

CSC 550: Introduction to Artificial Intelligence

Fall 2008

See online syllabus at:

<http://dave-reed.com/csc550>

Course goals:

- survey the field of Artificial Intelligence, including major areas of study and research
- study the foundational concepts and theories that underlie AI, including search, knowledge representation, and sub-symbolic models
- contrast the main approaches to AI: symbolic vs. emergent
- provide practical experience developing AI systems using Scheme

1

What is the field of Artificial Intelligence?

General definition: AI is the branch of computer science that is concerned with the automation of intelligent behavior.

- what is intelligent behavior?
- is intelligent behavior the same for a computer and a human?
e.g., Weizenbaum's ELIZA program

Tighter definition: AI is the science of making machines do things that would require intelligence if done by people. (Minsky)

- at least we have experience with human intelligence

possible definition: intelligence is the ability to form plans to achieve goals by interacting with an information-rich environment

intelligence encompasses abilities such as:

understanding language, reasoning, perception, learning, ...

2

What is AI? (cont.)

Self-defeating definition: AI is the science of automating intelligent behaviors currently achievable by humans only.

- this is a common perception by the general public
- as each problem is solved, the mystery goes away and it's no longer "AI"
successes go away, leaving only unsolved problems

Self-fulfilling definition: AI is the collection of problems and methodologies studied by AI researchers.

- AI ranges across many disciplines
computer science, engineering, cognitive science, logic, ...
- research often defies classification, requires a broad context

3

Pre-history of AI

the quest for understanding & automating intelligence has deep roots

4th cent. B.C.: Aristotle studied mind & thought, defined formal logic

14th-16th cent.: Renaissance thought built on the idea that all natural or artificial processes could be analyzed and understood

19th cent.: advances in science made the idea of artificial life seem plausible

- Shelley's *Frankenstein* raised moral and ethical questions
- Babbage's Analytical Engine proposed programmable machine -- metaphor for brain

19th-20th cent.: advances in logic formalisms, e.g., Boolean algebra, predicate calculus

20th cent.: advent of digital computers in late 1940's made AI viable

- Turing wrote seminal paper on thinking machines (1950)

birth of AI occurred when Marvin Minsky & John McCarthy organized the Dartmouth Conference in 1956

- brought together researchers interested in "intelligent machines"
- for next 20 years, virtually all advances in AI were by attendees
Minsky (MIT), McCarthy (MIT/Stanford), Newell & Simon (Carnegie),...

4

History of AI

the history of AI research is a continual cycle of

optimism & hype → reality check & backlash → refocus & progress → ...

1950's – birth of AI, optimism on many fronts

general purpose reasoning, machine translation, neural computing, ...

- first neural net simulator (Minsky): could learn to traverse a maze
- GPS (Newell & Simon): general problem-solver/planner, means-end analysis
- Geometry Theorem Prover (Gelertner): input diagrams, backward reasoning
- SAINT(Slagle): symbolic integration, could pass MIT calculus exam

1960's – failed to meet claims of 50's, problems turned out to be hard!

so, backed up and focused on "micro-worlds"

within limited domains, success in: reasoning, perception, understanding, ...

- ANALOGY (Evans & Minsky): could solve IQ test puzzle
- STUDENT (Bobrow & Minsky): could solve algebraic word problems
- SHRDLU (Winograd): could manipulate blocks using robotic arm, explain self
- Minsky & Papert demonstrated the limitations of neural nets

5

History of AI (cont.)

1970's – results from micro-worlds did not easily scale up

so, backed up and focused on theoretical foundations, learning/understanding

- conceptual dependency theory (Schunk)
- frames (Minsky)
- machine learning: ID3 (Quinlan), AM (Lenat)

practical success: expert systems

- DENDRAL (Feigenbaum): identified molecular structure
- MYCIN (Shortliffe & Buchanan): diagnosed infectious blood diseases

1980's – BOOM TOWN!

cheaper computing made AI software feasible

success with expert systems, neural nets revisited, 5th Generation Project

- XCON (McDermott): saved DEC ~ \$40M per year
- neural computing: back-propagation (Werbos), associative memory (Hopfield)
- logic programming, specialized AI technology seen as future

1990's – again, failed to meet high expectations

so, backed up and focused : embedded intelligent systems, agents, ...

hybrid approaches: logic + neural nets + genetic algorithms + fuzzy + ...

- CYC (Lenat): far-reaching project to capture common-sense reasoning
- Society of Mind (Minsky): intelligence is product of complex interactions of simple agents
- Deep Blue (formerly Deep Thought): defeated Kasparov in Speed Chess in 1997

6

Philosophical extremes in AI

Neats vs. Scruffies

- Neats focus on smaller, simplified problems that can be well-understood, then attempt to generalize lessons learned
- Scruffies tackle big, hard problems directly using less formal approaches

GOFAs vs. Emergents

- GOFAI (Good Old-Fashioned AI) works on the assumption that intelligence can and should be modeled at the symbolic level
- Emergents believe intelligence emerges out of the complex interaction of simple, sub-symbolic processes

Weak AI vs. Strong AI

- Weak AI believes that machine intelligence need only mimic the behavior of human intelligence
- Strong AI demands that machine intelligence must mimic the internal processes of human intelligence, not just the external behavior

7

Criteria for success

long term: Turing Test (for Weak AI)

- as proposed by Alan Turing (1950), if a computer can make people think it is human (i.e., intelligent) via an unrestricted conversation, then it is intelligent
- Turing predicted fully intelligent machines by 2000, *not even close*
- Loebner Prize competition, extremely controversial

short term: more modest success in limited domains

- performance equal or better than humans
e.g., game playing (Deep Blue), expert systems (MYCIN)
- real-world practicality \$\$\$
e.g., expert systems (XCON, Prospector), fuzzy logic (cruise control)

8

Criteria for success (cont.)

AI is still a long way from its long term goal

but in ~50 years, it has matured into a legitimate branch of science

- has realized its problems are hard
- has factored its problems into subfields
- is attacking simple problems first, but thinking big

surprisingly, AI has done better at "expert tasks" as opposed to "mundane tasks" that require common sense & experience

- hard for humans, not for AI:
e.g., chess, rule-based reasoning & diagnosis,
- easy for humans, not for AI:
e.g., language understanding, vision, mobility, ...

9

Course outline

1. Problem-solving as search
 - review of Scheme programming
 - state spaces, search strategies, heuristics
 - game playing
2. Knowledge representation & reasoning
 - representation structures (semantic nets, frames, scripts, ...)
 - expert systems, uncertainty
3. Machine learning: decision trees
 - decision tree structure, IDE learning algorithm
4. Machine learning: neural networks
 - connectionist models, backprop, associative memory
5. Machine learning: genetic algorithms
 - emergent models: evolutionary paradigm, artificial life
6. Selected AI topics
 - student presentations

10