

Artificial Intelligence

1. Introduction

Shashi Prabh

School of Engineering and Applied Science
Ahmedabad University

What is Artificial Intelligence (AI)?

Intelligence - Ability to apply knowledge to one's environment, to think abstractly, to learn, to understand or to predict

Artificial Intelligence

1. Understanding and building intelligent entities (machines) that can act effectively and safely in a wide variety of unseen situations
 2. The study of computations that make it possible to perceive, reason and act
- Since AI is relevant to any intellectual task, it is a truly universal field
 - The term was coined by McCarthy (also, Shannon and Minsky) in 1956 in a memo to organize a workshop at Dartmouth College

VOL. LIX. No. 236.]

[October, 1950


MIND
A QUARTERLY REVIEW
OF
PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND
INTELLIGENCE

BY A. M. TURING


1. *The Imitation Game.*

I PROPOSE to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning



What is AI?

- Intelligence can be characterized by
 - Internal **thought process and reasoning** – internal characterization
 - Intelligent **behavior** – external characterization
- AI **as Human-Like System** (Turing test)
 - Psychology, neuroscience, human behaviour
 - Planes do not mimic birds
- AI **as Rational System** i.e., doing the right thing
 - Mathematics, statistics, engineering, control theory, economics
 - Logic and probability are the main tools



AI is about doing the “right thing,” or rational behavior

Rational Agent

- **Agent** [Latin: *agere*, to do] is something that acts
 - AI agents are expected to operate autonomously over a long time
 - Perceive environment, adapt to changes, create & pursue goals
- **Rational agent** acts to achieve the best **expected outcome**
 - Can be defined mathematically
 - Designs can be proven to meet the objectives/specifications
 - Is the “Standard Model” of AI
- In the early days of AI, rational agents were built using logic
 - Methods based on probability theory and machine learning are used to create agents that can make decisions in the presence of uncertainty
 - Maximize the **expected utility**

Human-Friendly Agents

- A major limitation of the rational agent approach is that specifying the objective completely and correctly is usually difficult
- The objectives must be aligned with those of humans
 - A chess playing robot can try to harm the opponent to win
- The standard model is inadequate
 - We don't want AI that pursues its own objective
 - A new formulation is needed that captures the objectives correctly to create agents that are **provably beneficial to humans**



AI2027: Is this how AI might destroy humanity? -




BBC World Service

BBC World Service ✓ 230K views

A research paper predicting that artificial intelligence will go rogue in 2027 and lead to humanity's extinction within a decade is making waves in the tech world. Click here to...

Foundations of AI



History of AI

- **As old as CS:** motivation to build computers was to build machines that can think
 - Calculators already existed
 - **Georgetown-IBM experiment:** “Electronic brain” to translate Russian to English, 1954
- **1940-50 : Early days**
 - 1943 : McCulloch and Pitts – Boolean circuit model of brain
 - 1949 : Hebb – neuron learning rule
 - 1950 : Turing – Computing Machinery and Intelligence


History of AI

- 1950-70 : Initial development, over-optimism
 - 1950s : Samuel's Checker player, Newell and Simon's Logic Theorist, early AI programs
 - 1956 : Dartmouth Meeting, "Artificial Intelligence" coined
 - 1965 : Robinson's complete algorithm for logical reasoning
- 1970-90: Knowledge-based approaches
 - 1969-79 : Early development of KB systems, Dendral (Molecular structure), Mycin (blood infection diagnosis)
 - 1980-88 : Expert system industry boom
 - 1988-93 : Expert system industry bust, AI winter (extending to 2000)

History of AI

- 1988 - 2000 : Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - General increase in technical depth
 - AI spring – agents and learning systems
- 2012 - : Excitement, exponential growth
 - Big data, deep learning, large language models, generative AI
 - Reunification of fields
 - AI adopted in many industries
 - AI reached mass population

AI Winters



Applications of AI


- Learn from the vast corpus of knowledge and data available on the Internet to answer questions and generate new content
 - Education, editing, programming, farming, making people more intelligent, daily life work, ...
- Assist medical professionals with diagnosis, surgery, treatment plan, patient care
- Assist in manufacturing with automated and intelligent assembly line, inspection, maintenance and supply-chain management
- Household jobs, cook and personal assistant
- Running an office, taking impartial and optimal decisions
- Autonomous vehicles
- Endless possibilities!

Course Summary

- This course is about systems that act **rationally**
- **Intelligence from computation**
 - Search-Based Agents, Constraint Satisfaction Problem, Logical Agents (+ Knowledge)
- **Intelligence from data**
 - Bayesian Networks, Hidden Markov Models, Sequential Decision Problems (MDP, POMDP), Multi-Agent Decision Problems, Machine Learning, Reinforcement Learning

Administrative Matters

- Office: 125 GICT
- Office hour: Wed 1:30 PM to 3:30 PM
- Website: shashi-prabh.github.io/teaching
- Textbook: Russell and Norvig, 4th Edition
- **Pre-requisites:** Data Structures, Algorithms, Programming, Probability



Artificial Intelligence

CSE 518, Monsoon 2026

Lectures: Tue and Thu 4:00 PM - 5:30 PM. Room: TBD

Instructor: Shashi Prabh

Office: GICT 125

Office hour: Wed 1:30-3:30 PM, or by appointment

Email: shashi.prabh@ahduni

Prerequisites: Introduction to Computation and Programming (CSE 100), Data Structures and Algorithms (CSE 210), Probability and Random Processes (MAT 202), Linear Algebra (MAT 200 or equivalent). Programming ability is a must!

Course description

Artificial Intelligence (AI) is reshaping nearly every domain of human activity, and the rapid rise of foundation models has made first-hand fluency with both classical and modern AI essential. This course builds that foundation. The syllabus covers intelligent agents, problem solving by search (uninformed, informed, local), constraint satisfaction problems, logic and knowledge representation, probabilistic reasoning with Bayesian networks, probabilistic reasoning over time (Markov chains, Hidden Markov Models, forward and Viterbi algorithms), utility theory and decision networks, Markov decision processes, reinforcement learning, machine learning and neural networks, deep learning, transformers and large language models, and AI ethics, safety and alignment.

Course objectives

- Build an in-depth understanding of the fundamentals of Artificial Intelligence and prepare students for more advanced AI courses.
- Develop the ability to represent knowledge and reason under uncertainty.
- Develop problem-solving skills using AI techniques, and the judgement to differentiate between techniques and apply the right one to a given problem.
- Build hands-on competence with widely used AI tools and libraries, and the ability to design and implement (in Python) efficient autonomous agents — including LLM-augmented agents built with modern tooling.
- Develop the ability to evaluate AI applications and their limitations, including the ethical and societal implications of modern AI systems.


Session	Topic	Topic details	Readings	Slides	Exercises
1	Introduction	What is AI?, goals of AI, AI applications, foundations of AI, history of AI	1		• Install AIMA Python package • Problem Set 1
2	Intelligent Agents	Agents and Environment, Nature of Environment, Structure of Agent, Rationality, Performance Measures	2		• Problem Set 2 • Notebook & exercises
3, 4	Problem Solving Agents and Search	Problem Solving Agents, Best-First Search, Uninformed Search: Breadth First Search, Depth First Search, Iterative Deepening DFS	3.1-3.4		• Problem Set 3 • Notebook • Programming exercises
5	Informed Search	Heuristic Functions, Greedy Best First search, A* Search	3.5-3.6		
6	Search in Complex Environments	Local Search, Search with Non-deterministic Actions	4.1, 4.3-4.5		

Administrative Matters

- Evaluation
 - Mid-term exam: 35%
 - Final exam: 35%
 - Quizzes (2): 10%
 - Assignments: 10%
 - Project: 10%
- Relative grading
- No make-up
 - Scaling possible in emergency or for important family events
 - Supporting documents must be provided on the final exam day

Project

- Default topic or a topic of your choice
 - Approach me for approval and feedback
- Due 2 weeks before the last day of classes
 - Start early



What this course is

- This course is about fundamentals of AI
- This is a mathematics and programming intensive technical course
- Provides foundational knowledge of AI
 - Prepares you for learning more advanced topics
- Gives hands-on experience with building small AI systems

What this course is not

- This is **NOT** a course that
 - Teaches AI related policy or detailed discussions on ethical implications
 - Try SAS or Management School courses
 - Teaches you AI fast
 - Makes you an LLM expert in a semester



Learn AI in Just 3 HOURS 🚀 | ChatGPT & Generative AI |
#shorts
670K views • 1 year ago
Ishan Sharma ✓
BEST FREE AI Course For EVERYONE | Ishan Sharma Instagram: <https://bit.ly/ishans>
Discord ...

Become an
AI ENGINEER
in 30 days



How To Learn AI So FAST It Feels ILLEGAL

33K views • 2 months ago

Tech With Soleyman

How To Learn AI So FAST It Feels ILLEGAL The AI revolution is here, and it's creating 2 types of people; those who..

4K

4 chapters Intro | Quick Money | Paths | Build Expertise

Helpful Suggestions

- Take notes! Remain attentive!
 - Learn to take notes efficiently: search the web
- Clear your doubts during the lecture itself
 - Do not hesitate to ask questions
- Before coming to sessions
 - Review your notes
 - Scan the portion of the textbook to be covered
- Do exercises on your own
- Don't hesitate to approach me if you need help!

Expectations from students

- Spend 10-12 hours per week
- Remain up-to-date
 - Course material will get difficult with time
 - Catching-up will get difficult
- If absent, find out missed content and announcements from other students
- Noise and distractions policy
 - Come to lectures on time
 - Use of cellphone, laptop or tablet during lectures is not permitted.
 - Keep cellphones in silent mode in your bag.
- Academic integrity must be maintained

Quiz

- 1 **[Rationality]** What does it mean for an agent to be rational?
 - A) It always knows the outcomes regardless of the environment.
 - B) It acts to maximize its expected performance based on its knowledge and abilities.
 - C) It always selects the action that is the fastest to execute.

- 2 **[AI & Human Performance]** What best characterizes the goal of AI?
 - A) AI mimics human behavior.
 - B) AI seeks to exceed human performance.
 - C) AI focuses on building systems that act intelligently, whether human-like or not.

Quiz

1 [Rationality] What does it mean for an agent to be rational?

- A) It always knows the outcomes regardless of the environment.
- B) It acts to maximize its expected performance based on its knowledge and abilities. [Correct]
- C) It always selects the action that is the fastest to execute.

2 [AI & Human Performance] What best characterizes the goal of AI?

- A) AI mimics human behavior.
- B) AI seeks to exceed human performance.
- C) AI focuses on building systems that act intelligently, whether human-like or not. [Correct]

Summary

- Artificial Intelligence is the study of designing intelligent agents
- AI can be categorized as systems that think like humans, act like humans, think rationally, or act rationally.
 - Doing the “right” thing, or, rational action is the practical approach that is also suitable for mathematical treatment.
- A rational agent selects actions that maximize its expected performance based on its percept sequence and knowledge, operating in diverse environments, which may be partially observable or uncertain.
 - Specifying objectives correctly and completely is a limitation of this approach which can lead to behavior that is not aligned with humans
- The foundations of AI draw from a very diverse set of disciplines
- AI progress has gone through AI Winters

Activities

1. Brush-up Python programming skills
 2. Install AIMA Version 4 package
 3. Work on PS 1
 4. Reading assignment: Chapters 1 and 2 (scan)
- Next lecture: Intelligent Agents [Chapter 2]