For class use only, do not distribute

COMP150
Behavior-Based Robotics

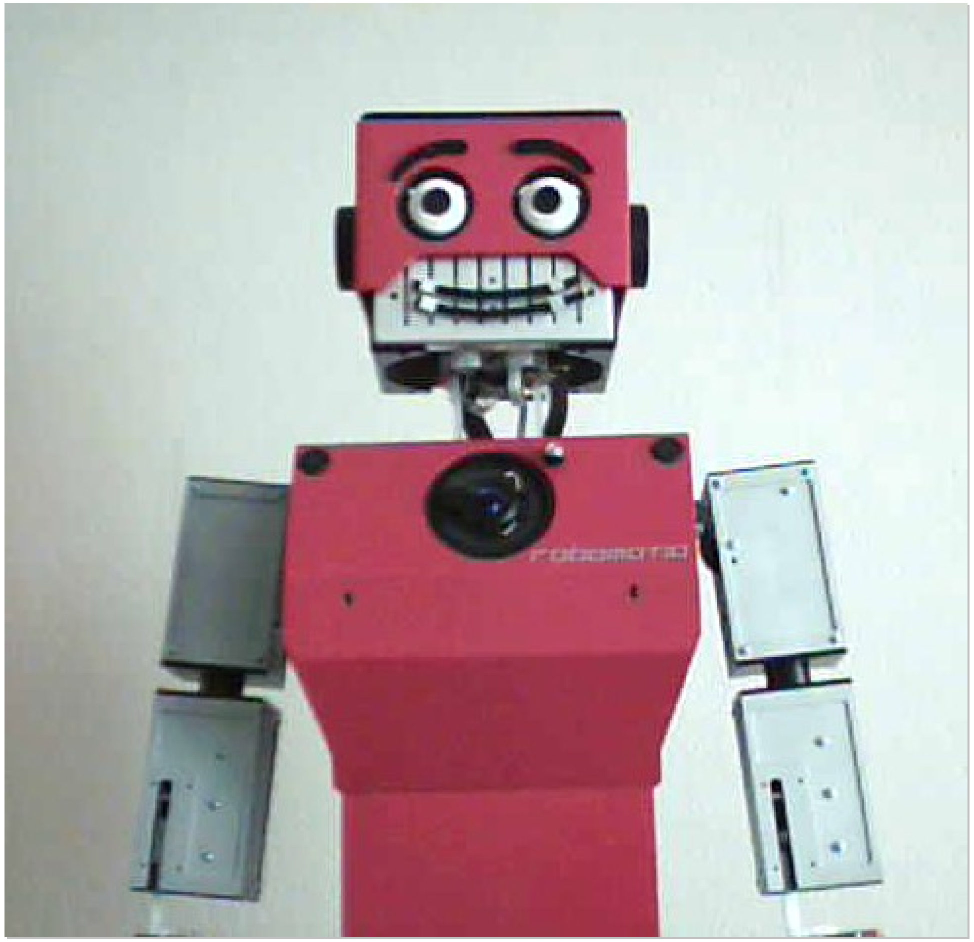

http://www.cs.tufts.edu/comp/150BBR/timetable.html
http://www.cs.tufts.edu/comp/150BBR/syllabus.html

Plan for this week

- Last progress report due on Thursday (please submit on time)
- Rating of the grad student class presentations (use “N/A” if you did not attend a particular presentation)
- We will first talk about artificial intelligence research (original goals, status quo, etc.) with some examples “game playing” and logical reasoning
- Then we will revisit the state-of-the-art in robotics and talk about ethical questions (resulting both from AI and robotics research)
- Will also briefly talk about the final project presentations and the final papers (on Thursday)

The Turing Test

- Talked already about what “intelligence” means and the problems with defining it
- Turing (1950) “Computing machinery and intelligence”: “Can machines think?” “Can machines behave intelligently?”
- Operational test for intelligent behavior: the *Imitation Game*

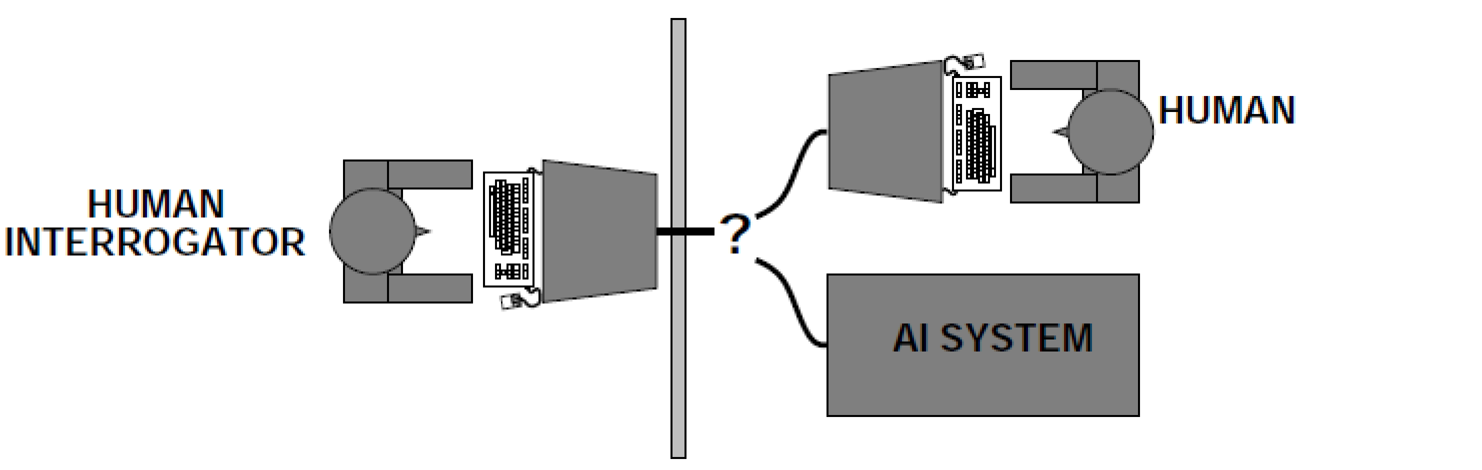

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes (e.g., see “Loebner prize”)

“Artificial” Intelligence

- Good workable definition: “Systems that are intelligent but non-natural (i.e., non-biological, engineered, etc.)”
- Two main directions: **thinking** vs. **action**
- Two main aims: **human-like** vs. **rational**
- Note: *rational* as *normative* (or prescriptive) rather than descriptive
- 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
  - Rational behavior: doing the right thing
- The right thing: that which is expected to maximize goal achievement, given the available information

Subfields of AI

- Distinguish “classical AI” (GOFAI) from “new AI” (NFAI)
- Classical areas of focus: symbolic knowledge representation, planning, reasoning, and learning
- Typical apparatus: logical language and search algorithms (Problem solving = knowledge representation + search)
- Classical systems: cognitive architectures (SOAR, ACT, EPIC, PRS, etc.), expert systems (MYCIN, Geologist, etc.)
- Non-classical areas: distributed subsymbolic representations, non-logical rea soning methods (spreading activation, Bayesian reasoning, Markov logics, etc.)
- Classical systems: neural networks, genetic algorithms, agent-based computing, articial life, Bayesian networks, etc.

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
- **Psychology:** adaptation, phenomena of perception and motor control, experimental techniques (psychophysics, etc.)
- **Economics:** formal theory of rational decisions, game theory
- **Linguistics:** knowledge representation, grammar, semantics
- **Neuroscience:** plastic physical substrate for mental activity
- **Control theory:** homeostatic systems, stability, simple optimal agent designs

Spotted history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's “Computing Machinery and Intelligence”
- 1952-69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program,
  Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 Dartmouth meeting: \Articial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 19660-74 AI discovers computational complexity
- Neural network research almost disappears
- 1969-79 Early development of knowledge-based systems
- 1980-88 Expert systems industry booms
- 1988-93 Expert systems industry busts: \AI Winter"
- 1985-95 Neural networks return to popularity
- 1988- Resurgence of probability; general increase in technical depth
  “Nouvelle AI": ALife, GAs, soft computing
- 1995- Agents, agents, everywhere...
- 2003- Human-level AI back on the agenda

Basic approaches in AI

- The role of learning:
  - If you want an intelligent machine, program in the intelligence.
  - If you want an intelligent machine, make it a good learner and send it out into the world.
- The role of the hardware:
  - Intelligence is a software problem. An intelligent program can be run on a brain or a computer.
  - The hardware is relevant: we should look for intelligence that is based on the properties of nervous systems.

Domain-independent methods in AI

- Neat AI (mostly overlapping with Haugeland's GOFAI-”good old-fashioned AI”):
  - Theories should be elegant and parsimonious.
  - We should understand precisely what our theories can do and how they behave.
  - Most (or all) of intelligence is governed by general principles.
- Scruffy AI (mostly overlapping with Haugeland's NFAI-”new-fangled AI”):
  - The mind is a kludge. To make things ecient, inelegant shortcuts are often appropriate.

    It may be impossible to come to a precise understanding of our theories.
  - There are only a few general principles that apply across domains. Intelligence comes from domain knowledge

AI “Spin-offs” in CS

- Neural networks
- “AI robotics" (behavior-based robotics, SLAM, HRI, ...)
- Machine learning, data mining
- Machine vision and pattern detection
- Agent-based (distributed) computing
- (Statistical) natural language processing (NLP)
- Cognitive systems

Which of the following can be done todyay?

- Play a decent game of table tennis
- Drive safely along a curving mountain road
- Drive safely along Mass Ave
- Buy a week's worth of groceries on the web
- Buy a week's worth of groceries at Whole Foods
- Play a decent game of bridge
- Discover and prove a new mathematical theorem
- Write an intentionally funny story
- Give competent legal advice in a specialized area of law
- Translate spoken English into spoken Swedish in real time
- Converse successfully with another person for an hour
- Perform a complex surgical operation
- Unload any dishwasher and put everything away

Which of the following can be done todyay?

- Play a decent game of table tennis
- Drive safely along a curving mountain road
- Drive safely along Mass Ave
- Buy a week's worth of groceries on the web
- Buy a week's worth of groceries at Whole Foods
- Play a decent game of bridge
- Discover and prove a new mathematical theorem
- Write an intentionally funny story
- Give competent legal advice in a specialized area of law
- Translate spoken English into spoken Swedish in real time
- Converse successfully with another person for an hour
- Perform a complex surgical operation
- Unload any dishwasher and put everything away

Unintentionally funny stories

- One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving told him there was a beehive in the oak tree. Joe threatened to hit Irving if he didn't tell him where some honey was. The End.

- Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned. The End.

- Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese. The fox walked over to the crow. The End.

Unintentionally funny stories

- Joe Bear was hungry. He asked Irving Bird where some honey was. Irving refused to tell him, so Joe oered to bring him a worm if he'd tell him where some honey was. Irving agreed. But Joe didn't know where any worms were, so he asked Irving, who refused to say. So Joe oered to bring him a worm if he'd tell him where a worm was. Irving agreed. But Joe didn't know where any worms were, so he asked Irving, who refused to say. So Joe oered to bring him a worm if he'd tell him where a worm was...

Example: Game Playing

- "Unpredictable" opponent: solution is a strategy specifying a move for every possible opponent reply
- Time limits: unlikely to find goal, must approximate
- Plan of attack:
  - Computer considers possible lines of play (Babbage, 1846)
  - Algorithm for perfect play (Zermelo, 1912; Von Neumann, 1944)
  - Finite horizon, approximate evaluation (Zuse, 1945; Wiener, 1948; Shannon, 1950)
  - First chess program (Turing, 1951)
  - Machine learning to improve evaluation accuracy (Samuel, 1952-57)
  - Pruning to allow deeper search (McCarthy, 1956)

Example: Game Playing

- **Checkers:** Chinook ended 40-year-reign of human world champion Marion Tinsley in 1994 (used an endgame database dening perfect play for all positions involving 8 or fewer pieces on the board, a total of 443,748,401,247 positions)
- **Chess:** Deep Blue defeated human world champion Gary Kasparov in a six-game match in 1997 (it searches 200 million positions per second, uses very sophisticated evaluation, and undisclosed methods for extending some lines of search up to 40 ply)
- **Othello:** human champions refuse to compete against computers, who are too good!
- **Go:** human champions refuse to compete against computers, who are too bad (in Go, the branching factor > 300, so most programs use pattern knowledge bases to suggest plausible moves)

Types of games

	deterministic	chance
perfect information	chess, checkers, go, othello	backgammon monopoly
imperfect information	battleships, blind tictactoe	bridge, poker, scrabble nuclear war

- Perfect play for deterministic, perfect-information games
- Idea: choose move to position with highest minimax value

  = best achievable payoff against best play

Deterministic games

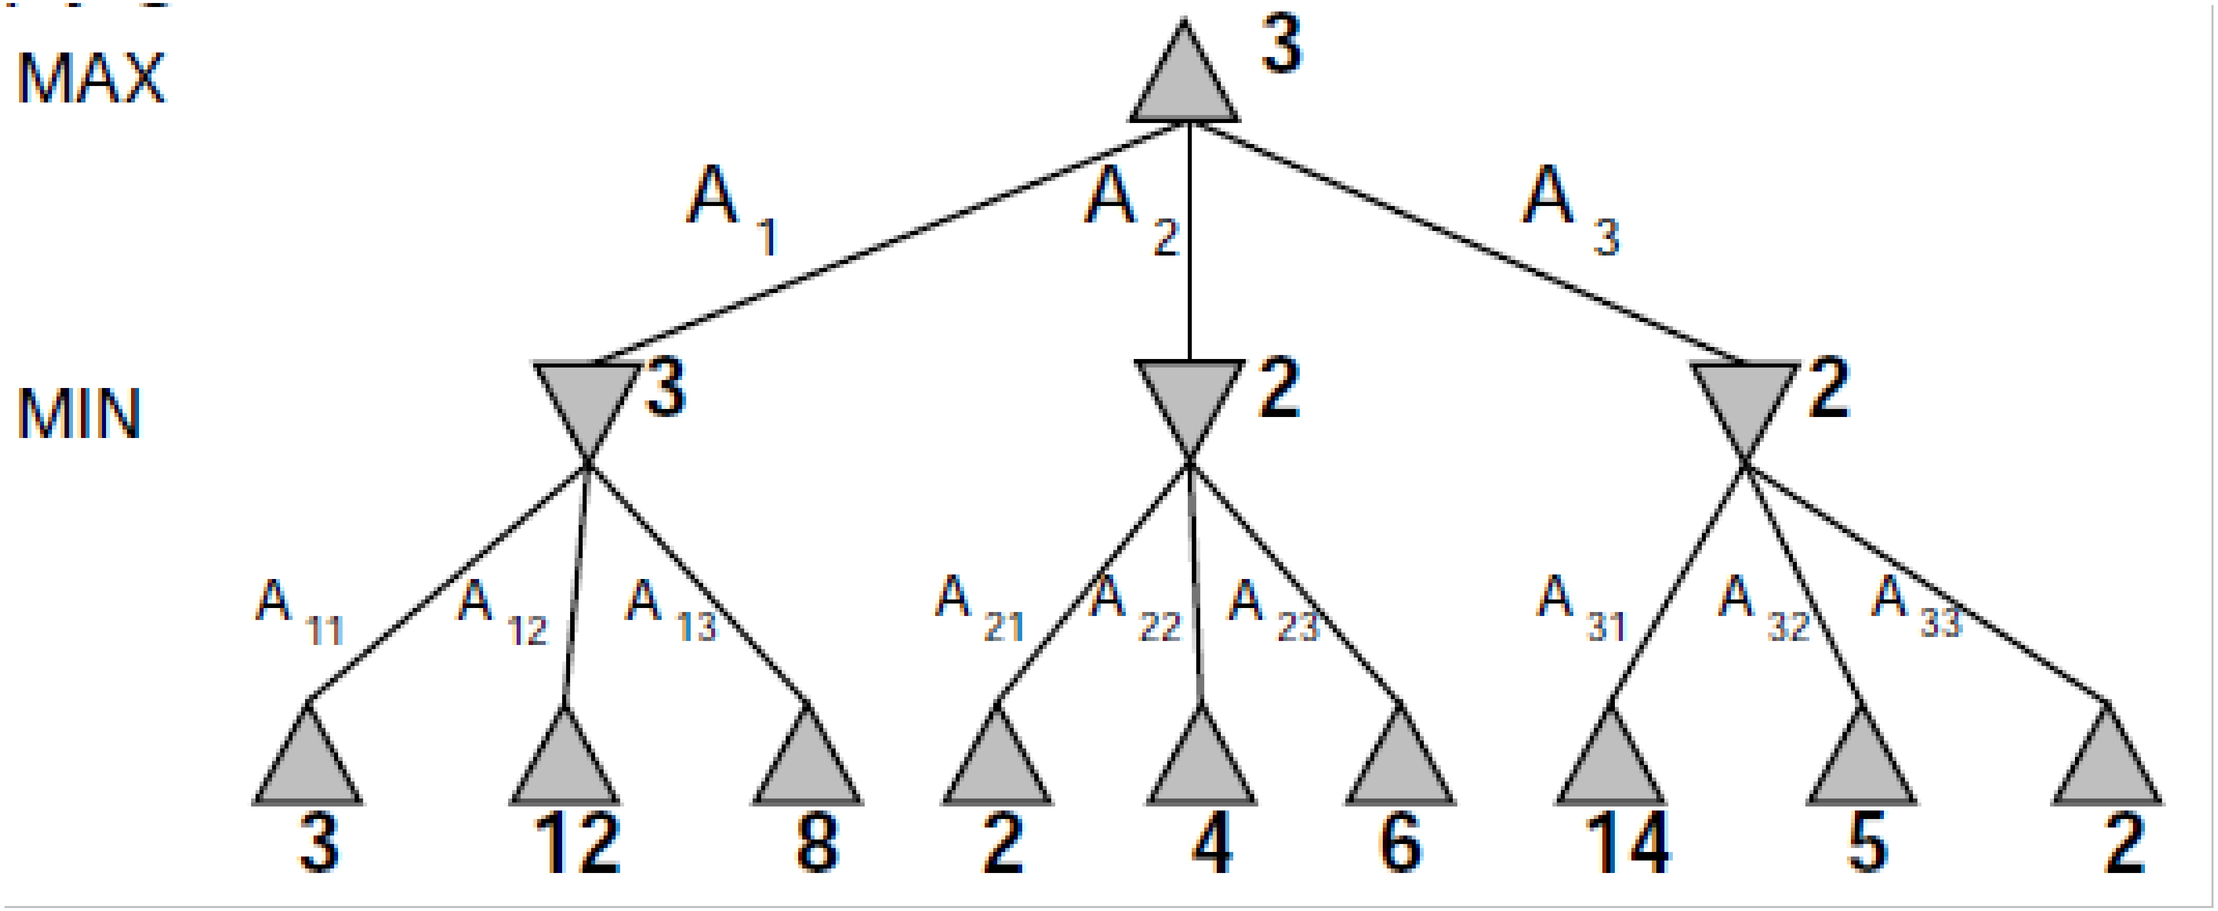

Nondeterministic games