Introduction – Artificial Intelligence a Modern Approach

Russell and Norvig: 1

"The exciting new effort to make computers think . . . *machines with minds*, in the full and literal sense" (Haugeland, 1985) "[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . ." (Bellman, 1978)	"The study of mental faculties through the use of computational models" (Charniak and McDermott, 1985) "The study of the computations that make it possible to perceive, reason, and act" (Winston, 1992)
"The art of creating machines that perform functions that require intelligence when performed by people" (Kurzweil, 1990) "The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)	"A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes" (Schalkoff, 1990) "The branch of computer science that is concerned with the automation of intelligent behavior" (Luger and Stubblefield, 1993)

Figure 1.1 Some definitions of AI. They are organized into four categories:

Systems that think like humans.	Systems that think rationally.
Systems that act like humans.	Systems that act rationally.

What is AI?

Views of AI fall into four categories:

Thought Processes Like Humans	Rational Thought Processes
Act Like Humans	Act Rationally

The textbook advocates "acting rationally"

Thinking humanly: cognitive modeling

- Cognitive Science – must figure out how human's think
  - [introspection – experimental investigation]
  - Requires scientific theories of internal activities of the brain
  - Express these theories as computer programs
- How to validate? Requires
  1. Predicting and testing behavior of human subjects (top-down)
  2. Direct identification from neurological data (bottom-up)

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- Operational test for intelligent behavior: the Imitation Game

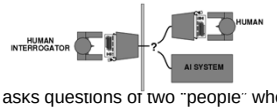

- Interrogator asks questions of two "people" who are out of sight
- 30 minutes to ask whatever he or she wants
- Task: to determine only through the questions and answers which is which
- Computer deemed intelligent if interrogator can't distinguish between person and computer.

Artificial intelligence is the enterprise of constructing an artifact that can pass the Turing text

Acting humanly: Turing Test (cont)

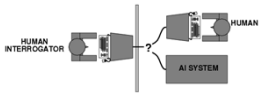

- What major components were important
  - Natural language processing
  - Knowledge representation
  - Automated reasoning
  - Machine learning
- What additional for total Turing Test
  - Computer vision
  - Robotics
- Note: looking at I/O behavior only

Thinking rationally: "laws of thought"

- Aristotle: what are correct arguments/thought processes?
- Several Greek schools developed various forms of *logic*: *notation* and *rules of derivation* for thoughts; may or may not have proceeded to the idea of mechanization
- Direct line through mathematics and philosophy to modern AI
- Problems:
  1. Not all intelligent behavior is mediated by logical deliberation
  2. Some knowledge is very hard to encode – informal, uncertain
  3. In practice, computationally intractable

Acting rationally: rational agent

- Correct thinking is good but:
  - Sometimes you must do something and there is no provably correct thing to do
  - Sometimes you must react quicker without time for reasoning
- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g.,

Acting rationally: rational agent (cont)

- Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information
- Doesn't necessarily involve thinking – e.g., blinking reflex – but thinking should be in the service of rational action
- This is the view taken by the book

Rational agents

- An agent is an entity that perceives and acts
- This course is about designing rational agents
- Abstractly, an agent is a function from percept histories to actions:

  $[f: \mathcal{P}^* \rightarrow \mathcal{A}]$

  For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance
- Caveat: computational limitations make perfect rationality unachievable
  → design best program for given machine resources

AI prehistory

• Philosophy	Logic, methods of reasoning, mind as physical system foundations of learning, language, rationality
• Mathematics	Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability
• Economics	utility, decision theory
• Neuroscience	physical substrate for mental activity
• Psychology	phenomena of perception and motor control, experimental techniques
• Computer engineering	building fast computers
• Control theory	design systems that maximize an objective function over time
• Linguistics	knowledge representation, grammar

Abridged history of AI

• 1943	McCulloch & Pitts: Boolean circuit model of brain
• 1950	Turing's "Computing Machinery and Intelligence"
• 1956	Dartmouth meeting: "Artificial Intelligence" adopted
• 1952—69	Look, Ma, no hands!
• 1950s	Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
• 1965	Robinson's complete algorithm for logical reasoning
• 1966—73	AI discovers computational complexity Neural network research almost disappears
• 1969—79	Early development of knowledge-based systems
• 1980--	AI becomes an industry
• 1986--	Neural networks return to popularity
• 1987--	AI becomes a science
• 1995--	The emergence of intelligent agents

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Proved a mathematical conjecture (Robbins conjecture) unsolved for decades
- No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego)
- During the 1991 Gulf War, US forces deployed an AI logistics planning and scheduling program that involved up to 50,000 vehicles, cargo, and people
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft
- `Proverb` solves crossword puzzles better than most humans