Artificial Intelligence

Lecture 01 - Introduction

What Is Artificial Intelligence?

- Difficult question
  - What is intelligence?
- Answer depends on who you ask
- John McCarthy, who coined the term Artificial Intelligence in 1956, defines it as:
  - "the science and engineering of making intelligent machines", especially intelligent computer programs

Intelligence

- Definition:
  - the ability to learn, understand and make judgements or have opinions that are based on reason
- Relate to tasks involving higher mental processes
  - Examples:
    - creativity, solving problems, pattern recognition, classification, learning, induction, deduction, building analogies, optimization, language processing, knowledge and many more
- Intelligence is the computational part of the ability to achieve goals

Intelligent Behaviour

- **Perceiving** one's environment,
- **Acting** in complex environments,
- **Learning** and understanding from experience,
- **Reasoning** to solve problems and discover hidden knowledge,
- **Knowledge** applying successfully in new situations,
- **Thinking** abstractly, using analogies,
- **Communicating** with others, and more like
  - **Creativity**, **Ingenuity**, **Expressiveness**, **Curiosity**

More Definitions of AI

a) "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)

b) "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)

c) "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)

d) "A field of study that seeks to explain and emulate intelligent behaviour in terms of computational processes" (Schalkoff, 1990)

"The branch of computer science that is concerned with the automation of intelligent behaviour" (Luger and Stubblefield, 1993)

More Definitions of AI

- The definitions on the top,
  - (a) and (b) are concerned with reasoning, whereas those on the bottom,
  - (c) and (d) address behaviour
- The definitions on the left,
  - (a) and (c) measure success in terms of human performance, whereas those on the right,
  - (b) and (d) measure ideal concept of intelligence called rationality
- Note: A system is **rational** if it does the right thing

Strong AI

- Generally, artificial intelligence research aims to create AI that can **replicate human intelligence** completely
- Strong AI refers to a machine that approaches or supersedes human intelligence,
  - If it can do typically human tasks,
  - If it can apply a wide range of background knowledge and
  - If it has some degree of self-consciousness
- Strong AI aims to build machines whose overall intellectual ability is indistinguishable from that of a human being

Weak AI

- Weak AI refers to the use of software to study or accomplish specific problem solving or reasoning tasks that do not encompass the full range of human cognitive abilities.
- Example: a chess program such as Deep Blue
- Weak AI does not achieve self-awareness; it demonstrates wide range of human-level cognitive abilities; it is merely an intelligent, a specific problem-solver

Goals of AI

- The definitions of AI gives four possible goals to pursue:
  1. Systems that think like humans.
  2. Systems that think rationally
  3. Systems that act like humans
  4. Systems that act rationally

- Modern AI focuses on the last definition
  - we will also focus on this “engineering” approach
  - success is judged by how well the agent performs

Goals of AI

- General AI Goal
  - Replicate human intelligence: still a distant goal
  - Solve knowledge intensive tasks
  - Make an intelligent connection between perception and action
  - Enhance human-human, human-computer and computer to computer interaction / communication
- Engineering based AI Goal
  - Develop concepts, theory and practice of building intelligent machines
  - Emphasis is on system building
- Science based AI Goal
  - Develop concepts, mechanisms and vocabulary to understand biological intelligent behaviour
  - Emphasis is on understanding

AI Approaches

- Think human-like
- Think Rationally
- Act Human-like
- Act Rationally

Think human-like

- An exciting new effort to make computers think; machines that it is, the with minds, in the full and literal sense
- Focus is not just on behaviour and I/O, but looks at reasoning process
- Computational model as to how results were obtained
- Goal is not just to produce human-like behaviour but to produce a sequence of steps of the reasoning process, similar to the steps followed by a human in solving the same task

Cognitive Science

- Aims to develop, explore and evaluate theories of how the mind works through the use of computational models.
- The important is not what is done but how it is done; means intelligent behaviour is not enough, the program must operate in an intelligent manner
- Example: The Chess programs are successful, but say little about the ways humans play chess

Think Rationally

- The study of mental faculties through the use of computational models; that it is, the study of the computations that make it possible to perceive, reason, and act
- Focus is on inference mechanisms that are provably correct and guarantee an optimal solution
- Develop systems of representation to allow inferences to be like “Socrates is a man. All men are mortal. Therefore Socrates is mortal.”
- Goal is to formalize the reasoning process as a system of logical rules and procedures for inference
- The issue is, not all problems can be solved just by reasoning and inferences

Act Human-like

- The art of creating machines that perform functions requiring **intelligence** when performed by people; that it is the study of, how to make computers do things which at the moment people do better
- Focus is on action, and not intelligent behaviour centred around representation of the world
- A Behaviourist approach, is not concerned with how to get results but to the similarity to what human results are

Act Human-like

- Example: Turing Test
  - 3 rooms contain: a person, a computer, and an interrogator
  - The interrogator can communicate with the other 2 by teletype (to avoid the machine imitate the appearance or voice of the person)
  - The interrogator tries to determine which is the person and which is the machine

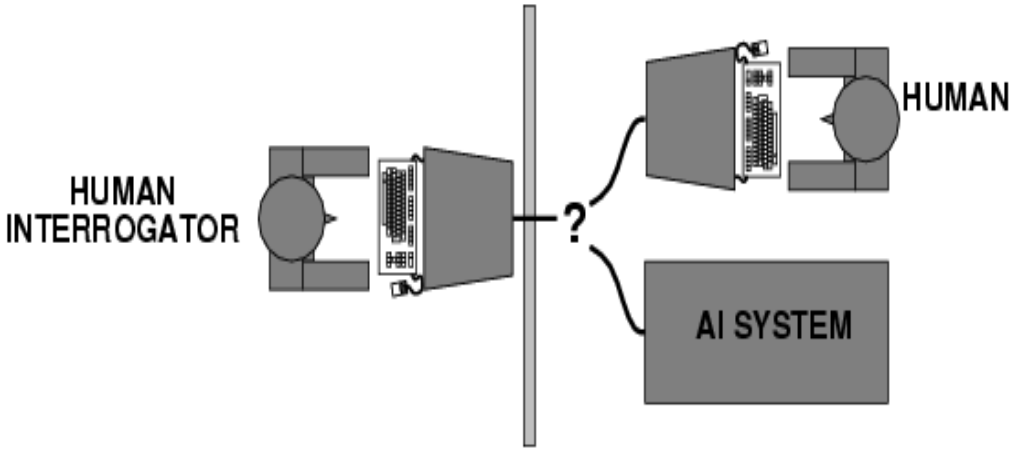

Act Human-like

- The machine tries to fool the interrogator to believe that it is the human, and the person also tries to convince the interrogator that it is the human
- If the machine succeeds in fooling the interrogator, then conclude that the machine is intelligent

- Goal is to develop systems that are human-like

Act Rationally

- Tries to explain and emulate i**ntelligent behaviour** in terms of computational processes; that it is concerned with the automation of intelligence
- Focus is on systems that act sufficiently if not optimally in all situations;
- It is passable to have imperfect reasoning if the job gets done
- Goal is to develop systems that are rational and sufficient

Branches of AI

- Major branches
  - Logical AI
  - Search
  - Pattern Recognition
  - Knowledge Representation
  - Inference
  - Common Sense Knowledge and Reasoning

Branches of AI

- Major branches continued
  - Machine Learning
  - Planning
  - Epistemology
  - Ontology
  - Heuristics
  - Genetic Programming

Logical AI

- Logic is a language for reasoning; a collection of rules used while doing logical reasoning
- Types of logic
  - Propositional logic - logic of sentences
  - predicate logic - logic of objects
  - logic involving uncertainties
  - Fuzzy logic - dealing with fuzziness
  - Temporal logic, etc

Logical AI

- Propositional logic and Predicate logic are fundamental to all logic
- Propositional logic
  - Propositions are “Sentences”; either true or false but not both.
  - A sentence is smallest unit in propositional logic
  - If proposition is true, then truth value is "true"; else “false”
  - Example : Sentence"Grass is green"; Truth value “ true”; Proposition is “yes”

Logical AI

- Predicate logic
  - Predicate is a function may be true or false for arguments
  - Predicate logic are rules that govern quantifiers
  - Predicate logic is propositional logic added with quantifiers
  - Examples:
    - “The car Tom is driving is blue”,
    - “The sky is blue”,
    - “The cover of this book is blue”
  - Predicate is blue, give a name B;
  - Sentence represented as B(x); read B(x) as "x is blue" ;
  - Object represented as x

Search

- Search is a problem-solving technique that systematically consider all possible action to find a path from initial state to target state
  - Many search techniques; the most fundamental are
    - Hill climbing
    - Breadth first
    - Depth first
    - Least cost

Search

- Search components
  - Initial state - First location
  - Available actions - Successor function : reachable states
  - Goal test - Conditions for goal satisfaction
  - Path cost - Cost of sequence from initial state to reachable state
- Search objective
  - Transform initial state into goal state - find a sequence of actions.
- Search solution
  - Path from initial state to goal - optimal if lowest cost

Pattern Recognition

- When a program makes observations of some kind, it is often programmed to compare what it sees with a pattern
- Pattern recognition problems
  - Machine vision
  - Character recognition
  - Computer aided diagnosis
  - Speech recognition

Knowledge Representation

- How do we represent what we know?
  - Knowledge is a collection of **facts**
  - To manipulate these facts by a program, a suitable representation is required
  - A good representation facilitates problem solving
- Knowledge representation techniques
  - Different types of knowledge require different types of representation
    - Predicate Logic
    - Semantic Networks
    - Frames and Scripts
    - Production Rules

Inference

- Inference is the act or process of deriving a conclusion based solely on what one already knows
- It is deduction of new facts from old ones
- Logic captures inference
  - Deductive Inference
  - Inductive Inference

Common Sense Knowledge and Reasoning

- Common sense is the mental skills that most people have
- It is the ability to analyse a situation based on its context
- People can think
  - because the brain contain vast libraries of common sense knowledge and has means for organizing, acquiring, and using such knowledge
- Computer can not think
  - the computers programs do many things, they can play chess at the level of best players but cannot match capabilities of a 3 year old child at recognizing objects
- Currently, computers lack common sense

Common Sense Knowledge and Reasoning

- Teaching computers common sense
  - Project “OpenMind” at MIT - Here the goal is to teach a computer things that human take them for granted; here the knowledge is represented in the form of Semantic net, Probabilistic graphical models, and Story scripts.
  - Project “Cyc” - It is an attempt to manually build a database of human common sense knowledge; it has 1.5 million collection of common sense facts, but still far away from several hundred million needed

Learning

- Programs learn from what the facts or the behaviours can represent
- Major Paradigms of Machine Learning
  - Rote: Learning by memorization
  - Induction: Learning by example
  - Analogy: Learning from similarities
  - Genetic Algorithms: Learning by mimicking processes nature uses
  - Reinforcement: Learning from actions

Planning

- A plan is a representation of a course of action
  - Planning is a problem solving technique
  - Planning is a reasonable series of actions to accomplish a goal
- Planning programs
  - Start with facts about the world, particularly
    - facts about the effects of actions,
    - facts about the particular situation, and
    - statement of a goal

Planning

- Benefits of planning
  - reducing search,
  - resolving goal conflicts, and
  - providing a basis for error recovery.
- Strategy for planning
  - A strategy is just a sequence of actions
  - From facts the program generates a strategy for achieving the goal

Epistemology

- There are various kinds of knowledge:
  - knowing how to do something (e.g., how to ride a bicycle),
  - knowing someone in person, and
  - knowing a place or a city
- Epistemology is the study of the kinds of knowledge that are required for solving problems in the world

Ontology

- Ontology is concerned with existence; a study of the categories of things that exist or may exist in some domain
  - Ontology is a data model, represents a domain and is used to reason about the objects in that domain and the relations between them
  - Ontology is used in artificial intelligence, as a form of knowledge representation about the world or some part of it

Heuristics

- Heuristics are in common use as **rules of thumb**
- In computer science, a heuristic is an algorithm with provably good run times and with provably good or optimal solution
- Heuristics are intended to gain computational performance or conceptual simplicity, potentially at the cost of accuracy or precision
- People use heuristics to make decisions, come to judgements, and solve problems, when facing complex problems or incomplete information
- These rules work well under most circumstances

Genetic programming (GP)

- Genetic programming is an automated method for creating program from a high-level problem statement
- GP starts from a high-level statement of the requirements of a problem and attempts to produce a computer program that solves the problem
- The user (human) communicates the high-level statement of the problem to the GP system by performing certain well-defined preparatory steps

History of AI

- 1943: early beginnings
  - McCulloch & Pitts: Boolean circuit model of brain
- 1950: Turing
  - Turing's "Computing Machinery and Intelligence“
- 1956: birth of AI
  - Dartmouth meeting: "Artificial Intelligence“ name adopted
- 1950s: initial promise
  - Early AI programs, including
    - Samuel's checkers program
    - Newell & Simon's Logic Theorist

History of AI

- 1955-65: “great enthusiasm”
  - Newell and Simon: GPS, general problem solver
  - Gelertner: Geometry Theorem Prover
  - McCarthy: invention of LISP
- 1966-73: Reality dawns
  - Realization that many AI problems are intractable
  - Limitations of existing neural network methods identified
  - Neural network research almost disappears

History of AI

- 1969-85: Adding domain knowledge
  - Development of knowledge-based systems
  - Success of rule-based expert systems,
    - e.g., DENDRAL, MYCIN
    - But were brittle and did not scale well in practice
- 1986- Rise of machine learning
  - Neural networks return to popularity
  - Major advances in machine learning algorithms and applications

History of AI

- 1990- Role of uncertainty
  - Bayesian networks as a knowledge representation framework
- 1995- AI as Science
  - Integration of learning, reasoning, knowledge representation
  - AI methods used in vision, language, data mining, etc.

AI Applications

- Game Playing
- Speech Recognition
- Understanding Natural Language
- Computer Vision
- Expert Systems

Game Playing

- Games are Interactive computer program, an emerging area in which the goals of human-level AI are pursued
- Games are made by creating human level artificially intelligent entities,
  - e.g. enemies, partners, and support characters that act just like humans

Game Playing

- Game play is a search problem defined by:
  - Initial state - board
  - Expand function - build all successor states
  - Cost function - pay off of the state
  - Goal test - ultimate state with maximal pay off
- Game playing is characterized by:
  - "Unpredictable" opponent
  - Need to specify move for every possible opponent reply
  - Time limits - games become boring if there is no action for too long a time; opponents are unlikely to find goal, must approximate

Game Playing

- Computer Games
  - Computers perform at champion level games,
    - Examples: Checkers, Chess, Othello, Backgammon
  - Computers perform well games
    - Example: Bridge
  - Computers still do badly
    - Example: Go, Hex
- The Deep Blue Chess program won over world champion Gary Kasparov

Speech Recognition

- A process of converting a speech signal to a sequence of words
- In 1990s, computer speech recognition reached a practical level for limited purposes
- Using computers recognizing speech is quite convenient, but most users find the keyboard and the mouse still more convenient

Speech Recognition

- The typical usages are:
  - Voice dialling (Call home)
  - Call routing (collect call)
  - Data entry (credit card number)
  - Speaker recognition
- The spoken language interface PEGASUS in the American Airlines' EAASY SABRE reservation system, allows users to obtain flight information and make reservations over the telephone

Understanding Natural Language

- Natural language processing (NLP) does automated generation and understanding of natural human languages
  - Natural language generation system
    - Converts information from computer databases into normal-sounding human language
  - Natural language understanding system
    - Converts samples of human language into more formal representations that are easier for computer programs to manipulate

Understanding Natural Language

- Some major tasks in NLP:
  - Text-to-Speech (TTS) system: converts normal language text into speech
  - Speech recognition (SR) system: process of converting a speech signal to a sequence of words
  - Machine translation (MT) system: translate text or speech from one natural language to another
  - Information retrieval (IR) system: search for information from databases such as Internet or World Wide Web or Intranets

Computer Vision

- It is a combination of concepts, techniques and ideas from: Digital Image Processing, Pattern Recognition, Artificial Intelligence and Computer Graphics
- The world is composed of 3-D objects, but the inputs to the human eye and computers' TV cameras are 2-D
- Some useful programs can work solely in 2-D, but full computer vision requires partial 3-D information that is not just a set of 2-D views
- At present there are only limited ways of representing 3-D information directly, and they are not as good as what humans evidently use

Computer Vision

- Examples
  - Face recognition
  - Autonomous driving
    - The ALVINN system, autonomously drove a van from Washington, D.C. to San Diego, averaging 63 mph day and night, and in all weather conditions.
  - Other usages:
    - Handwriting recognition, Baggage inspection, Manufacturing inspection, Photo interpretation

Expert Systems

- Systems in which human expertise is held in the form of rules
  - It enable the system to diagnose situations without the human expert being present
  - A man-machine system with specialized problem-solving expertise
    - The "expertise" consists of knowledge about a particular domain, understanding of problems within that domain, and "skill" at solving some of these problems
  - Knowledge base
    - A knowledge engineer interviews experts in a certain domain and tries to embody their knowledge in a computer program for carrying out some task

Expert Systems

- Expert systems rely on knowledge of human experts

- One of the first expert systems was MYCIN in 1974, which diagnosed bacterial infections of the blood and suggested treatments

Expert Systems

- Some applications areas:
  - Diagnosis and Troubleshooting:
    - deduces faults and suggest corrective actions for a malfunctioning device or process
  - Planning and Scheduling:
    - analysing a set of goals to determine and ordering a set of actions taking into account the constraints; e.g. airline scheduling of flights
  - Financial Decision Making:
    - advisory programs assists bankers to make loans, Insurance companies to assess the risk presented by the customer, etc.
  - Process Monitoring and Control:
    - analyses real-time data, noticing anomalies, predicting trends, and controlling optimality and do failure correction