Introduction to Artificial Intelligence

Unit 4: Concepts, Goals, History, Applications, Rational Agents & Perspectives

Dr. Gopal Chandra Jana

Dept. of CSE, Sharda University

Fundamentals of Computing and AI (CSAI1101) — Unit 4

September 3, 2026

Course Page:
https://www.gcjana.in/courses/shardauniversity/2601/fundamentals-of-computing-and-artificial-intelligence/

Outline

Learning Outcomes of Unit 4

After this unit, a student will be able to...

1. **Explain** what Artificial Intelligence is, its goals and its guiding principles.
2. **Trace** the historical evolution and key milestones of AI.
3. **Describe** the core domains of AI and a learning pathway to become an AI developer.
4. **Identify** real-world applications of AI across diverse fields.
5. **Distinguish** AI-based problem solving from traditional methods, and explain rational agents and intelligent behaviour.
6. **Discuss** the philosophical and practical perspectives — and the ethics — of AI systems.

Mapping to the Syllabus

A. Concepts, goals & history. **B.** AI curriculum, learning pathway & applications. **C.** AI vs traditional problem solving, rational agents & perspectives.

From Units 1–3 to Unit 4

Where we are

So far we learnt *how computers work* (Unit 1), *how to think about problems* (Unit 2) and *how to design and test logic* (Unit 3). Now we ask a bigger question:

Can we make machines intelligent?

- This unit introduces the **ideas** of AI — gently, no heavy maths.
- It sets the stage for **Unit 5: Machine Learning**.
- You are training to be a **B.Tech CSE (AI & ML)** engineer — this is your field!

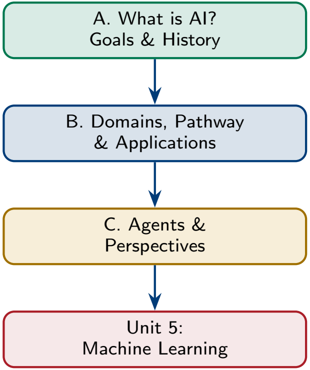

Figure: The road through Unit 4 and beyond.

What is Artificial Intelligence?

Definition

Artificial Intelligence (AI) is the branch of computer science that builds machines and programs able to perform tasks that normally require **human intelligence** — such as learning, reasoning, understanding language and recognising patterns.

- *Artificial* = made by humans.
- *Intelligence* = the ability to learn, reason and solve problems.
- John McCarthy (who coined the term) called it *"the science and engineering of making intelligent machines."*

AI is already around you

Voice assistants, map navigation, spam filters, film recommendations, face unlock — all use AI.

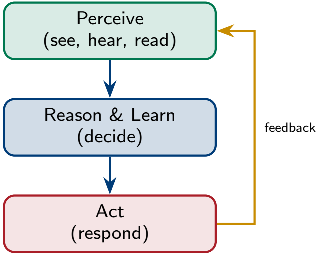

Figure: An intelligent system perceives, reasons and acts.

Four Ways to Define AI (Russell & Norvig)

Two questions × two standards

Should a machine *think* or *act*? And should it be measured against *humans* or against ideal *rationality*? This gives four views of AI.

	Like a Human	**Rationally (ideally)**
Thinking	Thinking humanly (model the human mind)	Thinking rationally (logic, correct reasoning)
Acting	Acting humanly (the Turing Test)	**Acting rationally** (rational agents — our focus)

Modern AI mostly follows the **acting rationally** view: build agents that *do the right thing*.

The Turing Test — Can a Machine Fool a Human?

Alan Turing's "Imitation Game" (1950)

A human **judge** chats (by text) with a hidden human and a hidden machine. If the judge cannot reliably tell which is the machine, the machine is said to show **intelligent behaviour**.

- It focuses on *behaviour*, not inner workings.
- "Can machines think?" becomes "Can machines *act* intelligently?"

A famous criticism

Fooling a judge is not the same as truly *understanding* — an objection we revisit with the *Chinese Room* later.

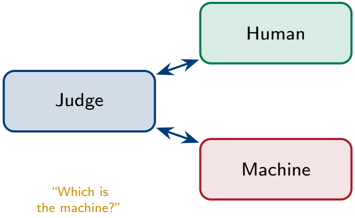

Figure: The Turing Test set-up.

Goals and Objectives of AI

What are we trying to achieve?

- Build systems that **learn** from data and experience.
- **Reason** and make good decisions.
- **Understand** language, images and speech.
- **Automate** difficult or repetitive tasks.
- Solve problems that have **no clear formula**.
- Assist humans — act as a helpful partner.

Objectives of an intelligent system

A good intelligent system should:

- **Perceive** its environment (sensors, data).
- **Reason** about what it perceives.
- **Learn** and improve over time.
- **Act** to achieve its goal.
- Handle **uncertainty** and incomplete information.

The dream: AGI

Artificial General Intelligence — a machine as broadly capable as a human — is still a research goal, not today's reality.

Types of AI — By Capability

Three levels

- **Narrow AI (ANI)**: good at *one* task. This is **all AI today** — chess, translation, face unlock.
- **General AI (AGI)**: human-level ability across *many* tasks. Still a *future* goal.
- **Super AI (ASI)**: hypothetically *beyond* human ability. Only theory / science fiction.

Reality check

Every AI you use — however clever — is **Narrow AI**. There is no general or super AI yet.

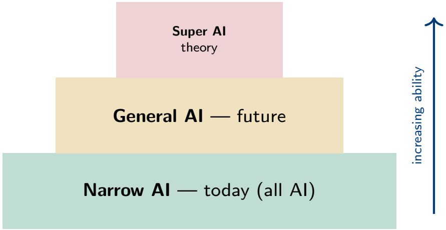

Figure: Narrow → General → Super AI.

Types of AI — By Functionality

Type	**What it does**	**Example**
Reactive machines	React to the present; no memory	Deep Blue (chess)
Limited memory	Use recent past data to decide	Self-driving cars
Theory of mind	Understand emotions/beliefs (research)	Social robots (future)
Self-aware	Conscious of itself (only theory)	Not yet built

Where we are

Today's AI is mostly **reactive** and **limited-memory**.

Still future

Theory of mind and **self-aware** AI do not exist yet — they are active research and philosophy.

AI, Machine Learning & Deep Learning

Nested circles — not the same thing

- **AI**: the big idea — any machine showing intelligence.
- **Machine Learning (ML)**: a *part* of AI where machines **learn from data** instead of being fully hand-coded.
- **Deep Learning (DL)**: a *part* of ML using *neural networks* with many layers.

Unit 5 preview

We study **Machine Learning** in detail in Unit 5 — the engine behind most modern AI.

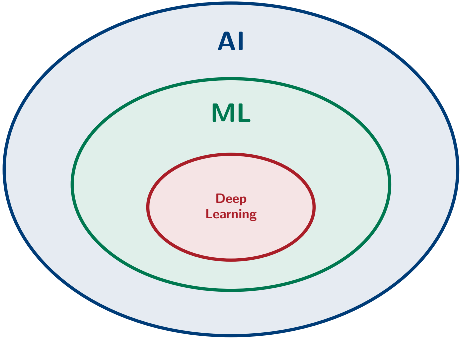

Figure: DL $\subset$ ML $\subset$ AI.

What Counts as "Intelligent"?

Signs of intelligent behaviour

We usually call behaviour *intelligent* if it involves:

- **Learning** from experience.
- **Reasoning** to reach conclusions.
- **Problem solving** in new situations.
- **Perception** — making sense of inputs.
- **Language** — understanding and communicating.
- **Adaptation** to a changing world.

The "AI effect"

Once a task is solved, people often say "that's not *real* AI"! (Calculators, spam filters were once "AI".) The goalposts keep moving.

Human vs machine strengths

Humans are better at: common sense, creativity, emotion, few-example learning.

Machines are better at: speed, memory, tireless repetition, spotting patterns in huge data.

The goal

The best results often come from **humans and AI working together**, each doing what it does best.

Common Myths about AI

Myth	Reality
“AI will take over the world soon.”	Today’s AI is **narrow** — no world takeover.
“AI understands like a human.”	It matches patterns; it does not truly **understand**.
“AI is always correct/objective.”	It can be **biased** if its data is biased.
“AI will replace all jobs.”	It **changes** jobs and creates new ones.
“Only geniuses can learn AI.”	Anyone willing to **practise** can learn AI!

Separate the science-fiction from the science — as future AI engineers, you should know the difference.

The Birth of AI — Foundations

Before AI had a name

- **1943**: McCulloch & Pitts model the first artificial *neuron*.
- **1950**: Alan Turing's paper *"Computing Machinery and Intelligence"* proposes the Turing Test.
- **1956**: The **Dartmouth Conference** — John McCarthy coins the term *"Artificial Intelligence"*. AI is born as a field.

The founders

McCarthy, Minsky, Rochester and Shannon organised Dartmouth. McCarthy is called the **"Father of AI"**.

Early programs

- **1955–56**: *Logic Theorist* (Newell & Simon) — the first AI program; proved maths theorems.
- **1957**: *Perceptron* (Rosenblatt) — an early learning neural network.
- **1958**: McCarthy creates *LISP*, the classic AI language.
- **1966**: *ELIZA* — an early chatbot that imitated a therapist.

Milestones in AI — A Timeline

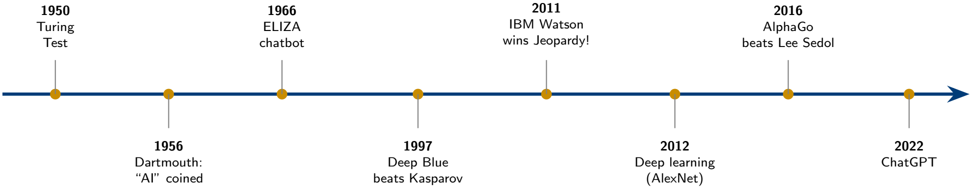

The pattern

AI has grown in waves — bursts of progress, then quieter periods. Since **2012**, **deep learning** and big data have driven rapid, world-changing advances.

AI Winters — When Progress Paused

Hype, then disappointment

An **AI winter** is a period when funding and interest fell because AI *over-promised* and *under-delivered*.

- **1974–1980**: early methods hit their limits; funding cut.
- **1987–1993**: expensive "expert systems" disappointed.

Each winter followed a wave of *too-high expectations*.

The lesson

Ideas often arrive *before* the computing power and data needed to make them work. Patience and honesty about limits matter.

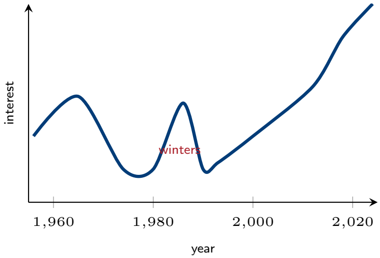

Figure: AI interest — with two "winter" dips.

Quick Check — Round 1 (Concepts & History)

[Q1] Who coined the term "Artificial Intelligence", and in which year and event?

[Q2] What does the Turing Test measure?

[Q3] Put in order (smallest to largest): AI, Deep Learning, Machine Learning.

[Q4] Which type of AI do we have today: Narrow, General or Super?

Think for 60 seconds — answers on the next slide.

Answers — Round 1

[A1]

John McCarthy, in **1956** at the **Dartmouth Conference**.

[A2]

Whether a machine can **behave** so intelligently in conversation that a human judge cannot tell it apart from a human.

[A3]

Deep Learning $\subset$ **Machine Learning** $\subset$ **AI**.

[A4]

Narrow AI — every AI system today is good at a specific task only.

The Core Domains of AI

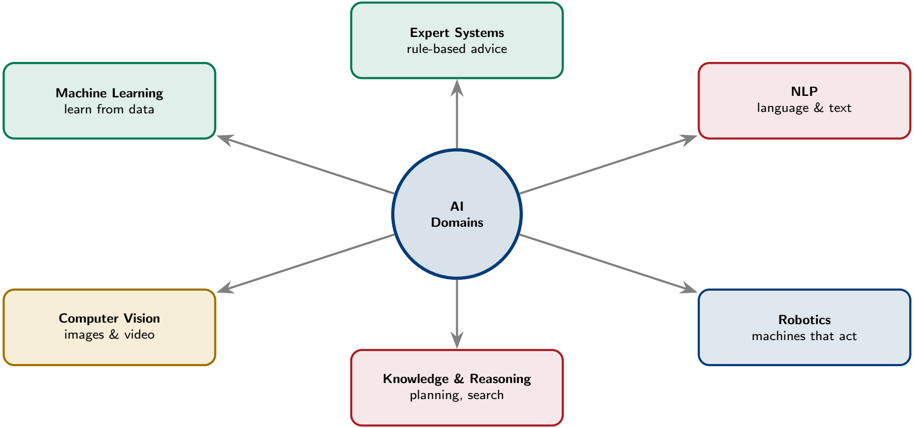

AI is a family of connected fields — most real systems combine several.

Core Domains — Explained

Learning & reasoning

- **Machine Learning**: learn patterns from data (Unit 5).
- **Knowledge Representation & Reasoning**: store facts and reason logically (search, planning).
- **Expert Systems**: capture a human expert's rules to give advice (e.g. medical diagnosis).

Perceiving & acting

- **Natural Language Processing (NLP)**: understand and generate human language (translation, chatbots).
- **Computer Vision**: understand images and video (face unlock, medical scans).
- **Robotics**: intelligent machines that sense and act in the real world.

They overlap

A self-driving car uses *vision* + *ML* + *robotics* + *reasoning* together.

Overview of the AI Curriculum (This Course)

What a foundation in AI covers

A well-rounded AI education builds on:

- **Maths**: linear algebra, probability, statistics, calculus.
- **Programming**: especially *Python*.
- **Data handling**: cleaning, visualising, analysing data.
- **Core AI**: search, knowledge, reasoning.
- **Machine & Deep Learning**.
- **Ethics** and responsible AI.

Where this course fits

Fundamentals of Computing and AI gives you the **foundation**:

- Units 1–3: computing, problem solving, logic.
- Unit 4: what AI is (you are here).
- Unit 5: an introduction to Machine Learning.

Later semesters go deeper into each area.

A Learning Pathway to Become an AI Developer

1. Foundations maths + Python → **2. Data Skills** analyse & plot → **3. Core AI** search, logic → **4. ML / DL** train models → **5. Projects** build & deploy

Golden advice

Build projects! You learn AI by *doing* — small projects teach more than watching videos.

Habits of good AI learners

Stay curious, practise coding daily, read about new tools, and always question your data and results.

Tools & Languages You Will Meet

Tool / Library	**Used for**
`Python`	The main language of AI — simple and powerful
`NumPy`, `Pandas`	Numbers, tables and data handling
`Matplotlib`	Plotting and visualising data
`scikit-learn`	Classic machine-learning models
`TensorFlow`, `PyTorch`	Deep learning / neural networks
`Jupyter Notebook`	Writing and running code interactively

You do not need all of these at once — start with **Python**, and add tools as you grow.

Real-World Applications — Everyday Life

AI in your pocket and home

- **Voice assistants**: Alexa, Siri, Google Assistant.
- **Recommendations**: YouTube, Netflix, Amazon, Spotify.
- **Maps & navigation**: live traffic, best routes.
- **Face unlock** and photo tagging.
- **Spam filters** and smart email replies.
- **Chatbots** for customer support.

Why AI here?

These tasks have **no fixed formula** and involve messy, real-world data — exactly where AI shines and traditional programming struggles.

Think about it

Count how many AI-powered features you used *today* before coming to class. The number may surprise you!

Real-World Applications — Across Industries

Health, finance, transport

- **Healthcare**: disease detection from scans, drug discovery, virtual health assistants.
- **Finance**: fraud detection, credit scoring, algorithmic trading.
- **Transport**: self-driving cars, traffic control, route optimisation.
- **Agriculture**: crop & pest detection, yield prediction.

Education, industry, more

- **Education**: personalised learning, auto-grading.
- **Manufacturing**: quality inspection, predictive maintenance, robots.
- **Retail**: demand forecasting, smart pricing.
- **Entertainment**: game AI, image & music generation.
- **Security**: face recognition, threat detection.

AI is a "horizontal" technology — it touches every industry.

Application Spotlight — How AI Sees the World

Example: recognising a handwritten digit

A camera captures a digit as a grid of pixels. An AI model outputs a **confidence** for each possible answer.

- Input: 28×28 pixel image.
- Output: probabilities, e.g. $P(2) = 0.93$, $P(7) = 0.05$.
- No human wrote an exact rule — the model **learned** it from examples.

The big idea

AI turns a hard "rule" problem into a **learning** problem — show many examples, let the machine find the pattern.

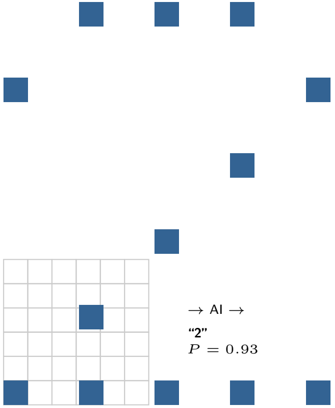

Figure: Pixels in, a confident answer out.

Spotlight — Natural Language Processing (NLP)

Teaching machines human language

NLP lets computers read, understand and produce text and speech.

- **Translation**: English $\leftrightarrow$ Hindi.
- **Chatbots** & virtual assistants.
- **Sentiment analysis**: is a review positive or negative?
- **Speech-to-text** and text-to-speech.
- **Summarisation** of long documents.

Why it is hard

Human language is full of *ambiguity*: "I saw the man with the telescope" has two meanings!

Spotlight — Computer Vision

Computer Vision lets computers understand images and video.

- **Image classification**: cat vs dog.
- **Object detection**: find cars in a photo.
- **Face recognition** and unlock.
- **Medical imaging**: spot tumours in scans.
- **OCR**: read text from images.

Together

A phone that reads a sign aloud in your language uses *vision* (read) + *NLP* (translate & speak).

Application Spotlight — AI in India

AI serving a billion+ people

- **Agriculture**: crop advice and pest alerts for farmers.
- **Healthcare**: AI screening for eye disease and TB in rural clinics.
- **Language**: translation across many Indian languages.
- **Governance**: document reading, chatbots for citizen services.
- **Fintech**: fraud detection for digital payments (UPI).

A growing field at home

India has a fast-growing AI ecosystem — national AI programmes, research labs, start-ups and supercomputers (e.g. *PARAM*, *AIRAWAT*).

Opportunity for you

Solving *local* problems with AI — in your own languages and communities — is a huge opportunity for the next generation of Indian AI engineers.

Quick Check — Round 2 (Domains & Applications)

[Q1] Which AI domain deals with understanding *language*? Which deals with *images*?

[Q2] Name the main programming language used in AI.

[Q3] Give one real-world AI application in healthcare and one in finance.

[Q4] Why is AI a good fit for tasks like face recognition but not for computing a bank EMI?

Answers — Round 2

[A1]

Language → **Natural Language Processing (NLP)**; images → **Computer Vision**.

[A2]

Python — simple, powerful, with many AI libraries.

[A3]

Healthcare: **disease detection from scans**; Finance: **fraud detection** (any valid example).

[A4]

Face recognition has *no exact formula* and messy data — ideal for AI. An EMI has an *exact formula* needing an *exact* answer — best done by traditional computing.

AI-Based vs Traditional Problem Solving

Traditional programming

The programmer writes **exact rules**. The computer follows them precisely.

Data + Rules → Answers

- Great when a clear formula exists.
- Same input always gives same output.

AI-based approach

The machine **learns the rules** from examples.

Data + Answers → Rules (model)

- Great when *no* clear formula exists.
- Improves as it sees more data.

Key flip: traditional coding is given the *rules*; AI is given the *answers* and finds the rules itself.

Traditional vs AI — Side by Side

Aspect	**Traditional**	**AI-based**
Rules	Written by the programmer	Learned from data
Best for	Exact, well-defined problems	Vague, pattern-rich problems
Data	Not essential	Essential (lots of it)
Output	Exact, repeatable	Approximate, probabilistic
Adaptivity	Fixed until re-coded	Improves with more data
Example	Calculating an EMI, sorting	Face recognition, translation

They complement each other — a real app often uses both: exact code *and* learned models.

AI as Search — Solving Problems by Exploring

Many AI problems are search problems

The AI explores possible **states** to reach a **goal**:

- **State**: a situation (e.g. a board position).
- **Actions**: legal moves from a state.
- **Goal test**: have we reached the target?
- **Path**: the sequence of actions to the goal.

Examples

Solving a maze, planning a route on a map, playing chess — all are *searches* through a space of possibilities.

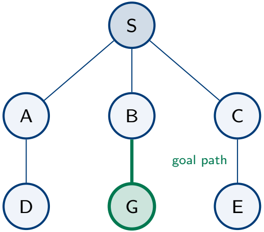

Figure: Searching a tree of states from Start (S) to Goal (G).

What is an Agent?

Definition

An **agent** is anything that **perceives** its environment through *sensors* and **acts** on it through *actuators*.

- **Percept**: what the agent senses.
- **Action**: what the agent does.
- **Agent function**: maps percepts to actions.

Examples

- *Human*: eyes/ears (sensors), hands/legs (actuators).
- *Robot*: cameras (sensors), motors (actuators).
- *Software agent*: reads data (sensors), shows output (actuators).

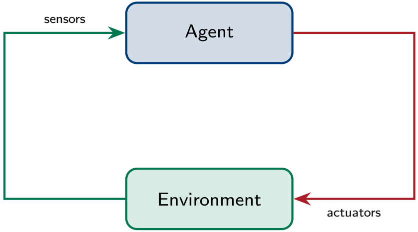

Figure: An agent senses and acts on its environment.

Rational Agents — Doing the Right Thing

Definition

A **rational agent** is one that acts to **achieve the best expected outcome** — "the right thing" — given what it knows and perceives.

- It has a **goal** or a **performance measure**.
- It chooses actions that maximise success.
- Rational $\neq$ perfect — it does the best it *can* with the information available.

Example

A chess agent's performance measure is *winning*. A rational move is the one most likely to lead to a win.

Rational vs. omniscient

A rational agent is **not** all-knowing. If a safe action leads to a bad outcome due to bad luck it could not foresee, the action was still rational.

PEAS

We describe an agent's task by **PEAS**: *Performance, Environment, Actuators, Sensors*.

Types of Intelligent Agents

Agent type	**How it decides**
Simple reflex	Acts only on the *current* percept (if–then rules)
Model-based reflex	Keeps an internal *model* of the world
Goal-based	Chooses actions that reach a *goal*
Utility-based	Prefers actions with the best *utility* (happiness score)
Learning agent	*Improves* its behaviour from experience

Simplest

A **simple reflex** agent: "IF obstacle THEN turn" — like a basic cleaning robot.

Most powerful

A **learning** agent gets better over time — the basis of modern AI.

Describing an Agent with PEAS — Worked Example

PEAS = Performance, Environment, Actuators, Sensors

Before building an agent, we describe its task clearly using PEAS.

PEAS	**Self-driving car**	**Vacuum-cleaning robot**
Performance	Safety, speed, comfort, legal	Cleanliness, battery use, time
Environment	Roads, traffic, pedestrians	Rooms, floor, furniture
Actuators	Steering, brake, accelerator	Wheels, suction motor
Sensors	Cameras, radar, GPS	Bump & dirt sensors

Try it: write the PEAS for a *chess-playing agent* or a *spam filter*.

Environments — Where Agents Live

Environments differ

The task is harder or easier depending on the environment:

- **Fully** vs **partially** observable.
- **Deterministic** vs **stochastic** (random).
- **Static** vs **dynamic** (changing).
- **Discrete** vs **continuous**.
- **Single-agent** vs **multi-agent**.

Examples

- *Chess*: fully observable, deterministic, multi-agent.
- *Self-driving*: partially observable, stochastic, dynamic, continuous.

Why it matters

The messier the environment, the more **intelligence** (and data) the agent needs.

Strong AI vs Weak AI

Weak AI (Narrow AI)

The machine *appears* intelligent and performs a specific task well — but it does not truly **understand** or have a mind.

"It acts intelligent."

This is all AI today.

Strong AI

The (hypothetical) claim that a machine could genuinely *think*, *understand* and be *conscious* — like a human mind.

"It really is intelligent."

Not achieved — and philosophically debated.

Key question: does *behaving* intelligently mean *truly* understanding? Philosophers disagree.

The Chinese Room — A Famous Thought Experiment

Searle's argument (1980)

Imagine a person in a room who knows *no* Chinese. They receive Chinese notes and use a **rule book** to send back correct Chinese replies. To those outside, the room "speaks Chinese".

- The person still does *not understand* Chinese.
- Searle argues a computer is the same: it *manipulates symbols* without real understanding.

The debate

It challenges *Strong AI*: passing a test (like Turing's) may not prove real understanding. Many counter-arguments exist — it is still debated today.

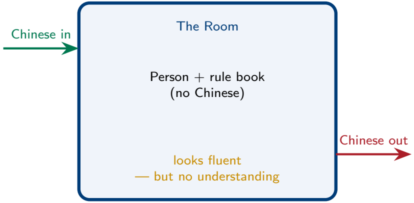

Figure: Symbols in, symbols out — but no meaning.

Practical Perspective — Benefits & Limitations

Benefits of AI

- Works **24/7** without tiring.
- Handles **huge data** quickly.
- Reduces human error in routine tasks.
- Finds patterns humans would miss.
- Assists doctors, teachers, scientists.

Limitations of AI

- **No common sense** or real understanding.
- Only as good as its **data** (biased data $\rightarrow$ biased AI).
- Can be a **black box** (hard to explain).
- Needs lots of data and computing power.
- Cannot feel or truly care.

Ethics & Responsible AI

Because AI affects real lives

As future AI engineers, you must build AI **responsibly**. Key concerns:

- **Bias & fairness**: AI must not discriminate.
- **Privacy**: protect people's data.
- **Transparency**: be able to explain decisions.
- **Accountability**: humans stay responsible.
- **Safety**: AI must not cause harm.
- **Jobs**: manage the impact on work.

A guiding principle

AI should be **human-centred** — designed to *help* people, with humans always in control of important decisions.

Your responsibility

"Because we *can* build it" is not the same as "we *should*". Good engineers ask both questions.

The Future of AI & Careers in AI

Where AI is heading

- **Generative AI**: creating text, images, code, music.
- **Explainable AI**: models that can justify their decisions.
- **Edge AI**: intelligence running on small devices.
- **AI + robotics**: smarter machines in the real world.
- **Responsible AI**: fairness, safety, regulation.

Careers you can aim for

- **ML / AI Engineer**
- **Data Scientist / Analyst**
- **Computer Vision / NLP Engineer**
- **Robotics Engineer**
- **AI Researcher**
- **AI Product Manager**

Your advantage

You are starting AI in your *first* semester — a strong head start for all of these paths.

Quick Check — Round 3 (Agents & Perspectives)

[Q1] In one line, how does AI-based problem solving differ from traditional programming?

[Q2] Define a rational agent.

[Q3] What is the difference between Weak AI and Strong AI?

[Q4] What does the Chinese Room argument try to show?

Think for 60 seconds — answers on the next slide.

Answers — Round 3

[A1]

Traditional: the programmer *writes the rules*. AI: the machine *learns the rules from data* (data + answers → rules).

[A2]

A **rational agent** acts to achieve the *best expected outcome* given its goal and what it perceives — it "does the right thing".

[A3]

Weak AI only *appears* intelligent at a task; **Strong AI** would *truly* think and understand (not yet achieved).

[A4]

That manipulating symbols correctly (passing a test) does *not* prove real **understanding** — a challenge to Strong AI.

Summary of Unit 4

What we covered

1. **What AI is** — machines that learn, reason and act; the four views, the Turing Test, and AI $\supset$ ML $\supset$ DL.
2. **Goals & types** — narrow, general, super; reactive to self-aware; the objectives of intelligent systems.
3. **History** — from 1943/1950/1956 through the AI winters to the deep-learning boom since 2012.
4. **Domains & pathway** — ML, NLP, vision, robotics, reasoning; a step-by-step path to becoming an AI developer.
5. **Applications** — everyday life and every industry.
6. **Agents & perspectives** — AI vs traditional problem solving, rational agents, Strong vs Weak AI, the Chinese Room, and ethics.

AI gives machines the ability to perceive, reason, learn and act — responsibly, with humans in control.

Key Terms to Remember

Concepts & history

- **AI / ML / DL** nested fields
- **Turing Test** = behavioural test
- **Dartmouth 1956**, McCarthy = father of AI
- **AI winter** = a funding slump
- **Narrow / General / Super** AI

Agents & perspectives

- **Agent** = sensors + actuators
- **Rational agent** = does the right thing
- **PEAS** description
- **Weak vs Strong AI**
- **Chinese Room** = symbols $\neq$ understanding

Practice Questions (Exam Oriented)

Short answer (2–5 marks)

1. Define AI. State two goals of intelligent systems.
2. What is the Turing Test? What does it measure?
3. Differentiate Narrow, General and Super AI.
4. Name and briefly explain any four core domains of AI.
5. What is a rational agent? Give one example.

Long answer (8–10 marks)

6. Trace the historical evolution of AI with at least six milestones.
7. Compare AI-based and traditional problem solving on at least five points, with examples.
8. Explain the types of intelligent agents with an example of each.
9. Describe five real-world applications of AI across different industries.
10. Explain Strong vs Weak AI and the Chinese Room argument. What does it tell us about machine understanding?

Think & Explore (Take-Home)

Mini tasks

- List 10 AI features you use in a week; name the domain of each (vision, NLP, etc.).
- Pick one industry and describe two ways AI helps it.
- Write a PEAS description for a self-driving car.
- Debate: "Can a machine ever truly understand?" — write your view in five lines.

Reading & reflection

- Read a short summary of Turing's 1950 paper.
- Try a free AI tool (an image classifier or chatbot) and note one thing it does well and one it gets wrong.
- Prepare for Unit 5: *Introduction to Machine Learning*.

Reminder

Post any doubts on the course page. Stay curious — AI is the field you are training to lead.

References

Textbooks (as per the course page)

1. S. Russell & P. Norvig, *Artificial Intelligence: A Modern Approach*, Pearson. **[Main text]**
2. E. Rich, K. Knight & S. B. Nair, *Artificial Intelligence*, Tata McGraw Hill.
3. Peter Norton, *Introduction to Computers*, Tata McGraw Hill.
4. D. W. Patterson, *Introduction to Artificial Intelligence and Expert Systems*, PHI.

Seminal papers & further reading

5. A. M. Turing, "Computing Machinery and Intelligence," *Mind*, 1950.
6. J. McCarthy et al., *A Proposal for the Dartmouth Summer Research Project on AI*, 1955.
7. J. Searle, "Minds, Brains, and Programs," 1980 (the Chinese Room).

Course page: `https://www.gcjana.in/courses/shardauniversity/2601/fundamentals-of-computing-and-artificial-intelligence/`

Thank You

Any Questions?

"Artificial intelligence is the new electricity."

— Andrew Ng

Post your doubts on the course page $\rightarrow$
https://www.gcjana.in/courses/shardauniversity/2601/fundamentals-of-computing-and-artificial-intelligence/#Post-doubts