ARTIFICIAL INTELLIGENCE

Introduction

Outline

- What is AI?
- A brief history
- The state of the art

What is Intelligence

- A wish-list of general characteristics of intelligence
  - **Perception**: manipulation, interpretation of data provided by sensors
  - **Action**: control, and use of effectors to accomplish a variety of tasks
  - **Reasoning**: deductive (logical) inference, inductive inference,
  - **Learning**: adapting behaviour to better cope with changing environments, discovery of patterns, learning to reason, plan, and act.
  - **Communication**: with other intelligent agents including humans using signals, signs, icons, …
  - **Planning**: formulation of plans -- sequences or agenda of actions to accomplish externally or internally determined goals
  - ...

What is AI?

- The exciting new effort to make computers think .. *machines with minds*
- AI is the art of creating machines that perform functions that require intelligence when performed by humans
- AI is the study of the computations that make it possible to perceive, reason, and act
- AI is the enterprise of design and analysis of intelligent agents.

Acting humanly: Turing Test

- Alan Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" → "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game

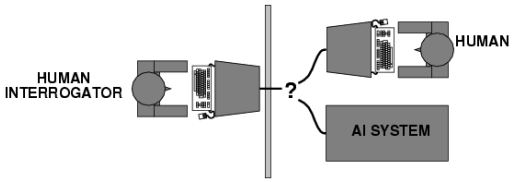

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Annual Loebner prize competition (since 1990): the first prize of $100,000 to be awarded to the first program that passes the "unrestricted" Turing test

What is AI?

Thinking humanly	Thinking rationally
Acting humanly	Acting rationally

- Are you concerned with *thought processes/reasoning* or *behavior*?
- Do you want to model humans or measure against an *ideal* concept of intelligence, **rationality**

Acting rationally: rational agent

- This course is about designing rational agents
  - Advocated by the textbook
- Rational behavior: doing the right thing
  - The right thing: that which is expected to maximize goal achievement, given the available information
- An agent is an entity that perceives and acts
- A **rational agent** is one that acts so as to achieve the best outcome

Brief history of AI

- 1943 McCulloch & Pitts: model of artificial neurons
- 1950 Turing's "Computing Machinery and Intelligence“
- 1956 McCarthy, Minsky, Newell, Simon, Turing, Uhr, et al., Dartmouth workshop: "Artificial Intelligence" adopted
- 1952—69 Early enthusiasm, great expectations, optimism fueled by early success on some problems thought to be hard (e.g., theorem proving)
- 1966—73 Collapse in AI research: Progress was slower than expected, Unrealistic predictions, Herbert Simon (1957) chess champion in 10 years
  AI discovers computational complexity

Brief history of AI

- 1969—79 Early development of knowledge-based systems
- 1980--- AI becomes an industry: Expert systems industry booms, then busts (88-93): “AI Winter“
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science: revolution in the content and methodology of work in AI, great advance, active research field since

Brief History of AI

- Mid 1990s-present : The emergence of intelligent agents
  - AI technologies continue to find applications in
    - information retrieval,
    - data mining and knowledge discovery,
    - customizable software systems,
    - smart devices (e.g., homes, automobiles),
    - agile manufacturing systems,
    - autonomous vehicles,
    - Bioinformatics
    - Internet tools: search engines, recommender systems
    - etc.
  - Steady progress on fundamental AI research problems continues.

State of the art

- Deep Blue defeated the reigning world chess champion Garry Kasparov in 1997
- Chinook defeated human checkers champion in 1994, can't lose at checkers 2007
- `Proverb` (1999) solves crossword puzzles better than most humans
- The IBM supercomputer Watson beat human champions on 'Jeopardy!‘ (2011)

State of the art

- Proved a mathematical conjecture (Robbins conjecture, 1996) unsolved for 60 years
- 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, save the US more money than spent on all AI research since 1950
- NASA's on-board autonomous planning program controlled the scheduling of operations for a spacecraft (2000), Mars Exploration Rovers (2004).

State of the art

- Driverless cars
  - No hands across America (driving autonomously 98% of the time from Pittsburgh to San Diego, 1995)
  - DARPA Grand Challenge prize competition for driverless cars (2007): 60-mile urban area course completed in less than 6 hours obeying all traffic regulations, $2 million winner Tartan Racing
  - Driverless vans completed 8000-mile trip from Italy to China (July-Oct 2010)
  - Google's driverless cars 140,000 miles driven in testing
- Robotic vacuum cleaners
- Speech recognition, machine translation, …
- Robotic scientists make scientific discoveries by itself (Science 2009): formulation of hypotheses and designing of experiments to test them