Artificial Intelligence

Dr. Qaiser Abbas
Department of Computer Science & IT,
University of Sargodha
qaiser.abbas@uos.edu.pk

Outline

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

Course Outcomes

- CO 1 **Investigate** the concepts and terminologies of intelligent agents
- CO 2 **Create** problem solving algorithms using uninformed searches
- CO 3 **Design** problem solving algorithms using heuristic searches
- CO 4 **Gauge** the performance of local and optimized search methods
- CO 5 **Perform** evolutionary strategies for solving a given problem
- CO 6 **Implement** constraint satisfaction for a given set of problems
- CO 7 **Compare** the performance of adversarial search algorithms
- CO 8 **Design** an intelligent agent to solve the problem at hand

Course Assessment Methods:

- Homework Assignments There will be homework assignments, which will measure the level of understanding and the skills attained by the students related to the assignment topic. At least two assignments will be the part of assessment.
- Quizzes There will be at least two quizzes. The purpose of the quizzes is to motivate the students to study the course material timely and make sure that the learning process is reinforced. Through this way, it will also measure the level of understanding of the students related to course material.
- Examinations There will be a mid term and a final term exams. The purpose of the exams is to measure the level of understanding of the students related to the course material.
- Lab Projects There will be at least two Lab Projects. The purpose of the projects is to motivate the students to get hands on experience with the AI methods and procedures.

Text and Reference Books

- S. Russell, and P. Norvig, ***Artificial Intelligence: A Modern Approach*** (3rd edition)., Prentice Hall 2010.

- M. Tim Jones, ***Artificial Intelligence: A Systems Approach***, Jones and Bartlett Publishers, Inc; 1st Edition

What is Intelligence?

- The dream of AI has been to build. . .
  "... machines that can think, that learn and that create."
- "The question of whether Machines Can Think. . .
  . . . is about as relevant as the question whether Submarines Can Swim."

  Dijkstra (1984)

Strong and Weak AI

- One may dream about. . .

  . . . that computers can be made to think on a level at least equal to humans, that they can be conscious and experience emotions.

  Strong AI

- This course is about. . .

  . . . adding “thinking-like” features to computers to make them more useful tools. That is, “not obviously machine like”.

  Weak AI

Weak AI

- Weak AI is a category that is flexible, as soon as we understand how an AI-program works, it appears less “intelligent”.
- And as soon as a part of AI is successful, it becomes an own research area! E.g. large parts of advanced search, parts of language understanding, parts of machine learning and probabilistic learning etc.
- And AI is left with the remaining hard-to-solve problems!

What is AI?

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

Acting humanly: The Turing test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" −→ "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation (the action of using someone or something as a model or simulation) Game

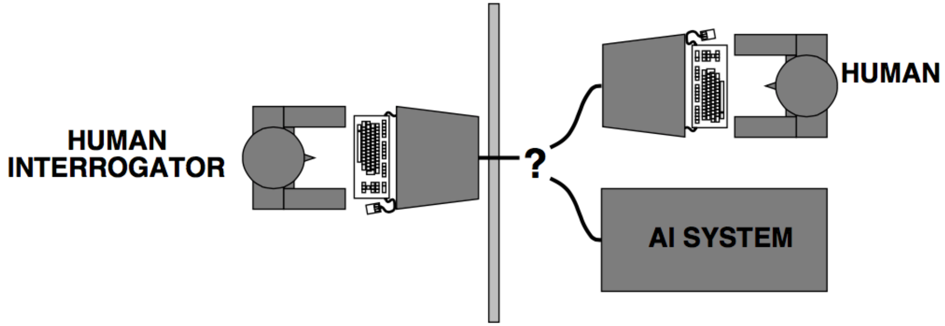

- Predicted that by the year 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated (expect, predict) all major arguments against AI in the following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning, computer vision, robotics

Problem: Turing test is not reproducible, constructive, or amenable (controllable) to mathematical analysis

Thinking humanly: Cognitive Science

- If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. We need to get *inside* the actual workings of human minds.

- There are three ways to do this:
  - through introspection—trying to catch our own thoughts as they go by;
  - through psychological experiments—observing a person in action;
  - and through brain imaging—observing the brain in action.
-

  Once we have a sufficiently precise theory of the mind, it becomes possible to express the theory as a computer program. If the program's input–output behavior matches corresponding human behavior, that is evidence that some of the program's mechanisms could also be operating in humans.

Thinking humanly: Cognitive Science

1960s: "Cognitive revolution"

- Requires scientific theories of internal activities of the brain
  - What level of abstraction? "Knowledge" or "circuits"?
  - How to validate? Requires
    - Predicting and testing behavior of human subjects (top-down)
    - Or Direct identification from neurological data (bottom-up)
- Both approaches (roughly, Cognitive Science and Cognitive Neuroscience) are now distinct from AI
- Both share with AI the following characteristic:

  the available theories do not explain (or engender (cause or give rise to)) anything resembling human-level general intelligence

- Hence, all three fields share one principal direction!

Thinking rationally: Laws of Thought

Normative or prescriptive (relating to some standard) 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
- Direct line through mathematics and philosophy to modern AI

- Problems:
  - Not all intelligent behavior is mediated by logical deliberation (careful discussion or consideration).
  - What is the purpose of thinking? What thoughts should I have out of all the thoughts (logical or otherwise) that I could have?

Thinking rationally: Laws of Thought

- There are two main obstacles to this approach.
  - First, it is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain.
  - Second, there is a big difference between solving a problem "in principle" and solving it in practice. Even problems with just a few hundred facts can exhaust the computational resources of any computer unless it has some guidance as to which reasoning steps to try first.

Acting rationally: Rational agents

- Rational behavior: "doing the right thing", i.e., 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

- An agent is an entity that perceives and acts.
  - This course (and the course book) is about designing rational agents

- Abstractly, an agent is a function from percept histories to actions:

  $f : P* \rightarrow A$, For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

- Caveat (warning): computational limitations make perfect rationality unachievable

  → design best program for given machine resources

Acting rationally: Rational agents

- Knowledge representation and reasoning enable agents to reach good decisions.

- The rational-agent approach has two advantages over the other approaches.
  - First, it is more general than the "laws of thought" approach because correct inference is just one of several possible mechanisms for achieving rationality.
  - Second, it is more amenable to scientific development than are approaches based on human behavior or human thought.

Contribution to AI

• 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 (substance) for mental activity
• Psychology	phenomena of perception and motor control, experimental techniques
• Computer engg.	building fast computers
• Control theory	design systems that maximize an objective function over time
• Linguistics	knowledge representation, grammar

Potted 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: e.g., Samuel's checkers program, Gelernter's Geometry Engine, Newell & Simon's Logic Theorist and General Problem Solver
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity, Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1971 Terry Winograd's Shrdlu dialogue system
- 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 . . .
- 1997 IBM Deep Blue beats the World Chess Champion
- 2001– Very large datasets: Google gigaword corpus, Wikipedia
- 2003– Human-level AI back on the agenda
- 2011 IBM Watson wins Jeopardy
- 2012 US state of Nevada permits driverless cars

State of the art

- What can AI do today? A concise answer is difficult because there are so many activities in so many subfields. Here we sample a few applications;
- Understanding and Processing Natural Languages:
- Driverless vehicles:
- Speech recognition:
- Autonomous planning and scheduling:
- Game playing:
- Spam fighting:
- Logistics (the detailed organization and implementation of a complex operation) planning:
- Robotics:
- Machine Translation:

Homework-1

1.7 Examine the AI literature to discover whether the following tasks can currently be solved by computers:

a. Playing a decent game of table tennis (ping-pong).

b. Driving in the center of Cairo.

c. Buying a week's worth of groceries at the market.

d. Buying a week's worth of groceries on the web.

e. Playing a decent game of bridge at a competitive level.

f. Discovering and proving new mathematical theorems.

g. Writing an intentionally funny story.

h. Giving competent legal advice in a specialized area of law.

i. Translating spoken English into spoken Swedish in real time.

j. Performing a complex surgical operation.

For the currently infeasible tasks, try to find out what the difficulties are and predict when, if ever, they will be overcome.