪ “official birth year” of the research area

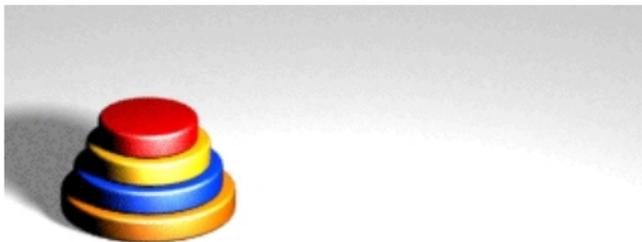
early enthusiasm:

Herbert Simon (1957)

It is not my aim to surprise or shock you – but the simplest way I can summarize is to say that 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: General Problem Solver (GPS)

- **GPS**: developed in 1957 by **Herbert Simon** and **Allen Newell**
- **goal**: build a universal problem solving machine by imitating human problem solving strategies
- ↪ first representative of “thinking humanly” approach to AI
- every formalized symbolic problem solvable **in principle** by GPS
- **practice**: GPS solves simple tasks like towers of Hanoi, but does not scale to **realistic application problems**



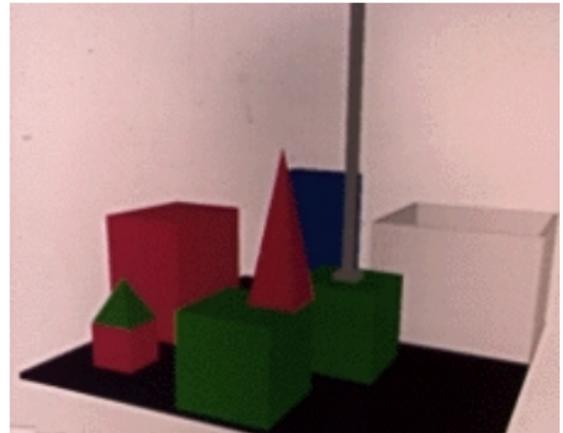
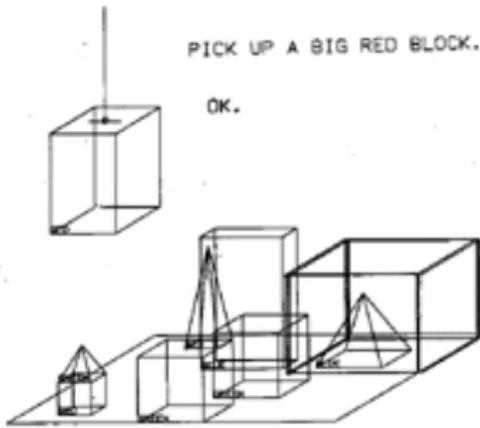
60 Years of AI: 1960s and 1970s

- 1960s: “intelligent behavior” in many demonstration systems for **micro worlds** (e.g., **blocks world**)
- 1970s: systems for micro worlds fail to scale
- ↪ “intelligent behavior” requires **knowledge**
- ↪ **knowledge-based systems**

German: wissensbasierte Systeme, Mikrowelten

Micro Worlds: Blocks World (SHRDLU)

- **SHRDLU** understands natural language
- developed 1968–1970 by Terry Winograd at MIT AI Lab
- user dialog (via teletype) on **blocks world**
- graphical representation on screen



↪ <http://hci.stanford.edu/~winograd/shrdlu/>

A SHRDLU Dialog

Human: PICK UP A BIG RED BLOCK.

Computer: OK. [does it]

Human: GRASP THE PYRAMID.

Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.

Human: FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE
HOLDING AND PUT IT INTO THE BOX.

Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS
TALLER THAN THE ONE I AM HOLDING.

Computer: OK. [does it]

Human: WHAT DOES THE BOX CONTAIN?

Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.

Human: WHAT IS THE PYRAMID SUPPORTED BY?

Computer: THE BOX.

Human: HOW MANY BLOCKS ARE NOT IN THE BOX?

Computer: FOUR OF THEM.

Human: IS AT LEAST ONE OF THEM NARROWER THAN THE ONE
WHICH I TOLD YOU TO PICK UP?

Computer: YES, THE RED CUBE.

...

60 Years of AI: 1980s

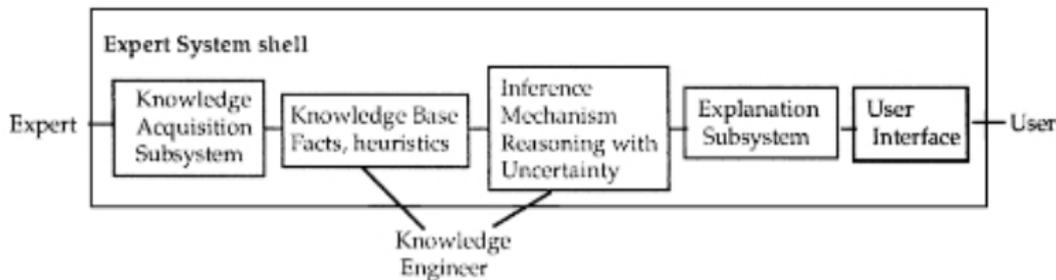
1980s: AI gold rush

- commercial success of rule-based **expert systems** (e.g. R1)
- research boom, e.g., **Fifth Generation Computer Systems** project in Japan
- (second) heyday of **neural networks**
- **end of 1980s**: expert systems and other AI approaches less successful than hoped

↪ **AI winter**

German: Expertensysteme, neuronale Netze, KI-Winter

Expert Systems



- **knowledge acquisition:** support for modeling expert knowledge
- **knowledge base:** database of facts and heuristic rules
- **inference engine:** reasoning mechanism; processes information in knowledge base to solve problems
- **explanation system:** explains conclusions of inference engine

German: Wissensakquise, Wissensbasis, Inferenz-Engine, Erklärungssystem

Expert Systems: R1/XCON

- **application area:** configuration of VAX computers according to customer specifications
- developed by John McDermott et al. (1978–1981)
- **input:** desired properties of the computer system
- **output:** specification of the computer system
- **inference engine:** simple forward chaining of rules

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

Commercial Success of R1/XCON

- DEC failed in development of “conventional” program
- asked McDermott to build an AI system ~→ R1/XCON
- until 1986: processed 80,000 jobs; success rate 95–98%
- estimated savings for DEC: 25,000,000 US\$/year
- **but:** suffered disadvantages of simple rule-based systems
 - rule-base quickly grew very large
 - ~→ became unreliable, hard to understand and maintain

60 Years of AI: 1990s

1990s: AI comes of age

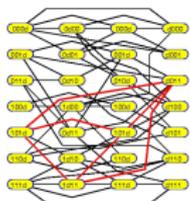
- advent of probabilistic methods
- agent-oriented approaches
- formalization of AI techniques
- better understanding of theoretical complexity
- increased use of mathematical methods

Russell & Norvig (1995)

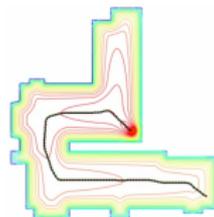
Gentle revolutions have occurred in robotics, computer vision, machine learning, and knowledge representation. A better understanding of the problems and their complexity properties, combined with increased mathematical sophistication, has led to workable research agendas and robust methods.

60 Years of AI: Today

3-SAT TRANSFORMS to CLIQUE

 $(x_1 + \bar{x}_2 + x_3) (x_1 + x_3 + \bar{x}_4) (x_1 + \bar{x}_2 + x_4) (x_2 + \bar{x}_3 + x_4)$ 

	A	B	C	D	E	F	G	H
1	Q							
2			Q					
3					Q			
4		Q						
5				Q				
6	1	3,4	2,5	4,5	3,5	1	2	3
7								
8								



- many coexisting paradigms
 - reactive vs. deliberative
 - probabilistic vs. symbolic
 - often hybrid approaches
- many methods, often borrowing from other research areas
 - logic, decision theory, statistics, ...
- many approaches
 - theoretical, algorithmic/experimental, system-oriented, ...
- many success stories no longer count as “pure” AI
 - board games, logic programming, search methods, ...

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**

AI Systems Past and Present

Example System (1996): Chess

Deep Blue vs. Kasparov (1996):

- first win of a chess computer against reigning world champion under tournament conditions (time controls)

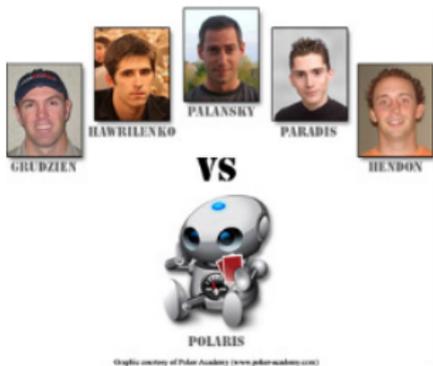


Deep Blue in Futurama:

<http://www.cc.com/video-clips/p1jk76/futurama-action-rangers>

Example System (2015): Poker

Cepheus (2015), developed at the University of Alberta,
“solves” heads-up limit Hold'em



Polaris vs. world-class players (2008)



Univ. of Alberta computer poker research group

<http://poker.cs.ualberta.ca/>

Example System (1998): Driving Cars

ALVINN (1998), developed by Dean Pomerleau et al., CMU, keeps the lane for more than 4000 km

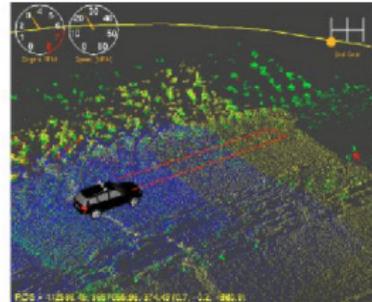
- semi-autonomic driving
- camera images, artificial neural networks



Example System (2005): Driving Cars

Stanley (2005) by Sebastian Thrun et al., Stanford University, wins DARPA Grand Challenge (2'000'000 US\$ prize)

- drives autonomously through Mojave desert (212 km, off-road)
- winning time: less than 7 hours



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: AI comes of age; research becomes more rigorous and mathematical; mature methods
- 2000s/2010s: AI systems enter mainstream