Title: Historical notes

AIMA: Chapter 1

Introduction to Artificial Intelligence

CSCE 476-876, Fall 2022

URL: `www.cse.unl.edu/~cse476`

URL: `cse-linux-01.unl.edu/~c-bchoueiry2/CSCE476_1238`

Berthe Y. Choueiry (Shu-we-ri)

Summary

- Goal of AI:
  - Understand intelligent entities
  - Build intelligent entities
- Multidisciplinary:
  - philosophy, psychology, cognitive science
  - mathematics, engineering, computer science
- Study:
  - general mechanisms
  - specific tasks
- Dimensions:
  - Concern: thought+mental processes vs. action+behavior
  - Performance: human-like vs. rational

What is AI?

- Acting humanly:
  Turing test: NLP, KR, Automated Reasoning, ML
  Total Turing test: computer vision, robotics
- Thinking humanly:
  Build a theory of the mind
  Cognitive science = computer models + psychological experiments
- Thinking rationally: Logic, probability
- Acting rationally: Rational agent
  Does the 'right' think
  Limited/bounded rationality

Beneficial machines: value alignment problem: machine's and humans objectives

$\longrightarrow$ Machines provably beneficial to humans

Disciplines

- Philosophy: from the physical brain to the mind, source of knowledge, how knowledge lead to action
- Mathematics: inference, computability, reasoning under uncertainty
- Economics: handling preferences, collaborating/competing, decision/game theory, multiagents, optimization (OR)
- Neuroscience
- Psychology: cognitive science
- Computer Engineering
- Control theory, cybernetics
- Linguistics: NLP/computational linguistics

Main milestones

- Minsky, McCarthy: Foundations, Knowledge representation and reasoning
- Feigenbaum, Reddy: Expert systems to solve real-world problems
- Pearl: Reasoning about uncertainty
- Bengio, Hinton, LeCun: Deep Learning

Quick historical note (I)

1956: McCarthy organizes a two-month workshop in Dartmouth
no breakthrough, united major players, term coined

1952-1969: Early enthusiasm and great expectations
General Problem Solver (Newell & Simon),
Chess program (learning disproved "computers do what they are told to do"),
LISP in 1958, time sharing, principles of knowledge representation and reasoning,
Split: logic (neat) vs. anti-logic (scruffy, clumsy),
Microworlds (*e.g.*, block world: $\rightarrow$ vision, constraint propagation, NL understanding, planner),
Neural Nets, etc.

Quick historical note (II)

1966-1973: hard reality. Too big claims, wildly optimistic

1- Systems work on 1 or 2 examples, failed otherwise

NLP: Russian → English.

The spirit is willing but the flesh is weak

→ the vodka is good but the meat is rotten

2- Intractability: difficulty to scale up, handle combinatorial explosion

UK → report Lighthill in 1973, etc.

Neural Net almost disappears

1969-1979: Knowledge-based systems

Knowledge-based system (DENDRAL) expert knowledge

expert systems (MYCIN), certainty factors, frames (OO!)

Quick historical note (III)

1980-present: AI becomes an industry

R1 at DEC, Fifth Generation project.

S/W: Carnegie Group, Inference, Intellicorp, Teknowledge

H/W: Lisp Machines, TI, Symbolics, Xerox.

1986-present: The return of neural networks.

Late 80's-early 90's: AI winter

1987-present: Big changes. AI becomes a science

Claims more rigorously supported: empirically or theoretically

Resurgence of probabilistic & decision theory (UAI).

Machine Learning

"Nouvelle AI:" ALife, GAs, soft computing

1995-present: Emergence of intelligent agents

2001-present: Big Data: Image recognition, IBM's Watson (2011)

2011-present: Deep Learning (CNN)

State of the art: huge progress

- Academic/research: booming
- Industrial/commercial: Robotics, legged locomotion, planning and scheduling, machine translation, speech recognition, recommendation systems, image understanding, medicine, climate science, etc.

Risks and benefits of AI, ethics debates

- Lethal autonomous weapons
- Surveilland and persuasion
- Biased decision making
- Impact on employment
- Safety-critical applications
- Cybersecurity

$\longrightarrow$ In this class, we focus on the basics: representation and reasoning