



Artificial Intelligence and Machine Learning- BDS602

Course Coordinator:

Dr. Neethi M V

Assistant Professor

Dept. of CSE-DS

ATMECE, Mysuru



Module 1 – Introduction

Session 1

- What is AI?
- Foundations and History of AI



INTRODUCTION

- We call ourselves *Homo sapiens*—man the wise—because our intelligence is so important to us. For thousands of years, we have tried to understand how we think; that is, how a mere handful of matter can perceive, understand, predict, and manipulate a world far larger and more complicated than itself.
- The field of artificial intelligence, or AI, goes further still: it attempts not just to understand but also to build intelligent entities.



INTRODUCTION => WHAT IS AI?

Thinking Humanly

“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)

Acting Humanly

“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)

Thinking Rationally

“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)

Acting Rationally

“Computational Intelligence is the study of the design of intelligent agents.” (Poole *et al.*, 1998)

“AI . . . is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)

- In Figure 1.1 we see eight definitions of AI, laid out along two dimensions.
- The definitions on top are concerned with THOUGHT PROCESSES and REASONING, whereas the ones on the bottom address BEHAVIOR.
- The definitions on the left measure success in terms of fidelity to HUMAN performance, whereas the ones on the right measure against an ideal performance measure, called RATIONALITY.
- A system is rational if it does the “right thing,” given what it knows.


Figure 1.1 Some definitions of artificial intelligence, organized into four categories.

INTRODUCTION → WHAT IS AI? → Acting humanly: The Turing Test approach

- The Turing Test, proposed by Alan Turing (1950), was designed to provide a satisfactory operational definition of intelligence. A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer
- The computer would need to possess the following capabilities:
- **Natural Language Processing** to enable it to communicate successfully in English.
- **Knowledge Representation** to store what it knows or hears;
- **Automated Reasoning** to use the stored information to answer questions and to draw new conclusions;
- **Machine Learning** to adapt to new circumstances and to detect and extrapolate patterns.


INTRODUCTION → WHAT IS AI? → Acting humanly: The Turing Test approach

- Total Turing Test includes a video signal so that the interrogator can test the subject's perceptual abilities, as well as the opportunity for the interrogator to pass physical objects "through the hatch." To pass the total Turing Test, the computer will need
- **Computer Vision** to perceive objects, and
- **Robotics** to manipulate objects and move about.



INTRODUCTION → WHAT IS AI? → Thinking humanly: The cognitive modeling approach

- 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.
- The interdisciplinary field of cognitive science brings together computer models from AI and experimental techniques from psychology to construct precise and testable theories of the human mind.



INTRODUCTION → WHAT IS AI? → Thinking rationally: The “laws of thought” approach


- The Greek philosopher Aristotle was one of the first to attempt to codify “right thinking,” that is, irrefutable reasoning processes. His syllogisms provided patterns for argument structures that always yielded correct conclusions when given correct premises.
- For example, “Socrates is a man; all men are mortal; therefore, Socrates is mortal.”
- These laws of thought were supposed to govern the operation of the mind; their study initiated the field called logic.
- The so-called logicist tradition within artificial intelligence hopes to build on such programs to create intelligent systems.
- 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.

INTRODUCTION → WHAT IS AI? → *Acting rationally: The rational agent approach*

- An agent is just something that acts (agent comes from the Latin *agere*, to do).
- Computer agents are expected to do:
 - operate autonomously,
 - perceive their environment,
 - persist over a prolonged time period,
 - adapt to change,
 - create and pursue goals.
- A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.
- In the “laws of thought” approach to AI, the emphasis was on correct inferences (conclusions).
- Making correct inferences is sometimes part of being a rational agent, because one way to act rationally is to reason logically to the conclusion that a given action will achieve one’s goals and then to act on that conclusion.
- On the other hand, correct inference is not all of rationality; in some situations, there is no provably correct thing to do, but something must still be done.
- There are also ways of acting rationally that cannot be said to involve inference.


THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE

Philosophy	Knowledge Rep., Logic, Foundation of AI (is AI possible?)
Maths	Search, Analysis of search algos, logic
Economics	Expert Systems, Decision Theory, Principles of Rational Behavior
Psychology	Behavioristic insights into AI programs
Neuroscience (Brain Science)	Learning, Neural Nets
Control theory and Cybernetics	Information Theory & AI, Entropy, Robotics
Computer Sc. & Engg.	Systems for AI



INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Philosophy*

- 1. Philosophy (423 BC - present):
- Can formal rules be used to draw valid conclusions?
- How does the mind arise from a physical brain?
- Where does knowledge come from?
- How does knowledge lead to action?
 - Logic, methods of reasoning.
 - Mind as a physical system.
- Foundations of learning, language, and rationality.





INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Mathematics*

- 2. Mathematics (c.800 - present):
 - What are the formal rules to draw valid conclusions?
 - What can be computed?
 - How do we reason with uncertain information?
 - Three areas developed- Logic, Computation and Probability
 - Formal representation and proof
 - Development of Formal logic – 1) Propositional or Boolean logic – 2) Development of First-Order logic by extending the Boolean logic to include objects and relations
 - Algorithms, computation, decidability, tractability First algorithm was developed - Euclid's algorithm for computing GCD
 - Probability – Baye's rule is the underlying approach for uncertain reasoning in AI systems.

INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Economics*


- 3. Economics :
 - How should we make decisions so as to maximize payoff?
 - How should we do this when others may not go along?
 - How should we do this when the pay off may be far in the future?
- **Decision theory**, which combines probability theory with utility theory, provides a formal and complete framework for decisions (economic or otherwise) made under uncertainty.
- **Game theory** included the surprising result that, for some games, a rational agent should adopt policies that are randomized.
- **Operational Research** how to make rational decisions when payoffs from actions are not immediate but instead result from several actions taken in sequence.
- **Satisficing**—making decisions that are “good enough,” rather than laboriously calculating an optimal decision

INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Neuroscience*

- 4. Neuroscience:
- The study of nervous system, particularly the brain.
- The exact way in which the brain enables thought is unknown.
- However, it does enable thought has the evidence.
- “A strong blow to the head can lead to mental incapacitation”
- **Neurons** – Brain consists of nerve cells or neurons.
- There is no theory on how an individual memory is stored.
- A collection of simple cells can lead to thought, action, consciousness that “brains cause minds”
- The parts of a nerve cell or neuron.
- Each neuron consists of a cell body or soma, that contains a cell nucleus.
- Branching out from the cell body are a number of fibers called dendrites and a single long fiber called the axon.
- A neuron makes connections with 10 to 100,000 other neurons at junctions called synapses

INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Neuroscience*

- Signals are propagated from neuron to neuron by a complicated electrochemical reaction.
- The signals control brain activity in the short term and also enable long-term changes in the connectivity of neuron.
- These mechanisms are thought to form the basis for learning in the brain.





INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Psychology*

- 5. Psychology (1879 - present):
- How do humans and animals think and act?
- Adaptation.
- Phenomena of perception and motor control.
- Experimental techniques.

INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Computer Engineering*

- 6. Computer Engineering:
- How can we build an efficient computer?
- For artificial intelligence to succeed, we need two things: intelligence and an artifact.
- The computer has been the artifact of choice.
- Each generation of computer hardware has brought an increase in speed and capacity and a decrease in price
- In 2005 power dissipation problems led to multiply the number of CPU cores rather than the clock speed.
- Current expectations are that future increases in power will come from massive parallelism – Convergence with the properties of the brain.
- The software side of computer science has supplied AI with:
- The **operating systems, programming languages, and tools** needed to write modern programs.

INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Control Theory and Cybernetics*

- 7. Control Theory and Cybernetics:
- How can artifacts operate under their own control?
- Previous Assumption: Only living things could modify their behavior in response to changes in the environment.
- Machines can modify their behavior in response to the environment (sense/action loop)
- Ex: Water-flow regulator, steam engine governor, thermostat.
- The theory of stable feedback systems (1894).
- Build systems that transition from initial state to goal state with minimum energy. In 1950, control theory could only describe linear systems.
- AI largely rose as a response to this shortcoming.



INTRODUCTION → THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE → *Linguistics*

- 8. Linguistics (1957 – present):
- How does language relate to thought?
- Knowledge representation: the study of how to put knowledge into a form that a computer can reason with Grammar.
- Speech demonstrates human intelligence:
- Analysis of human language reveals thought taking place in ways not understood in other settings. Language and thought are believed to be tightly intertwined.
- Modern linguistics and AI intersect in a hybrid field called Computational Linguistics or Natural Language Processing (NLP).

THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE

- **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

THE HISTORY OF ARTIFICIAL INTELLIGENCE

- **The gestation of AI (1943 - 1956):**
- 1943: McCulloch & Pitts: Boolean circuit model of brain.
- 1950: Turing's "Computing Machinery and Intelligence".
- 1956: McCarthy's name "Artificial Intelligence" adopted.
- Why was it necessary for AI to become a separate field ?
- Why couldn't all the work done in AI have taken place under the name of control theory or operations research or decision theory, which have objectives similar to those of AI?
- Why isn't AI a branch of mathematics?
- The first answer is that AI from the start embraced the idea of duplicating human faculties such as creativity, self-improvement, and language use.
- None of the other fields were addressing these issues.
- The second answer is methodology. AI is the only one of these fields that is clearly a branch of computer science (USING COMPUTER SIMULATIONS).
- AI is the only field to attempt to build machines that will function autonomously in complex, changing environments.

THE HISTORY OF ARTIFICIAL INTELLIGENCE

- **Early enthusiasm, great expectations (1952 - 1969):**
- Early successful AI programs:
- Newell & Simon's Logic Theorist.
- Newell & Simon's General Problem Solver (GPS) to imitate human problem solving.
- Solved puzzles and considered sub goals and possible actions similar to which humans approached the same problems.
- Thus, GPS was probably the first program to embody the "thinking humanly" approach.
- Newell & Simon's Physical symbol system hypothesis:
- which states that "a physical symbol system has the necessary and sufficient means for general intelligent action."
- Any system (human or machine) exhibiting intelligence must operate by manipulating data structures composed of symbols.




Figure 1.4 A scene from the blocks world. SHRDLU (Winograd, 1972) has just completed the command "Find a block which is taller than the one you are holding and put it in the box."

INTRODUCTION → THE HISTORY OF ARTIFICIAL INTELLIGENCE

- **A dose of reality (1966 - 1973):**
- The first kind of difficulty arose because most early programs knew nothing of their subject matter.
- They succeeded by means of simple syntactic manipulations.
- Early machine translation efforts to speed up the translation of Russian scientific papers in the wake of the Sputnik launch in 1957.
- It was thought initially that simple syntactic transformations based on the grammars of Russian and English, and word replacement from an electronic dictionary, would suffice to preserve the exact meanings of sentences.
- The fact is that accurate translation requires background knowledge in order to resolve ambiguity and establish the content of the sentence.
- The second kind of difficulty was the intractability of many of the problems.
- Most of the early AI programs solved problems by trying out different combinations of steps until the solution was found.
- This strategy worked initially with limited objects and hence very few possible actions and very short solution sequences.

INTRODUCTION → THE HISTORY OF ARTIFICIAL INTELLIGENCE

- **Knowledge-based systems (1969 - 1979):**
- Uses more powerful, domain-specific knowledge.
- Allows larger reasoning steps and can more easily handle typical expertise.
- Ex1 :1969: DENDRAL by Buchanan et al.
generated all possible structures consistent with the formula of the molecule
- The significance of DENDRAL was that it was the first successful knowledge-intensive system.
- Its expertise derived from large numbers of special-purpose rules.
- Ex2: 1976: MYCIN by Shortliffe.
- Expert systems in medical diagnosis With about 450 rules, MYCIN was better than junior doctors.

INTRODUCTION → THE HISTORY OF ARTIFICIAL INTELLIGENCE

- **AI becomes an industry (1980 - present):**
 - Expert systems industry booms.
 - 1981: Japan's 10-year Fifth Generation project to build intelligent computers.
- **The return of NNs (1986 - present):**
 - Mid 80's: Back-propagation learning algorithm reinvented.
 - 1988: Resurgence of probability.
- **AI adopts the scientific method (1987 - present):**
 - In terms of methodology, AI has finally come firmly under the scientific method.
 - To be accepted, hypotheses must be subjected to rigorous empirical experiments
 - and the results must be analyzed statistically for their importance.
 - It is now possible to replicate experiments by using shared repositories of test data and code.



INTRODUCTION → THE HISTORY OF ARTIFICIAL INTELLIGENCE

- The emergence of intelligent agents (1995- present):
- One of the most important environments for intelligent agents is the Internet.
- AI systems have become so common in Web-based applications that the "-bot" suffix has entered everyday language.
- Internet tools, such as search engines, recommender systems, and Web site aggregators use AI technologies.
- The availability of very large data sets(2001-present):
- In the 60-year history of computer science, the emphasis has been on the algorithm as the main subject of study.
- But recent work in AI suggests that for many problems, it makes more sense to worry about the data.
- This is true because of the increasing availability of very large data sources.
- Examples: trillions of words of English, billions of images from the Web, or billions of base pairs of genomic sequences.

I

INTRODUCTION → THE HISTORY OF ARTIFICIAL INTELLIGENCE

- 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! – long list of Xs and a belief that “ a machine can never do X”
- 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



THE STATE OF THE ART

- **Robotic vehicles:** A driverless robotic car named STANLEY sped through the rough terrain of the Mojave dessert at 22 mph, finishing the 132-mile course first to win the 2005 DARPA Grand Challenge.
- **Speech recognition:** A traveler calling United Airlines to book a flight can have the en_x0002_tire conversation guided by an automated speech recognition and dialog management system
- **Autonomous planning and scheduling:** A hundred million miles from Earth, NASA's Remote Agent program became the first on-board autonomous planning program to control the scheduling of operations for a spacecraft
- **Game playing:** IBM's DEEP BLUE became the first computer program to defeat the world champion in a chess match when it bested Garry Kasparov by a score of 3.5 to 2.5 in an exhibition match

- **Spam fighting:** Each day, learning algorithms classify over a billion messages as spam, saving the recipient from having to waste time deleting what, for many users, could comprise 80% or 90% of all messages, if not classified away by algorithms.
- **Logistics planning:** During the Persian Gulf crisis of 1991, U.S. forces deployed a Dynamic Analysis and Replanning Tool, DART (Cross and Walker, 1994), to do automated logistics planning and scheduling for transportation.
- **Robotics:** The iRobot Corporation has sold over two million Roomba robotic vacuum cleaners for home use.
- **Machine Translation:** A computer program automatically translates from Arabic to English, allowing an English speaker to see the headline “Ardogan Confirms That Turkey Would Not Accept Any Pressure, Urging Them to Recognize Cyprus.”



Discussion and Interaction





Thank
you