An Introductory Module on

Computational Thinking and Artificial Intelligence

for

Secondary Stage (Grades 9 – 12)



© NCERT not to be republished

DIKSHA

Download DIKSHA app from Google Playstore and follow the steps given below and access the e-Resources through your smartphone or tablet using DIKSHA



1. Select preferred language
2. Choose your role: Student or Teacher
3. Grant access and allow app permissions
4. Tap to scan the QR code
5. Focus camera on the QR code in textbook
6. Click to Play QR code specific e-resource(s)

For accessing the e-Resources using DIKSHA on desktop or laptop follow the step stated below:

Go to https://diksha.gov.in/resources and enter the alphanumeric code given under the QR code

© NCERT not to be republished

LiFE
Lifestyle for Environment



An Introductory Module on

Computational Thinking and Artificial Intelligence

for

the Secondary Stage (Grades 9 –12)



UN514

© NCERT not to be republished

विद्यया S मृतमश्नुते

एन सी ई आर टी
NCERT

राष्ट्रीय शैक्षिक अनुसंधान और प्रशिक्षण परिषद्
NATIONAL COUNCIL OF EDUCATIONAL RESEARCH AND TRAINING

UN514 – An Introductory Module on Computational Thinking and Artificial Intelligence

978-93-5729-286-3

First Edition
October 2026 *Ashwina* 1948

PD XX SM

ALL RIGHTS RESERVED

- No part of this publication may be reproduced, stored in a retrieval system or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior permission of the publisher.
- This book is sold subject to the condition that it shall not, by way of trade, be lent, re-sold, hired out or otherwise disposed of without the publisher's consent, in any form of binding or cover other than that in which it is published.
- The correct price of this publication is the price printed on this page, Any revised price indicated by a rubber stamp or by a sticker or by any other means is incorrect and should be unacceptable.

© National Council of Educational Research and Training, 2026

OFFICES OF THE PUBLICATION DIVISION, NCERT

NCERT Campus
Sri Aurobindo Marg
New Delhi 110 016 **Phone : 011-26562708**

108, 100 Feet Road
Hosdakere Halli Extension
Banashankari III Stage
Bengaluru 560 085 **Phone : 080-26725740**

Navjivan Trust Building
P.O.Navjivan
Ahmedabad 380 014 **Phone : 079-27541446**

CWC Campus
Opp. Dhankal Bus Stop
Panihati
Kolkata 700 114 **Phone : 033-25530454**

CWC Complex
Maligaon
Guwahati 781 021 **Phone : 0361-2674869**

Publication Team

Head (In charge), Publication Division : Amitabh Kumar

Chief Editor : Bijnan Sutar

Chief Business Manager (In charge) : Bhoopendra Singh

Chief Production Officer (In charge) : Deepak Jaiswal

Editor : Shilpa Mohan

Printed on 80 GSM paper with NCERT watermark

Published at the Publication Division by the Secretary, National Council of Educational Research and Training, Sri Aurobindo Marg, New Delhi 110 016 and printed at ...

© NCERT not to be republished

Foreword

The National Education Policy (NEP) 2020 envisions an education system that prepares learners to participate meaningfully in a world shaped by rapid scientific and technological change, while remaining firmly grounded in constitutional values, ethical responsibility, and India's rich intellectual traditions. It emphasises critical thinking, creativity, problem-solving, digital literacy, and the ability to apply knowledge in real-life situations. The National Curriculum Framework for School Education (NCF-SE) 2023 gives curricular expression to this vision through learning that is experiential, inquiry-based, multidisciplinary, and connected with the lives of students.

The Secondary Stage is an important period in the intellectual and personal growth of learners. Their capacity for abstract thinking, reasoning, reflection, and independent judgement develops considerably during these years. They begin to engage more curiously with complex questions about knowledge, society, technology, identity, and responsibility. Education at this stage, therefore, needs to enable learners not only to use emerging technologies, but also to understand how they work, recognise their possibilities and limitations, and make thoughtful decisions about their use.

Computational Thinking and Artificial Intelligence assume particular significance in this context. Computational Thinking is not limited to coding or the use of computers. It is a way of approaching problems by observing patterns, breaking complex situations into manageable parts, identifying what matters, planning clear steps, testing possible solutions, and improving them through reflection. These habits are useful across subjects and in everyday life. They help learners organise information, reason systematically, communicate clearly, and respond creatively to unfamiliar situations. Artificial Intelligence is already present in many aspects of students' lives - in how they search for information, communicate across languages, receive recommendations, make digital payments, access services, create media, and learn. At the same time, AI systems may produce outputs that are incomplete, inaccurate, unfair, or inappropriate. Students therefore need opportunities to understand AI without being either overwhelmed by it or uncritically admiring of its output. They need to ask how a system functions, what it may have learned from, what information it uses, what it may have missed, and when its output needs to be checked.

The Introductory Module on Computational Thinking and Artificial Intelligence for the Secondary Stage—Grades 9 to 12—has been developed for this purpose. It assumes no previous knowledge of Artificial Intelligence. It does not seek to turn every learner into a specialist. Instead, it aims to help every student become a careful, confident, creative, and responsible user and evaluator of AI systems. It places understanding, practical experience, discussion, and reflection above the memorisation of technical definitions. The module introduces students to AI through situations they may already have encountered in daily life. It helps them distinguish AI from ordinary automation, understand in simple terms how machines learn from examples, examine why AI systems sometimes make mistakes, and recognise the importance of data, verification,

privacy, and responsible use. These ideas are presented as connected ways of thinking, not as isolated technical concepts.

The learning experiences are designed for active participation. Students observe, question, compare, test, verify, revise, and explain. A simple student notebook and classroom activities make their thinking visible and support ongoing assessment. The module also recognises the diversity of school contexts. Access to devices, connectivity, and technological infrastructure varies considerably. Activities can therefore be undertaken through individual or shared devices, teacher demonstrations, or classroom modes that use paper, everyday objects, observation, and discussion. The core learning experiences remain possible even where access to technology is limited.

Responsible and safe use of technology is woven throughout the module. Learners are encouraged to protect personal information, verify important AI-generated information, recognise misleading or harmful content, and seek help from a trusted adult when necessary. They also encounter questions of fairness, inclusion, privacy, accountability, and the wider social consequences of technology. These ideas support responsible digital citizenship without separating ethics from the learning of technology.

Teachers play a pivotal role in creating a classroom culture where students can ask questions, examine mistakes, and acknowledge uncertainty. They need not have all the answers in advance. By modelling curiosity, careful verification, ethical judgement, and a willingness to learn alongside students, teachers can help learners develop habits that remain valuable even as technologies continue to change.

The National Council of Educational Research and Training acknowledges the contributions of the Module Development Team, reviewers, subject experts, practising teachers, designers, and all others who have supported the development of this module. Their academic, pedagogical, and creative efforts have helped shape a resource that seeks to make a complex and rapidly evolving field understandable and meaningful for young learners. I hope that this Introductory Module enables students to approach Computational Thinking and Artificial Intelligence with curiosity, confidence, and critical understanding. It reinforces a central principle: technology can extend human capability, but judgement, creativity, responsibility, and ownership of learning remain with the learner. Suggestions and feedback from students, teachers, parents, and other users are welcome for further improvement in subsequent editions.

Dinesh Prasad Saklani
Director
National Council of Educational Research and Training

New Delhi
August 2026

About the Module

This is an introductory NCERT module on Computational Thinking and Artificial Intelligence for secondary school students. No prior knowledge of AI is assumed.

The module responds to a world in which students already encounter AI in their learning, communication, public services, finance, media, and everyday decisions. It supports the wider curricular goals of the National Education Policy 2020 and the National Curriculum Framework for School Education 2023: critical thinking, computational thinking, responsible digital citizenship, and the ethical use of technology. It favours understanding and practical experience over memorisation. It does not aim to turn every student into an AI engineer. It aims to help every student become a careful, confident, and responsible user and evaluator of AI systems. The CT–AI Module for Grades 9–12 is designed as an introductory module that can be taught across all four grades, providing a common foundation in Computational Thinking and Artificial Intelligence.

What Students Will Learn to Do

By the end, students should be able to recognise AI in daily life and tell it apart from ordinary automation, explain in plain terms how a machine learns from examples, check whether an AI output can be trusted, understand and protect their own data, and use AI tools as a helper rather than a replacement for their own thinking. These are not five separate topics so much as five aspects of one habit: thinking clearly and carefully about the machines we increasingly rely on.

How the Module has been Organised

Five units build on one another. Each unit opens with the words students will meet, moves through short, readable sections, and closes with a portfolio task and an assignment in three parts: Quick Checks, Apply It, and an optional Stretch Question. The unit map below shows the shape of the whole at a glance.

Unit	Central Question	Main Habits	Portfolio Outputs
1. AI is Already in Your Life	Where do we meet AI?	Pattern recognition	AI Diary
2. How Machines Learn from Examples	How does learning from data work?	Abstraction, data thinking	Train-and-Break log
3. When AI Gets It Wrong	How do we check confidence?	Evaluation	Verification checklist / Hallucination Hunt
4. Your Data, Your Rights, Your Responsibility	What should we share or protect?	Decomposition	Privacy Policy Dissection
5. Working with Modern AI Tools	How do we use AI well?	Algorithmic thinking, reflection	Prompt Library

The Curious Notebook

Running through the whole module is one common thread: the **Curious Notebook**. Each unit leaves a lasting record in it, listed in the table above, and by the end these five pieces are not assignments handed in and forgotten. They are a record of how students have learned to think about AI, capturing observation, reasoning, revision, and reflection, and can be revisited a year later.

Recurring Features

Short boxes recur throughout the units, each with a purpose. They fall into a few categories:

- **To begin:** Words You'll Meet.
- **To do:** Try This, Pause and Discuss, and the Curious Notebook task.
- **To connect with India:** AI in India and Meet a Scientist.
- **To think carefully:** Myth/Reality, the Computational Thinking Lens, the Inclusivity Lens.
- **To go deeper:** Stretch activities, for students who wish to learn more.
- **To stay safe:** A Note on Care.

Built for Every Classroom

Because classroom resources differ widely, every hands-on activity is offered in more than one form. From a device for each student to paper and discussion alone, so that a class with no devices and limited internet can still complete the core activities of the module fully. The Note for the Teacher sets out how these forms work and how to choose between them.

Using AI Tools Safely

This module asks students to work with real AI tools, and the following habits can help. They are simple, and they matter.

- **Keep personal data out.** Do not enter your name paired along with an ID, your Aadhaar or PAN number, passwords or one-time passwords, photographs of documents, or anyone else's personal details into an AI tool. What you type can travel further than the screen in front of you.
- **Do not train tools on people.** The classroom activities use everyday objects—a pen, an eraser, or a water bottle, rather than faces, classmates, Aadhaar or ID cards, or personal documents. Practising on objects teaches the same lesson and encourages safer habits.
- **Check before you trust.** An AI tool can sound confident and still be wrong. Verify anything that matters before you rely on it or pass it on, using the Five Checks from Unit 3.

© NCERT not to be republished
NCERT DRAFT

- **Tell an adult if something feels wrong.** If you meet content that is upsetting, or that targets you or someone you know, do not forward it. Tell a trusted adult, and where needed use the official channels identified in the module.

These are not reasons to avoid AI tools. These are habits that enable students to use AI with confidence. A careful user gets far more from a tool than a careless one, while also using it more safely.

विद्यया ऽ मृतमश्नुते
© NCERT not to be republished
एन सी ई आर टी
NCERT

Letter to the Students

You are growing up in times where artificial intelligence is no longer something, hidden in laboratories, or shown only in science fiction. It is already a part of how people search, learn, translate, create, pay, travel, protect information, and make decisions. Some of these systems are helpful. Some are confusing. Some are powerful enough that they need to be used with care. This module is written to help you understand the world of AI without being frightened by it or dazzled by it.

You are not starting from zero. Recognising patterns, distinguishing what fits and what does not, judging when something looks right and when it feels off, are things you already do all day without calling them skills. This module turns these everyday habits towards the machines that are now part of your life: the phone that unlocks when it recognises your face, the app that completes your sentence, and the system that briefly holds a payment to keep you safe. Some of these machines are genuinely clever. Some only appear to be. A large part of what you will learn here is how to tell the difference.

You do not need to know coding or mathematics beyond what you are learning at school. You begin as a careful observer, someone who can ask what a system is doing, what it may be missing, when it should be trusted, and when it should be checked. What that takes is curiosity and a willingness to look closely, and you already have both.

We wrote this module to be explored carefully. You will keep an AI Diary. You will learn to train a simple image classifier and then deliberately try to break it, because analysing what went wrong often reveals more than looking only at what worked. You will learn a quick five-step habit for checking whether an answer can be trusted, and you will practise it on real outputs. Everything you make will come together in one place, your Curious Notebook, which by the end becomes a record of how you learned to think about AI.

> NOTE
>
> This is not a module to be memorised. It is a set of habits to be practised.

In this book, you will meet a few students along the way, from Bhopal, Visakhapatnam, Bhubaneswar, Kohima, and Surat, each noticing AI in a different way. They are not experts. They are curious and careful problem-solvers. We hope you will learn alongside them.

One last thing, and it matters more than any fact in these pages. AI tools are powerful, and they will keep getting more so. Used well, an AI tool can explain an idea five different ways until one of them makes sense, or help you start when a blank page has stopped you from getting started. Used carelessly, it can do your thinking for you and leave you with nothing of your own. The difference is not in the tool. It lies in how you use it. Keep the thinking yours, stay curious, and check before you trust. Do that, and these tools become something worth having: a help, but never a replacement.

Welcome. Turn the page, and bring your curiosity with you.

Module Development Team

Note for the Teacher

Thank you for taking up this module. You are teaching a subject that was rarely part of school textbooks a generation ago, to students who will spend their working lives alongside these tools. That is a rare and important thing to be asked to do, and this module is built to make it feel possible in an ordinary classroom, whatever equipment you have or do not have. You do not need to be an AI expert to teach it well. You need to be curious alongside your students, willing to say "let us check that together", and ready to model the careful, honest habits the module asks of them.

The Teaching Approach

We chose not to define artificial intelligence in a box and ask students to memorise it. We take them instead through a day in their own lives and let the ideas surface from what they already know. The organising thread is a parallel: **the student learns from examples, and so does a learning system; the student can be confident and wrong, and so can a machine.** Holding that mirror up does real work. It turns a complex subject into something a student can reason about from their own experience. In many of the activities, the most valuable thing you can do is get something wrong on purpose and let the class catch it.

Using the Three Activity Modes

Access will vary, and we have planned for that. Every hands-on activity comes in three forms: **Full** (a device per student or pair), **Teacher-Demo** (one shared screen), and **Unplugged** (paper and discussion, no tools at all). All three teach the same lesson, so a classroom with no internet can complete the core of this module in full. Unit 5, where the earlier habits are put to use with AI tools, is the capstone the module builds towards. Where devices and reliable internet are available, use the Full mode; where they are not, the Teacher-Demo and Unplugged versions cover the same ideas. The choice is yours, and our suggestion is to change the mode, not to leave out the unit.

Planning the 20 Hours

The five units are meant to be taught in order, because each builds on the habit formed in the one before. Plan each session as a short introduction followed by the activity and the discussion that grows out of it, and close a unit by having students add its piece to their Curious Notebook. Most of the class time belongs to the activity and the talk around it rather than to explanation from the front, so ask students to read the sections on their own before class. The table below suggests a plan for 20 class hours. The table below has two columns: **Core Activities** and **Enrichment Activities.**

- **Core Activities** are carried out under the teacher's direct guidance and are included in the 20-hour delivery plan.
- **Enrichment Activities** are tasks that students may take up to explore the ideas further and enrich their learning.

© NCERT not to be republished

Hour	Unit and Section	Core Activities	Enrichment Activities
1	Unit 1, §1.1–1.3	AI Scavenger Hunt; AI or Not? sorting	AI in India case-card discussion
2	Unit 1, §1.4–1.6	Indian systems with AI or pattern-matching components; pattern-recognition reflection; My AI Diary launched.	Begin My AI Diary in class.
3	Unit 1, AI Diary workshop	My AI Diary started in class; teacher walks through one or two entries with the group.	Compare diary entries; discuss what surprised students.
4	Unit 2, bridge + §2.1–2.2	How You Learn bridge spread; rules vs learning-from-examples	Connect to Class IX/X Biology on neurons
5	Unit 2, §2.3–2.4	Patterns, features, labels, categories; four kinds of data (training, validation, test, live)	—
6	Unit 2, §2.5–2.6	Bias from uneven training data; categories as human constructs	Messy Data Audit on a sample dataset
7	Unit 2, §2.7–2.9 (prep)	Where the parallel breaks; three ways machines learn; Train-and-Break setup: design the object set, discuss what could go wrong	Plan how to try to break the classifier.
8	Unit 2, §2.9 (run)	Train-and-Break run: train the classifier (Teachable Machine, full, demo, or unplugged), then try to break it	What Changed? Retrain with different data.
9	Unit 3, §3.1–3.2	Confident but wrong; why models fail; medical-screening example	—
10	Unit 3, §3.3	Deepfakes and synthetic media; Spot the Manipulation gallery walk	—
11	Unit 3, §3.4–3.5	Where your own thinking enters; build personal verification checklist using the Five Checks.	—
12	Unit 3, §3.6–3.7 (demo)	What good evaluation means; Hallucination Hunt: teacher walks through one Hunt on the projector	—
13	Unit 3, §3.6–3.7 (student-led)	Hallucination Hunt: each student runs their own Hunt in a domain they know well	Error Journal: one surprising AI mistake from the week
14	Unit 4, §4.1–4.2	Four kinds of data; consent and the DPDP Act	—

© NCERT not to be republished

15	Unit 4, §4.3–4.4	Privacy Policy Dissection of one app each student uses; fairness, inclusion, accessibility	Safe or Unsafe to Upload? scenario sort
16	Unit 4, §4.5–4.6	AI in your learning — helper stance, not replacement; if something goes wrong	Class discussion: where can a complaint go?
17	Unit 5, §5.1–5.3	Asking better questions; helper stance; build Prompt Library (Rohan worked example, then student tasks)	—
18	Unit 5, §5.3 (peer review)	Students review each other's prompts in pairs and iterate together	Prompt Improvement Ladder
19	Unit 5, §5.4 + closing	Reflecting on what you made; Curious Notebook sharing and review	Full-module reflection
20	Closing session	Each student presents one of their five Curious Notebook artefacts; the class discusses what they noticed.	Cohort reflection on habits to carry forward

Helping Students with the Work

Each unit closes with an assignment in three parts. Quick checks confirm that a student followed what happened, Apply it asks them to carry an idea into a situation the unit did not cover, and the Stretch question is genuinely optional and is not meant to be set for the whole class. Apply it questions go better after one worked example on the board than after further explanation. When a student is stuck, the most useful move is to ask what they noticed rather than to supply the answer, because every question here is answered from observation and reasoning rather than from recall.

Assessment Through the Curious Notebook

A running portfolio called the Curious Notebook carries a student's work across the module, and it is where assessment naturally lives. It may be used for formative assessment: it captures observation, reasoning, revision, and reflection, not just correct answers. The Notebook is a working record and not a fair copy. Crossings-out, revised answers, and a first guess that turned out wrong are exactly the evidence worth having, and a student should not be marked down for them. Look through the Notebooks from time to time and leave a short comment rather than collecting them in for correction. Any exercise book or a set of stapled sheets serves perfectly well, so that no student is at a disadvantage for want of a special notebook. Every question in this module asks a student to apply an idea, judge a case, or explain a decision in their own words, and any questions you set later should follow the same pattern. The recurring boxes throughout the units are invitations to stop and think, not optional extras. Both the Notebook and the boxes are described in full in About the Module.

Handling Sensitive Topics

Some parts of this module touch on difficult subjects, such as deepfakes, data misuse, and online harm. If a discussion moves in that direction, or a student is in distress, please do not turn it into a classroom spectacle. Pause, protect the student, involve a trusted adult or school authority, and follow the reporting pathways indicated in the module. Treated with care, these moments teach something no other part of the course can.

The habits your students build here will outlast this module, and outlast whichever tools happen to be current. Thank you for carrying them across.

Module Development Team

© NCERT not to be republished

Module Development Team

Mayank Vatsa, *Professor*, IIT Jodhpur

Mitesh M. Khapra, *Professor*, IIT Madras

Rajesh D, *Associate Professor*, CIET, NCERT

Tarun Choubisa, *Chief Consultant*, Programme Office, NSTC

Pranita Gopal, *Assistant Professor*, CIET, NCERT

Anubhooti Jain, *Research Associate*, IIT Jodhpur

Chiranjeev, *Research Associate,* IIT Jodhpur

Muskan Dosi, *Research Associate,* IIT Jodhpur

Reviewers

Shekhar C. Mande, FNA, FASc, FNASc, Former *Director General,* CSIR; Distinguished *Professor*, Bioinformatics Centre, Savitribai Phule Pune University; Honorary Distinguished *Scientist*, National Centre for Cell Science, Pune

Anurag Behar, *CEO*, Azim Premji Foundation; Member, National Curriculum Framework Oversight Committee

Abhay Kumar, *Professor*, CIET, NCERT

© NCERT not to be republished

ACKNOWLEDGEMENTS

The Module Development Team acknowledges the contributions of Prateek Sharma, *Graphic Asstt. Gr. I (MPD)*; Sadia, Adiba, Amandeep Singh, Devahuti Sharma, and Brijesh in the layout and design of this module.

The NCERT acknowledges the valuable contributions of Surender Kumar, *In-charge*, DTP Cell, Publication Division, NCERT and his team—Sanju Sharma, Vivek Mandal and Sadiq Saeed, *DTP Operators* (Contractual), for their continuous assistance in the finalisation of this book. The Council also appreciates the efforts of Rajat Subhra Roy, *Proofreader* (Contractual), Publication Division, NCERT for his meticulous work and unwavering commitment throughout the publication process.

© NCERT not to be republished

CONTENTS

Foreword *iii*

About the Module *v*

Letter to the Students *viii*

Note for the Teacher *ix*

Unit 1: AI is Already in Your Life 1

Unit 2: How Machines Learn from Examples 13

Unit 3: When AI Gets It Wrong 39

Unit 4: Your Data, Your Rights, Your Responsibility 54

Unit 5: Working with Modern AI Tools: Prompts and the Helper Stance 69

Glossary 83

© NCERT not to be republished

THE CONSTITUTION OF INDIA

PREAMBLE

WE, THE PEOPLE OF INDIA, having solemnly resolved to constitute India into a [1]**[SOVEREIGN SOCIALIST SECULAR DEMOCRATIC REPUBLIC]** and to secure to all its citizens :

JUSTICE, social, economic and political;

LIBERTY of thought, expression, belief, faith and worship;

EQUALITY of status and of opportunity; and to promote among them all

FRATERNITY assuring the dignity of the individual and the [2][unity and integrity of the Nation];

IN OUR CONSTITUENT ASSEMBLY this twenty-sixth day of November, 1949 do **HEREBY ADOPT, ENACT AND GIVE TO OURSELVES THIS CONSTITUTION.**

---

1. Subs. by the Constitution (Forty-second Amendment) Act, 1976, Sec.2, for "Sovereign Democratic Republic" (w.e.f. 3.1.1977)
2. Subs. by the Constitution (Forty-second Amendment) Act, 1976, Sec.2, for "Unity of the Nation" (w.e.f. 3.1.1977)

UNIT

1

AI is Already in Your Life

UN514UN01



You walk into the school canteen and spot your friend across a crowd of fifty people in less than a second. The third umpire watches a slow-motion replay in a cricket match and decides whether the ball touched the bat; a tiny edge, or a tiny sound is enough. You hear three notes of a song and know the singer. Each of these is pattern recognition, and you have been doing it since long before you knew the word.

Some of these same pattern-recognition tasks are now done by machines; recognising faces in crowds, recognising songs from snippets, spotting unusual banking transactions in a stream of routine ones, and translating sentences from one language to another. None of this is science fiction, and most of it is not new. It is just that, in the last few years, the machines

have got noticeably better at some of these tasks, and worse at others, in ways worth knowing about.

This first unit is a tour. We are not going to define Artificial Intelligence (AI) in a box and ask you to memorise the definition. We are going to walk through a day in your life and try to notice where the AI or automated machines have shown up.

By the end of the unit, you will have a basic understanding of what people mean when they say 'AI', where it shows up in your daily life, and, most importantly, how to tell it apart from the older, simpler systems that often get called AI.

WORDS YOU WILL MEET IN THE CHAPTER

Words You'll Meet in this Unit

A few new terms before you start. Each gets a fuller definition when it appears in the unit, and the alphabetical glossary at the back of the module has them all.

- **Pattern Recognition:** Noticing that something new is similar to something you have seen before. You do it constantly without thinking about it.
- **Simple Device:** A machine that responds to a physical condition without making decisions or learning. Examples include a pressure cooker or an alarm clock.
- **Automation:** A system that follows fixed rules someone wrote in advance. Examples include a FASTag toll or a vending machine.
- **Artificial Intelligence (AI):** In this module, we mostly mean systems that learn patterns from examples rather than only following rules someone wrote in advance. When we need to distinguish these from simple devices and rule-following automation, we call them **learning-based AI**.

1.1 You are Already a Pattern-spotter!

Most of what you do during the day involves noticing patterns: reading a friend's mood from how they walked into class, predicting which of the two queues at the bus stop will move faster, knowing from one smell that something in the kitchen has just begun to burn, or guessing from that same smell that your favourite dish is cooking today. None of this feels like work; you do it without trying, often without noticing. But each one is a small act of pattern recognition, and you do thousands of them every day.



Fig. 1.1: Patterns in everyday life

This matters for our purposes because the same task—noticing patterns—is what AI systems are designed to do. They notice patterns in faces, voices, in bank transactions, in sentences, and in images. They are sometimes faster than you. Unlike you, they usually only know how to recognise one kind of pattern. They almost never have the context you have: who the person is, what they were doing yesterday, or why this might matter. But the underlying task—recognising patterns—is the same for you and the machine.



Pause and Discuss

You walk past the same shop every day for a year. One day, something feels different about it, but you cannot say what. You stop and look. The signboard is the same. The shopkeeper is the same. But something is off. Twenty minutes later, you realise; a small plant by the door has been moved. Could a machine have noticed this? What would it have needed in order to notice it?



What You have Learnt So Far

- Pattern recognition is something you do constantly, often without noticing.
- AI systems can perform similar pattern-recognition tasks. Sometimes faster, but usually within limited scope, and less context.

1.2 What AI is, and What it is Not

The phrase 'AI' gets used loosely nowadays. To make sense of the rest of this module, the two points below will help.

First, AI is a broad field with a long history. It includes systems built on hand-written rules; on search and planning; on logical inference; and, most prominently today, on learning from examples. In this module, when we say AI, we will mostly mean systems that learn patterns from large numbers of examples rather than just following rules someone wrote in advance.

Second, real AI systems often combine more than one approach: some of the rules are written by hand or in code, while some are learned.



Try This

List five things you spotted today that involved noticing a pattern. They need not to be big things; small ones count too, for example, 'Recognised my friend by the back of her head.' 'Guessed who was on the phone from the first ring.' 'Knew the bus was about to slow down before the brake lights came on.'

Beside each, write whether you think a machine could do the same task today. If **yes**, how well? If **no**, what does the machine lack: information, context, experience, or something else? Information, context, experience, or something else?

To get you started, here is one example completed for you:

'Recognised my friend by the back of her head.'

Could a machine do this?

Partly. A machine could use her hair, build, clothes, or the way she walks, but it would need many examples of your friend first, and even then it would not know to look for her the way you do. (A face recognition system cannot work from the back of the head.)

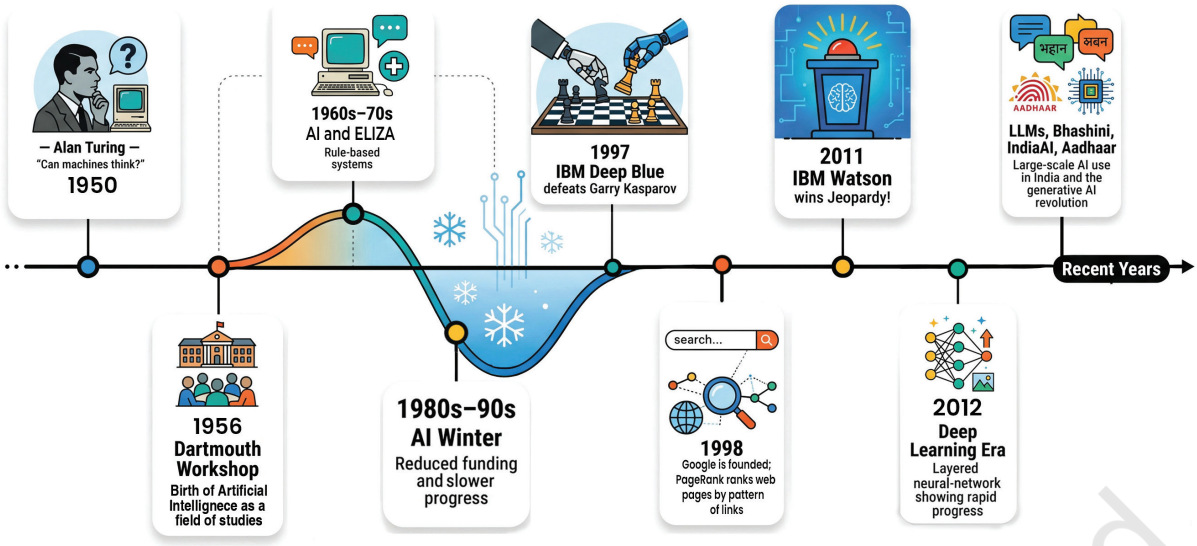


Fig. 1.2: The progress of AI over the years

AI: A Long History, in Brief

AI did not arrive yesterday. The field has roots going back over seventy years, with periods of slower progress and periods of remarkable progress. Below are a few stops along the road to give a feel for the shape of it.

- **1950:** Alan Turing publishes *Computing Machinery and Intelligence*, proposing the test that bears his name and asking, in a careful way, the question this whole field rests on: "Can machines think?" (Turing, Alan M. 1950. "Computing Machinery and Intelligence." Mind 59, no. 236. 433–60)
- **1956:** A summer workshop at Dartmouth College gathers the small group of researchers who would later be called the founders of this field. The term 'artificial intelligence' had been introduced in the 1955 proposal for this workshop.
- **1960s–70s:** The first wave brought symbolic AI and rule-based systems. Programs like ELIZA (which simulated a psychotherapist using simple rules) and the first medical-diagnosis expert systems appear. Most AI of this period relied on carefully written rules.
- **1980s–90s:** A long, hard period that the field itself calls the 'AI winter.' Many AI projects were discontinue because they do not deliver on big promises. The work continued quietly in a few places.
- **1997:** IBM's Deep Blue defeats the world chess champion Garry Kasparov in a six-game rematch in New York.
- **1998:** Google is founded, building on PageRank, a method developed at Stanford and ranks web pages using the link structure of the web.
- **2011:** IBM Watson defeats the top human champions on the American quiz show Jeopardy!, a contest that demands wide general knowledge and quick understanding of tricky, pun-filled clues.
- **2012 onwards:** A deep neural network using learning system dramatically outperforms the older approaches of recognising images. Over the next decade, deep learning transforms image recognition, speech recognition, translation, and many other tasks.
- **Recent years:** Large language models put text-generating AI in the hands of hundreds of millions of users. India developed initiatives, such as Bhashini for Indian-language translation and the IndiaAI Mission to build National AI Capacity. Aadhaar's biometric authentication operates at one of the largest scales of any deployed identity system in the world.

What looks like a sudden revolution has roots seven decades deep. The recent leaps are fast, but they are leaps in a long walk.

© NCERT not to be republished

The following three terms give us a useful way to understand what AI is and what it is not.

- **A simple device:** A pressure cooker whistles when the steam pressure inside crosses a set point. The cooker makes no decisions and does no learning; there is no controller executing instructions, just a simple physical mechanism doing the work. A doorbell, an alarm clock, or a thermometer are other everyday examples of a simple device.

- **Automation:** FASTag reads the tag on your vehicle, looks up your account, deducts the toll amount, and lets you through the plaza. A decision is made here, but the rules for reaching it were written by someone in advance. With automation, the same input in the same situation always triggers the same fixed rule, and so gives the same response every time. mKisan, the SMS service that sends crop advice to farmers, works in much the same way: given the crop, the district, and the date, send this message. Reliable, predictable, and mostly rigid.

- **AI (in the sense this module uses):** When typing on a smartphone, whether your own or someone else's, you might have noticed it predicting your next word and offering to complete it for you. The phone looks for patterns in the words and phrases people tend to use and suggests what is likely to come next.

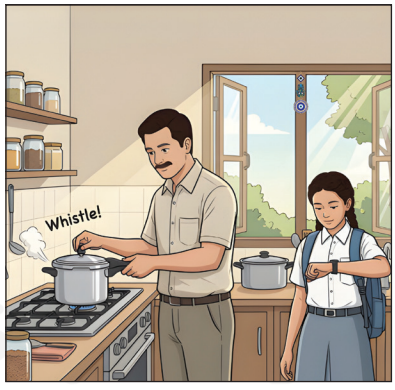


Fig. 1.3: A simple device

Simple Device

A device that whistles, rings, and lights up in response to a physical condition. No decisions, and no learning. (Pressure cooker, alarm clock, doorbell.)

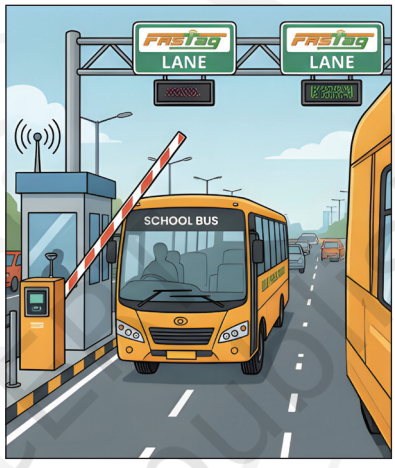


Fig. 1.4: Automation

Automation

Given the same input and the same system state, it follows the same fixed rule someone wrote in advance. (FASTag, mKisan, a vending machine.)



Fig. 1.5: Learning-based AI

Learning-based AI

Learns patterns from examples. Its output may change if context, data, thresholds, or the model version changes. (Payment-system safety checks, Bhashini translation, Photo-app face grouping.)

In cricket, you may have watched the third umpire review an LBW appeal using (Hawk-Eye). Multiple high-speed cameras observe the ball from different angles. Computer-vision methods help locate and track the ball, while triangulation reconstructs the three-dimensional path. A trajectory model then estimates where the ball would have travelled had the pad not been in the way. Modern Hawk-Eye systems use machine-learning methods in parts of the detection and tracking

Group Work

Below is a list. For each one, decide: simple device, automation, or AI? Some are easy. Some are genuinely hard, and disagreement is fine.

(a) A digital alarm clock that rings at 6:30 every morning.
(b) The autocorrect in the keyboard of any phone you use.
(c) The recommendations on a music app.
(d) The traffic-light timer at a busy junction near your school.
(e) Face recognition on any phone you use.
(f) A washing machine that runs a fifteen-minute cycle.
(g) An AI assistant such as ChatGPT, Gemini, or Claude.
(h) The IRCTC website's chatbot that answers questions about train bookings.

Discuss your answers in the group. If there is any disagreement, identify what evidence would settle the disagreement.

process, but they also depend on camera geometry, mathematical modelling, and physics. Hawk-Eye is therefore best understood as a hybrid system rather than a simple automation or purely learning-based AI.

A third example is something you may have used yourself. To keep UPI transactions secure, many payment platforms and banks run them through layered checks that combine fixed rules with learning systems. The learning part asks questions like: Does this look like the kind of transfer this person usually makes? Is this the time of day they usually transact? It estimates risk from patterns learned across many past transfers, so the same amount might pass today and be held tomorrow if the account context, recipient, timing, or other factors have changed. Unlike fixed automation, a learning system is not rigid. It can surprise us, and it is not always right, which is why we need to think carefully before accepting what AI gives us.

Why does this distinction matter? Because not everything that works automatically is AI, and calling everything 'AI' hides an important difference. A rule-based system often fails in predictable ways: if the rules do not cover a case, it usually cannot handle that case well. A learning system fails in less predictable ways. It may handle a case its builders never imagined, but it may also miss one they would have thought obvious. Knowing which kind of system you are dealing with tells you what kinds of mistakes to expect and how to check for them.

1.3 Where You Meet AI Every Day

Once you start looking, AI shows up in many small ways throughout an ordinary day. Not all of them are dramatic; most are routine moments you might easily miss. The list below is only a partial list; the purpose is simply to make the kinds of places visible.

- **Morning:** A phone unlocks when it sees a face. The keyboard suggests the next word as someone types. A weather app forecasts rain or a heatwave for your area, increasingly aided by AI, helping a family or a farmer decide when to irrigate or what to carry.



Fig. 1.6: Phone unlocking the user's face

- **While travelling:** If you take a shared tempo to the nearest town, or book an auto through an app at the bus stand, a system may match passengers, estimate arrival time, or suggest the fastest route

around a blocked road. A music app recommends the next song based on what you have listened to.

MYTH	REALITY
All three, pressure cooker, FASTag, and autocorrect, are different kinds of AI.	Only the third is learning-based AI. The pressure cooker responds to physical conditions (pressure and heat) without making decisions. FASTag follows fixed rules someone wrote in advance: read the tag, look up the account, and deduct the toll amount. Autocorrect is the only one that worked out its own patterns from examples, by being shown enormous amounts of typed text. Three different kinds of system, only one of them is learning.

- **During the day:** A voice assistant answers questions in your local language when typing is hard, because it has learned patterns in regional speech. A mandi-price app shows today's crop rates and, in some cases, learns from past prices to suggest a good time to sell. A messaging app suggests quick replies. A plant disease detection app reads a photo of a leaf and tells you what it may be.
- **Evening:** A video or music app suggests local-language songs, devotional content, or short farming and skilling videos based on what you have watched. A streaming app on a shared TV suggests the next show for the family. A bank or UPI app may hold an unusual transfer for a moment of verification, protecting a first-time digital user, or a small trader from fraud.

Most of these have been around for years, and none would seem unfamiliar to you. That is part of the point: AI in your daily life is rarely one big, dramatic event. It appears in many small ways, quietly working in the background, often successfully, and occasionally in surprising ways.

1.4 AI in India

Many of the examples in the previous section are global: the streaming app on your phone is the same one someone uses in another country, and so is the weather app. But India also develops its own AI systems, designed for and around Indian conditions and languages. Three are worth knowing about through this unit, because each tries, in its own way, to overcome barriers of access and equity that a global system was never built to address.

AI IN INDIA

Aadhaar biometric authentication: Aadhaar offers several authentication modes. The biometric ones use fingerprint, iris, or face, depending on the service and device. The system compares a fresh biometric sample with the record on file. It is built to operate across hundreds of millions of users and has been refined over the years. Like any pattern-matching system, its reliability depends on capture quality, lighting, the device used, and how closely the current sample matches the conditions under which the record was made. The system may not work equally well for everyone. Some users are harder to match, so the design must include alternative ways to verify them.

Bhashini: Bhashini is an AI-powered language translation platform, built by training the system on enormous numbers of sentence pairs in many Indian languages. Bhashini-enabled services, such as multilingual government chatbots and voice-based farmer-support tools, help people access information in their own language. It is the reason a farmer in Maharashtra using one of these services can ask a question in Marathi and get the answer in Hindi, or the other way around. The quality varies across language pairs, usually better for the pairs with more examples to learn from, and it keeps improving as more data becomes available.

UPI transaction-safety checks: UPI transactions in India occur at an enormous scale, in crores each day, according to the figures from the National Payments Corporation of India (NPCI). Many payment apps and banks run layered fraud-risk checks on transactions, and these may combine fixed rules with learning systems. The learning parts, when used, match a transaction against patterns of usual behaviour. Most transactions pass through in milliseconds, while a small fraction are held for a closer look because the transaction does not match those patterns. The exact mix of rules and learning models varies across banks and apps.



Fig. 1.7: Inaya's AI translation experience

Meet Inaya in Bhopal: Inaya is in Grade 11 in Bhopal. She speaks Urdu at home, Hindi at school, and English in her local community drama club. One day her cousin sends her a message: "Uff, mere bhai ne naak mein dum kar diya hai!" Inaya laughs. The translation app on her phone offers her the literal English: 'Oh! my brother has put life into my nose!', which makes no sense at all. The real meaning is closer to 'My brother is troubling me.' The app has translated the words. It has not translated why this is what someone says when they are annoyed. Inaya realises that the AI translation is correct word for word, but it misses the real meaning that many idioms carry.

1.5 What AI is Good at, and What it Struggles with

Table 1.1 shows, side by side, the kinds of tasks today's AI systems do well and the kinds they still find hard. Together, the two columns give a broad picture of where AI stands at present.

Table 1.1: Tasks AI does well, and tasks it finds hard

What current AI is good at	What current AI struggles with
Recognising patterns in large piles of similar data, photos, sentences, and transactions.	Reasoning about a situation that it has never seen, especially a situation that requires common sense.
Producing fluent text, smooth, natural-sounding sentences, on familiar topics, summarising long documents, and drafting emails.	Producing reliable specific facts, names, dates, exact quotations, especially for narrow or specialised topics where less text was available to train on.
Translation of well-resourced language pairs.	Translating jokes, poetry, idioms, and anything that depends on cultural context.
Sorting items into a small number of categories, when the categories are clear.	Handling edge cases that sit between two categories, or new categories that the system has not seen.
Answering questions on topics well represented in its training data, or in provided source material.	Knowing what it does not know, most current systems will produce a confident answer even when the right answer is "I don't know."
Assisting humans on tasks where the human can check the output.	Replacing humans on tasks where nobody is checking.

This list will change. The line between 'good at' and 'struggles with' is always moving, so what AI finds hard today may become easy in a few years. Tasks that were impossible ten years ago are routine now, and tasks that seem impossible today may be routine in five years. But the basic idea stays the same: some tasks suit the technology, others do not. So the question to ask when you meet a new AI system is not 'is it impressive?' but 'what kind of task is this, and is it the kind the technology handles well?'

In Unit 3 you will learn a step-by-step method for the harder version of this question: what to do when the system's answer is confident, fluent, and wrong.

What You have Learnt So Far

Aadhaar biometric authentication, Bhashini, and UPI transaction-safety checks are three Indian systems with important pattern-matching or AI-enabled components. Each is good at the task it was built for, but each has limits worth knowing.

WHAT YOU HAVE LEARNT SO FAR

Current AI is good at recognising patterns at scale (across very many examples), fluent at producing text, but weak on common sense and on knowing what it does not know.

When you encounter a new AI system, ask what kind of task it is doing, not how impressive it appears.

© NCERT not to be republished

1.6 Your turn: starting your AI Diary

Before we move to the next unit, let us start an AI Diary.



Computational Thinking Lens

You have been learning about and practising pattern recognition throughout this unit. You spotted AI across systems as different as a music recommender, a translation app, and a transaction-safety check. You noticed the common thread: beneath their surface differences, each system is doing the same underlying task—recognising patterns. Noticing that several different things share a hidden structure—that is pattern recognition—and it is the habit at the heart of this unit.



Inclusivity Lens

Voice assistants and screen readers are not new. For a student with low vision, a voice assistant is not a convenience; it is the way the phone becomes usable at all. The same pattern-recognition technology that powers a music recommender powers an accessible interface.

Bhashini, in particular, makes services available to users who are not comfortable in English or Hindi. They can speak in their own language and receive responses in their own language. Within a single language, the system may also perform differently across accents.

A speaker of Hindi from Mumbai and a speaker of Hindi from Patna may both be understood, but with different reliability. The same applies to English, as an Indian-English speaker, a UK-English speaker, and a US-English speaker get different results from many voice assistants.

The technology is the same. The difference it makes is not.

Why keep a diary?

Reading about AI in daily life is one thing; spotting it as it operates in your own day is another. The diary is a way to make the unit's ideas live in your own routine, not just on the page. After one day of careful noticing, you will have a clearer sense of how and where AI shapes your life, and that experience will teach you more than any number of definitions could.

How to do this activity?

Choose any one day—a school day or a Sunday, whichever you like. Throughout the day, keep a list of every system you use or come across (an app, a device, or an online service) that you think might involve AI. For each entry, write three details: what the system was, what you wanted it to do, and how well it worked. Aim for at least eight entries by the end of the day. Most students find more than they expected.

© NCERT not to be republished

For your Curious Notebook

AI Diary: One day in my life, observed.

S.No.	The system or app	What you were trying to do with it	How well it worked
1			
2			
3			
4			
5			
6			
7			
8			

A sentence is enough, but be specific. 'Keyboard suggested the wrong word three times in this message' is more useful than 'Autocorrect was bad.'

Reflection: At the end of the diary entry, write a short note of about five lines. Which entries surprised you? Was there a moment when you realised that a system was using AI, when you had thought of it as just a feature? Was there a moment when you assumed something was AI, but on reflection it turned out to be ordinary automation?

Bring it to the next class. Compare your diary entry with your friend's and see which AI systems turned up in most of them.



© NCERT not to be republished

REFLECT

Before you started this unit, what was your instinctive idea of what AI is? Describe it in two lines.

After reading this unit, has anything changed in that idea? Write two more lines. There is no right answer here; the point is to simply to notice the change.

Assignment

Quick checks

1. Pick one app on your phone that you use most days. Name two pattern-recognition tasks it likely runs in the background, and one task you suspect it cannot do well.
2. mKisan and an AI assistant such as Gemini or ChatGPT both involve a phone, a query, and a response. One is mostly AI, while the other is mostly not. Which is which, and how can you tell?
3. A friend says, "Aadhaar face match is just like a human recognising a face." In two lines, describe what the friend gets right and what the comparison misses.

Apply it

4. You are designing a new feature for a music app used in three Indian languages. Describe one pattern the app would need to recognise, and one mistake that pattern recognition is likely to make in the first month.
5. Over the next 24 hours, spot one system you first assumed was AI but that turned out to be ordinary automation, or the reverse. Write down the system, your initial guess, and what changed your mind.

Stretch (optional)

6. A payment app on your aunt's phone—PhonePe, Google Pay, or Paytm, depending on which she uses, holds her electricity-bill transfer for a moment of verification before letting it through. Write a half-page note that explains, in plain language, so that she can follow why an app might do this and what tradeoff makes such momentary holds hard to avoid completely.

Coming up in Unit 2: You now have a working sense of what AI is, and how to tell it apart from the systems that only look like it. The next unit goes one level deeper—how do machines learn from examples in the first place, and how does that compare to the way you learn?

© NCERT not to be republished

UNIT
2

How Machines Learn from Examples

UN514UN02



When we were two or three, we may not have known exactly what a cat was. Someone pointed to a furry, four-legged animal and said, 'cat.' Later, we might have seen a rabbit or a squirrel and called that a 'cat' too. That mistake was part of the learning. With more examples and corrections, we slowly learned that not every furry four-legged animal is a cat. We were not given a perfect definition. We built our understanding around the examples, mistakes, and feedback.

We have all learned many things this way. Without anyone writing down rules for us, we came to recognise faces, voices, songs, moods, even the mood of the school principal from the way she walks down the corridor. We often learned by being shown examples, making a few mistakes, and correcting them on the way. The machines in this unit try to do something similar:

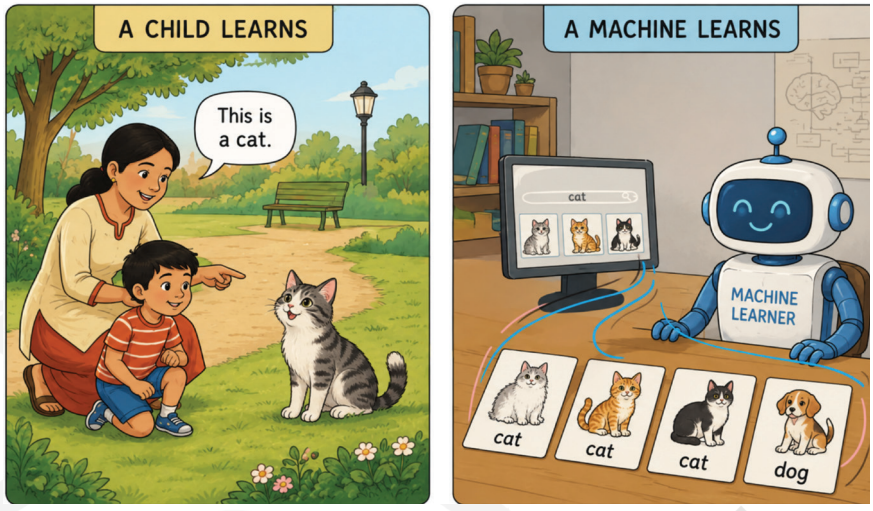


Fig. 2.1: A visual representation of a child and a machine, learning about a cat

they learn from examples rather than from rules someone wrote in advance. They are not as good at it as we are, in most ways, but in some ways they are faster. The point of this unit is to see how machines do it, where there is a parallel with our own learning, and where the parallel breaks.

WORDS YOU WILL MEET IN THE CHAPTER

Words You'll Meet in this Unit

This unit introduces several technical terms. Each one is explained when it first appears, and the alphabetical glossary at the back lists all of them, but it is worth knowing these words in advance:

- **Supervised, Unsupervised, and Reinforcement Learning:** Three ways a system can learn from examples, depending on what kind of feedback the examples give.
- **Example, Feature, and Label:** These are three building blocks of supervised learning. An 'example' is one labelled item used to help the system learn (20 pen photos are 20 examples). A 'feature' is a measurable property of an example (colour, shape, length, or the position of the clip). A label is the answer attached to an example (blue pen, red pen, eraser), the category that example belongs to.
- **Bias (in a Learning System):** An uneven pattern in the system's behaviour that can result from limitation in its training data. If the data had many examples of one kind and few of another, the system may perform unevenly, reflecting that imbalance. A face-recognition system trained mostly on adult faces, for example, may work less reliably on children's faces. Making the system smarter is not the answer. One important way to reduce such bias is to collect more varied training data. For a face-recognition system, that data should include adults, adolescents, young children, older people, and people from different parts of the country.
- **False Positive and False Negative:** Two kinds of error a system can make. In a false positive, the system says 'Yes' when it should have said 'No.' In a false negative, it says 'No' when it should have said 'Yes.'
- **Neural Network:** A kind of mathematical structure used in many learning systems today, loosely inspired by the brain but not the same as a brain.

2.1 How You Learn

Before we look at machines, let us look at our brain, the organ that plays a central role in our learning and memory. The picture in Fig. 2.2 is simplified. It captures the main idea, but how the brain actually works is far more complicated, and biologists are still making new discoveries about it. The picture is useful because it helps us to compare how our brain works with how AI systems work, the ones we are about to study, and the ones we already use every day.

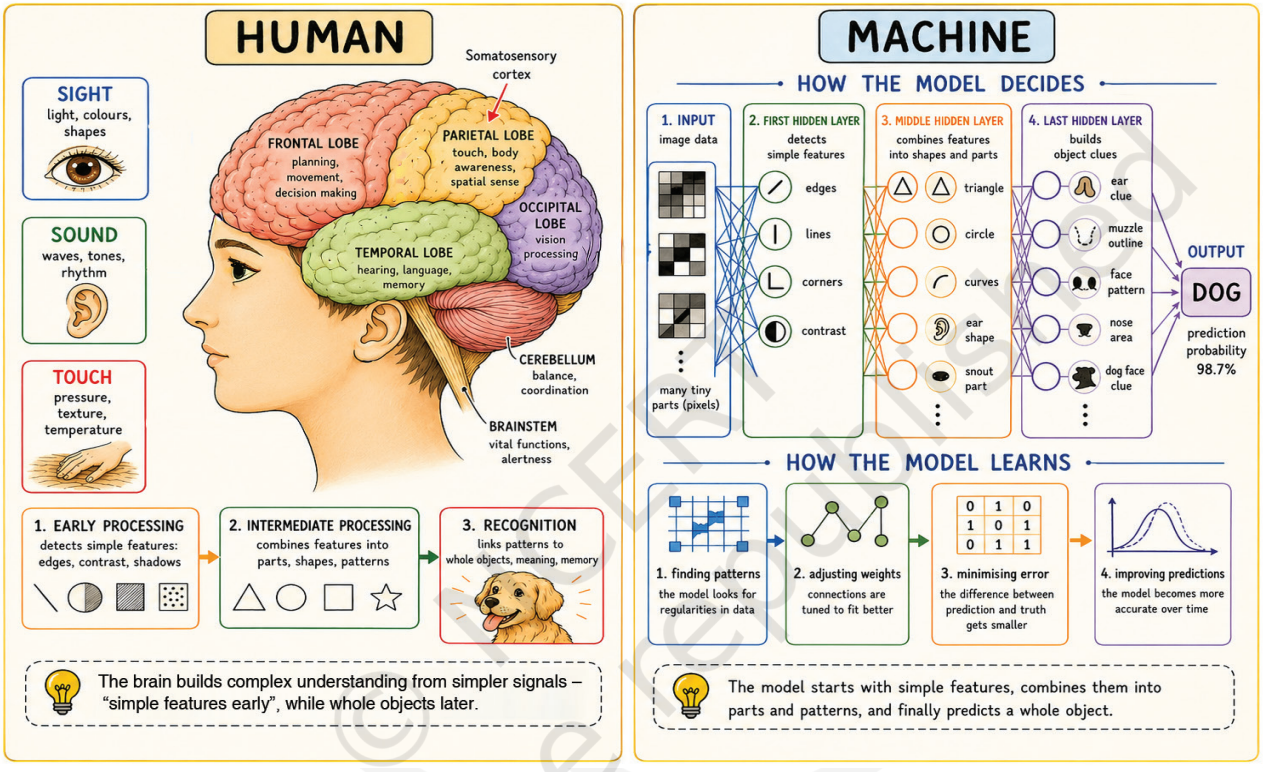


Fig. 2.2: Comparison of how the human brain and a machine learn

Think about how you recognise your friend Karthik across a crowded school canteen. It begins with your senses. You see his face across the room; you hear three notes of a song from the next room and know your favourite song is playing; in the kitchen, you feel the heat of the pressure cooker on your skin and know it is about to whistle. Each of these experiences begins with information reaching your brain, light reaches your eyes, sound reaches your ears, and heat reaches your skin, picked up by an organ that turns it into signals the brain can work with.

Inside the brain, cells pass these signals to one another. From your earlier science classes, you already know that the brain is made up of specialised cells called neurons. Neurons send

Try This Now

Close your eyes and listen for the next thirty seconds. Then write down three distinct sounds you heard, and what each one made you think or feel.

Now look back at your three answers.

For each sound, three things happened: you heard it, you recognised what it was, and you felt something about it. Which of these do you think a machine could do well? Which do you think it could not?

Let your answers guide you through the rest of the unit.

electrical impulses along their length and communicate with the other neurons at junctions called synapses, often through chemical signals called neurotransmitters. The brain contains billions of neurons, constantly passing signals among them.

The processing in the brain happens in layers. When you see something, your brain does not make sense of it all at once. The information from your eyes passes through several stages. In the first, the brain notices simple features, such as edges, light, shadows. In later stages, it combines these into shapes, curves, a hairline, the line of a chin. Finally, the clues come together and you recognise the person: 'That's Karthik.' The recognition is the output. It is the same for a name, a guess, a feeling, or a decision.

The journey, from input to output, happens so fast that it feels like a single moment of recognising. Behind that one moment, many small parts of your brain have been doing quick work. You never feel the layers; you only get the answer, 'that's my friend Karthik.' The brain takes in information through the sensory organs, billions of neurons process it through layers, and we get the answer.

How do we get a machine to do something similar?

What if a machine could attempt to carry the same process—take an input, push it through layers of processing, and produce an output, but using mathematics instead of neurons? That is the question the next sections take up. The systems that do this are called artificial neural networks, because they were loosely inspired by the picture of the brain you just read.

We should not confuse artificial neural networks with human brains. A neural network is a mathematical structure that shares a few features with the brain and misses most of what a real brain does. But the basic process is the same: like the brain, a machine that learns takes an input, process it through a series of steps, and produces an output that depends on the input.

But how does a machine actually do this: take an input, pass it through layers, and arrive at an output, without a brain to do the work? Let us read ahead.

What You have Learnt So Far

Your senses bring inputs into your brain. Layers of neurons process them, and you experience the output: "this is a dog," "this is a red rose," "my favourite song is playing." You do not experience all the work the neurons did to produce that output; you only get the answer. Machines that learn try to do something similar, loosely inspired by the brain, though they are not the same as the brain.

PAUSE AND DISCUSS

Your science textbook calls the connections, or junctions, between neurons "synapses." If a machine version has no neurotransmitter chemicals to send across a junction, what kind of connection might it use instead?

Meet a Scientist



In 1949, the Canadian psychologist Donald Hebb proposed a simple idea about how the brain learns: when two brain cells are repeatedly activated at the same time, the connection between them strengthens. This is now known as Hebb's rule, often summarised as *"Cells that fire together, wire together."* It also helps explain why practice works. Whenever you practise something, a school subject, a sport, a piece of music, the same brain connections are used again and again, and the learning grows stronger.

Hebb's rule helped researchers think about learning as a change in the strength of connections. Some modern neural networks work in a similar way: they learn by adjusting the numerical weights on their connections, strengthening some, and weakening others as they improve. The mathematics they use today is not Hebb's rule itself, but the underlying idea, that learning reshapes connections, is one Hebb helped put in place.

2.2 Rules versus Learning from Examples

There are two main ways to make a machine do a task that involves recognising things or making decisions.

- **The first way is to write rules.** A human expert thinks hard about the task, breaks it into steps, and writes the rules down. Then a programmer turns those rules into code. The machine follows the rules according to the instructions it has been given. This is how mKisan, the SMS advisory service for farmers, mostly works: the rules can specify the conditions, such as 'if the crop is cotton and the district is this and the date is in this range, then send this advice.' Given the same input and the same system state (district and date range), mKisan gives the same fixed response.
- **The second way is to show examples.** Instead of writing every rule, you collect a large number of examples, ideally thousands, sometimes millions, each labelled with the right answer. You feed these to a learning system, and it learns its own internal pattern of what each label means. This is how Bhashini learns to translate from Marathi to Hindi where lakhs of sentence pairs, in both languages, are fed to a learning system that works out the language patterns for itself.

The first way is precise but rigid—if the rules do not cover a case, the system cannot provide the intended response. The second way is flexible but messy. A learning system trained on a million examples can handle cases its builders never imagined, but it can also make mistakes in ways a rule-following system would never make.



What You have Learnt So Far

There are two ways to teach a system: we can write the rules, or we can provide it with examples. Writing rules are precise but rigid; learning from examples is flexible but depends heavily on the examples themselves. If the examples are few, or not relevant, the answers become less reliable.

AI in India

Bhashini and Indian languages: India has 22 languages listed in its Constitution and many more spoken in everyday life. Translating between them only by writing rules is extremely difficult and can become difficult while maintaining at scale, as the grammar of Marathi is not the grammar of Tamil, which is not the grammar of Manipuri, and the rules multiply faster than any team can write them. Bhashini was built using the second way. It was shown enormous numbers of sentence pairs in many Indian languages, and from those pairs it learned and developed its own translation patterns.

It works well for some pairs and is still being improved for others. The pairs it works best on are usually the pairs for which it had the most examples to learn from. The pairs it still struggles on are usually the pairs for which there were fewer examples available. The lesson generalises: how good a learning system is at any particular task depends on how good its examples for that task were.

2.3 What both systems need: examples, features, labels, and categories

Whether you are learning to recognise cats, or a machine is learning to recognise sentences in Marathi, the same few ideas keep showing up. We need to understand them.

- **Example:** An example is one item the system learns from, a single photo, a single sentence, a single sound. 20 photos of one blue pen are 20 examples. The three terms below all describe something about an example.

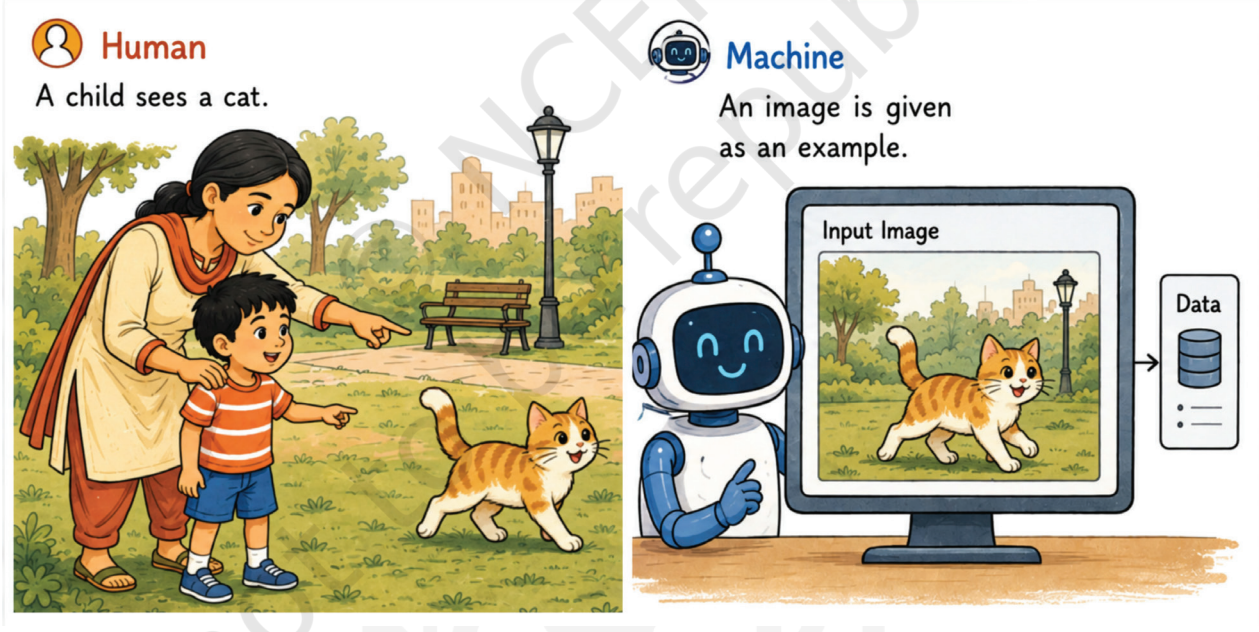


Fig. 2.3: Learning from examples: how humans and machines do it

- **Feature:** A feature is one particular detail you, or a machine, might pay attention to in order to recognise something. The pitch of a voice is a feature. The colour of a pen cap is a feature. The presence of whiskers is a feature when you are deciding whether a furry animal is a cat or a dog.

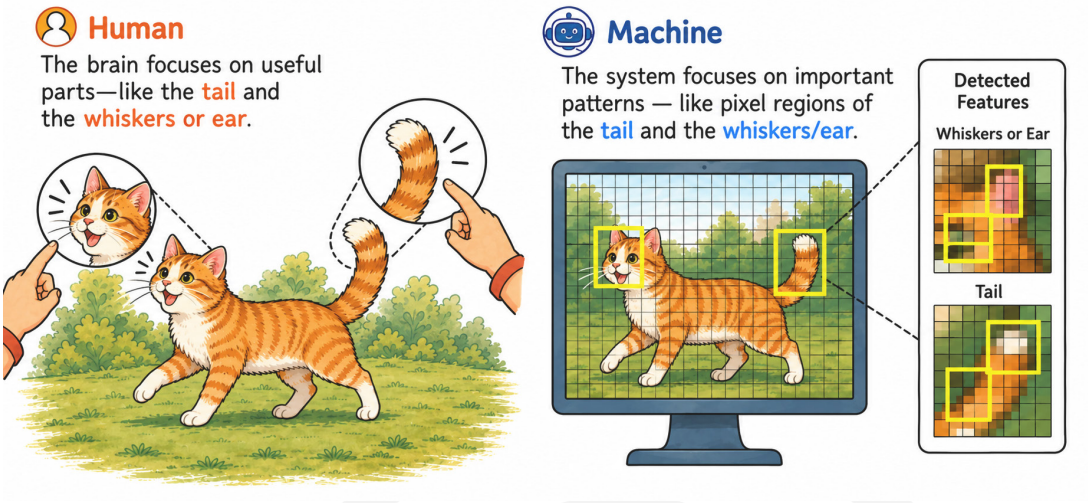


Fig. 2.4: Noticing important features: how humans and machines do it

- **Label:** A label is the right answer attached to an example. 'This sound is the letter A.' 'This pen is blue.' 'This furry animal is a cat.' Labels are how a learning system knows which answer is associated with example.

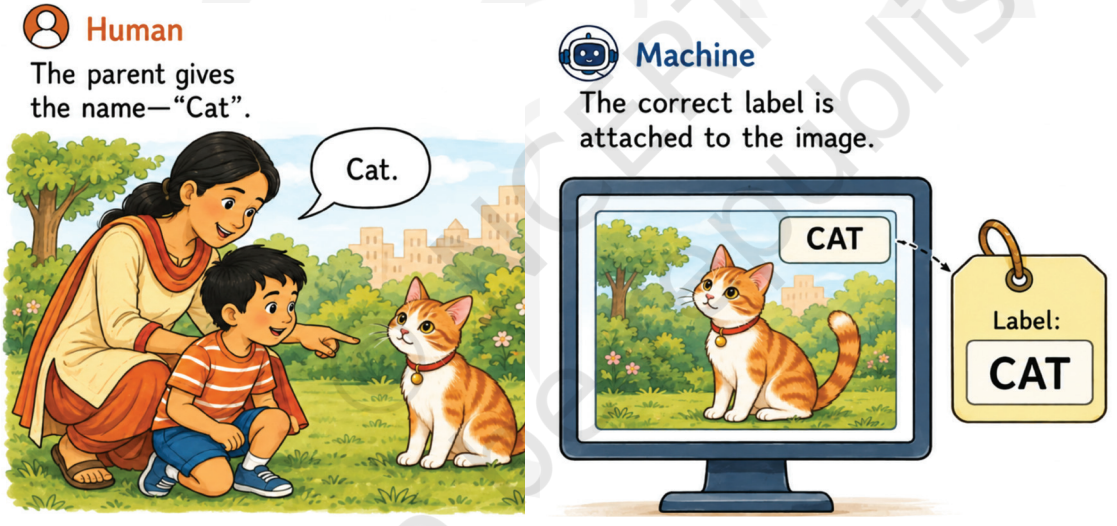


Fig. 2.5: Giving it a label: how humans and machines do it

- **Category:** A category is one of the groups you are sorting examples into. 'Cat' and 'dog' are categories. 'Spam' and 'not spam', the unwanted messages versus the genuine ones, are categories. The labels indicate which example goes into which category.

These four, 'example', 'feature', 'label', and 'category', interlock neatly. An example carries features; a label is attached to it; and the label tells you which category it belongs to. Because they fit together this way, we can line them up side by side and compare how each one looks for you and for a machine.

Table 2.1: The four building blocks, in the brain and in a learning system

Term	Your version	A learning system's version
Example	One cat you saw in a lane	One labelled image fed to the system
Feature	Whiskers, pointed ears, tail	Pixels, edges, colour patches
Label	"Cat", the word an adult attached to that animal for you	"Cat", the tag attached to the image in the training set
Category	Dog versus cat, as you hold them in your mind	Dog versus cat, the groups the system can output



What You have Learnt So Far

An example is one item the system learns from.

A feature is a detail you pay attention to in it.

A label is the right answer attached to it.

A category is one of the groups things get sorted into.

A pattern is the regularity that features reveal across many examples.

These ideas are used by many learning systems, and even by our brain, by every learner, including you.

And what do the building blocks reveal? A pattern. When a system sees many examples of the same category, it begins to notice a regularity, a recurring combination of features that tends to go with a given label. That regularity is called a pattern. Think of the letter 'अ' ('A' in Devanagari). It has a pattern: a vertical stroke on the right, a particular curve on the left. Across many handwritten samples the shape looks slightly different each time, but the underlying pattern stays the same. Finding that steady regularity inside many varied examples is exactly what a learning system does, and it is what lets it recognise a new example it has never seen before.



Try this Now

Pick something you can sort into categories: for example, your pens, your books, the students in your class, or the kitchen utensils.

- What categories did you sort them into?
- List two or three features you actually used to decide which category something belongs to.
- Are your features the same as your friend's?
- If you both sorted the same set of pens by colour, would you draw the line between "blue" and "green" in the same place?



Pause and Discuss

A learning system has been trained on photos labelled 'cat' and 'dog.' Someone shows it a photo of a fox. What do you think happens? Which features might the system use to decide, and which category, cat or dog, do you think it would pick?

2.4 What Kinds of Data: Training, Validation, Test, Live

So far we have said that a learning system learns from 'examples'. But do all examples play the same role? They do not. The same word, data, covers four different uses, and the people who build learning systems keep these uses separate. Let us look at each of them.

- **Training data** are the examples a learning system is shown so that it can find patterns. When Karthik (a student you will meet properly in Section 2.9) took twenty photos of his blue pen, twenty of his red pen, and twenty of his eraser, and labelled each one, those sixty photos were his training data. The system never sees a label on its own. It sees an example with a label attached, then another example with another label, and slowly works out which features tell the categories apart.
- **Validation data** are a separate set of examples kept aside to check how well the system is learning, while it is being trained. From time to time during training, the people building the system evaluate it on the validation data, to see how it is doing, and make small adjustments to the model, so that it learns better. Validation data is like the rehearsal audience for your school play: helpful and honest in its feedback, but not the final show.
- **Test data** are examples kept hidden from the system until the final evaluation. They are used only after training is complete, to judge how well the system has actually learned. If the system does well on training data but badly on test data, it has memorised the pattern rather than learning the patterns. This kind of failure is called overfitting, and we will return to it soon. A good rule is never to use test data for training and never to use it twice, just as your teachers never set the same exam paper twice.
- **Live data** is what the system meets once it is deployed in the world. It is rarely as tidy as training, validation, or test data. The pen Karthik's friend plans to give him next month, in a different colour and a different brand, is live data. The new accent a voice assistant meets in a new region is live data.
- **Performance** on live data is what matters in real use. It rarely matches performance on test data exactly, and when the live data becomes very different from the training data, the system can fail in ways the test data set

What is a Learning System?

A learning system is a computer program that is shown many examples and learns from patterns, rather than following rules a person wrote in advance. The examples during training shown are called its training data. The pattern it ends up with is called a model. Everything the system can do afterwards comes is based on those examples, which is why this unit keeps returning to what was in them.

could never have predicted. In Unit 3, we will meet this problem under the name distribution shift.

Karthik's pen-classifier project, by data kind: Let us see what each kind of data looked like in his project.

Table 2.2: The four kinds of data in Karthik's pen-classifier project

Kind of data	What Karthik used	What it was for
Training data	20 photos of the blue pen, 20 of the red pen, and 20 of the eraser, all labelled, taken on a table in good light	The system learns from these; the labels tell it which is which.
Validation data	5 more photos of each, kept aside during training	To check how the system was doing during training, and to adjust settings
Test data	5 photos of each, taken on a different day in different light, kept hidden until the end	Used once, after training, to judge how well the system actually learned
Live data	Anything Karthik or a friend later shows the trained classifier: a new pen, a different eraser, or an unrelated object	What the system meets in real use, after the project is over



What You have Learnt So Far

The same word, data, covers four different things: training data, validation data, test data, and live data, each with its own purpose.

Performance on test data is a careful estimate of how the system will do on live data, but it is not a guarantee.



Group Work

Think of any learning project you can describe in one line, a face-unlock system, a movie recommender, a handwriting reader. For your chosen project, name what its training data, validation data, test data, and live data might be.

Compare your answers with a classmate. If they differ, discuss why.

2.5 Why Training Data Matters

Whatever a learning system ends up learning, it knows because of the examples it was shown. Show it ten thousand photos of cats and dogs, all from the same kind of household, in the same light, taken from the same angle, and it will learn to tell those particular cats from those particular dogs, in that light, from that angle. Show it the same kind of cat in a different light or from a different angle, and it might get the answer wrong.

This is true for us too. If the only dogs you met growing up were the street dogs in your neighbourhood, the first time you saw a very small fluffy pet dog, a Pomeranian, perhaps, you might have hesitated. It was so different from your mental picture of 'dog' that you had to do extra thinking to recognise it

© NCERT not to be republished

as one. Your mental picture had been shaped by the dogs you had seen. You might say you made a mistake because it was an unusual-looking dog.

Mistaking an unusual case is an honest mistake, but it is still a mistake. For learning systems, this kind of mistake has a name.

Bias in a learning system, in plain language

If the examples a system was shown were not varied or representative enough, say, many examples of one kind and only a few of another, the system may become very good at recognising what it saw often and may struggle with what it saw rarely. We call this kind of imbalance bias. Bias here is not a moral judgement on the system; it is simply a sign that the system learned from unbalanced examples.

MYTH	REALITY
If a system was trained on a million examples, it must have seen everything. Surely it works for everyone.	A million examples is a lot, but if they all came from the same kind of source, the system has really just seen the same kind of example a million times. It can still fail on a kind it has never seen. 'How many' matters less than 'how varied.'

Let us now see how bias is visible in your own learning. Skills learned in one situation do not automatically carry over to a new one. For example:

- If the only weather you ever cycled in was dry, you may be unprepared the first time you ride in heavy rain.
- If you have only ever practised one type of question, a new type may confuse you.
- If you have only ever played cricket on a flat, even pitch, the day you face a turning pitch, where the ball grips and changes direction, will surprise you.

In each case, the dip in your performance is understandable and honest, just like the system's. And like the system, you may improve once you meet more variety.

Rohan in Bhubaneswar: Rohan is in Grade 10. His mother teaches a small group of young children in their neighbourhood. Watching her, he noticed something. A child who has only ever practised one kind of word problem, say, trains running between two stations, will answer confidently but wrongly when the next problem is about buses instead. The only pattern the child had learned involved trains.



What You have Learnt So Far

What a learning system knows is shaped by the examples it was shown, just as your own judgement is shaped by what you have seen.

Bias appears when a system learns from limited a range of examples. It is not a moral judgement on the system, and it can often be reduced by giving the system a wider range of examples.

Rohan realised the child needs not just more problems, but more varied ones. We will meet Rohan again in Unit 5, where he uses an AI tool to help his mother generate exactly that: more varied problems for her evening class.

AI in India

IMD weather and rainfall forecasting: Weather forecasting is a good example of how the real system combines more than one approach. The India Meteorological Department (IMD) uses observations (temperature, pressure, humidity, wind), satellite imagery, and weather-station data, all fed into physics-based models that simulate the atmosphere. Statistical methods refine the output, and increasingly, AI and machine-learning tools help with post-processing, pattern recognition, and local forecast interpretation. The result is the forecast you see on a weather app every morning.

When forecasts go wrong, as they sometimes do, the reasons often connect to the ideas in this section. If the monsoon arrives a week earlier than it has in the past twenty years, the historical patterns may no longer fit the new condition. If a small area near a hill has a weather pattern, the available satellites do not capture clearly, the data may be incomplete. This does not mean the forecasts are useless, only that they get more accurate as the data, models, and methods behind them improve.

Models and Simulations

A model is a simplified version of something real, built so that we can study it or compute with it. A globe is a model of the Earth. An equation for how fast a ball falls is a model of gravity. Models leave out details on purpose; a good model keeps exactly the details that matter for the task.

A simulation is a model set in motion. Instead of only describing a system, a computer steps the model forward through time to see what would happen next: if the air here is this warm and this humid today, where does the wind move, and what does tomorrow look like? Weather forecasts are produced in exactly this way. The IMD's physics-based models of the atmosphere are run forward as simulations, and their output is then refined with statistics and, increasingly, with AI. One caution about the word. In machine learning, model also has a second, related meaning: the trained system itself. The internal pattern a learning system builds from its examples is called a model too. When Karthik trained his pen classifier in Section 2.9, the thing he trained was a model. Both meanings share the same idea, a simplified stand-in for something in the world, built to be computed with.

2.6 Categories are human constructs, with consequences

Where do categories come from? Not from the world: the world does not arrive in pre-sorted categories. Categories come from people deciding which groups suit the task. ‘Cat’ or ‘dog’ is a choice. ‘Spam’ or ‘not spam’ is a choice. ‘Safe to drive’ or ‘unsafe to drive’ is a choice. Nature did not make these groupings; people did.

This matters because some examples do not clearly belong to one category or the other. Whether a message counts as ‘spam’ or ‘not spam’ depends on what a particular person finds unwanted, and people do not always agree. When a case is hard to classify, it shows that someone, at some point, decided

where one category ends and the next begins, and the system simply follows that decision. Think of an object whose colour is somewhere between blue and green. Two people might label it differently, not because the object is unclear, but because each of them decided the cut-off point differently.

Try this Now

- Find five things at home or in your school bag that are clearly blue.
- Find five things that are clearly green.
- Now find one thing that is hard to classify as blue or green, such as a teal pen, a turquoise-coloured dupatta, or a blue-green book
- Get three of your friends to sort the hard one into blue or green. Do they all agree?
- If a learning system had been trained on the clearly blue and clearly green ones, what label might it produce for the hard one?
- Does its answer make sense or does it lack proper reasoning?

A UPI safety check may apply a rule such as: 'a transfer above a certain amount, to a new recipient, should be held temporarily for verification.' If the rule is too strict, many genuine transfers get delayed; if it is too lenient, fraudulent ones slip through. Both kinds of error are possible, and someone has to decide which type of error is more acceptable than the other.

To see those errors clearly, imagine a classifier whose only job is to look at a photo and decide 'dog' or 'not dog.' There are four possible outcomes.

Table 2.3: The four possible outcomes of a 'dog or not dog' classifier

	Actually a dog	Actually not a dog
System says 'dog'	True positive (correct)	False positive (a cat classified a dog)
System says 'not dog'	False negative (a dog missed)	True negative (correct)

A false positive is when the system places something in a category it does not belong to, the classifier calling a cat a dog, or a UPI check holding a perfectly normal payment to a relative. A false negative is when the system misses something that does belong to that category, the classifier calling a real dog 'not dog,' or a fraudulent transfer slipping through the UPI check. A stricter line catches more fraud but delays more genuine cases; a more lenient line delays fewer genuine cases but misses more fraud. Real systems are designed to reduce whichever error would cause more harm.

What You have Learnt So Far

Categories are human choices, not facts in the world.

Edge cases, the ones that do not clearly belong to either category, show that someone decided where one category ends and the next begins.

Wherever a system sorts people or things, that decision shapes who is helped and who is delayed, and in high-stakes tasks, like medical screening or a loan decision, those outcomes matter much more.

© NCERT not to be republished

So far, the data we have looked at has been neat and well organised. Real-world data are rarely like that; they are often messy. What does 'messy' mean here? Consider a school register with rows for each student's name, age, height, and attendance. A teacher keeping it by hand may use their own shorthand, and if someone else reads it, they may find inconsistencies: dates written in different formats, ages as '12' in one row and 'twelve' in another, missing entries, heights in different units. Before this register could be used, each of these problems would need to be spotted and fixed. Table 2.4 shows a few of them, the problem with each, and how it was cleaned up.

Table 2.4: Examples of messy data and how each was cleaned up

Row in raw register	What's wrong?	After cleaning
Asha, 14, tall, 78%	Height is a description, not a number	Height entered as 152 cm after follow-up
Ravi, 14, 145, 102%	Attendance above 100% is impossible	Attendance corrected to 92% after re-checking
Karim, 14, 150, blank	Attendance value missing	Attendance entered as 87% after re-checking the register
Priya, 14, 154, 88%	Looks fine	Kept as is

Try this Now

Let us try inspecting some messy data: Imagine a classroom register with 30 rows, recording each student's first name, last name, age in years, height in centimetres, attendance percentage, and favourite subject. Glancing through, you notice: three rows have height entered as a word or symbol rather than a number ('tall,' 'medium,' '-'); four rows have favourite subject left blank; two rows show attendance above 100 percent; and two names are spelled in ways you have never seen.

Suppose this register were fed, exactly as it is, to a learning system meant to predict which students might need extra support.

- What kinds of mistakes would the system likely make because of the messy entries?
- The system would need a cut-off somewhere, for example, 'flag any student with attendance below a certain level.' How might the incomplete attendance values (the ones above 100%) affect that cut-off?

Discuss with a classmate.

The point of this activity: in every real learning project, the data starts out messy and usually needs a careful check and a clean-up before it can be fed to the system.

2.7 Where the Parallel Breaks: Recognising is not Understanding

So far in this unit, we have stayed close to the comparison between the human way of learning and the machine way of learning. You learn from examples; a learning system learns

from examples. You build patterns; it builds patterns. You are shaped by what you have seen; so is the machine. Up to here, the parallel holds. Now we will see where it breaks or deviates.

- **A learning system that has correctly classified ten million cat photos does not actually know much about cats.** It does not know that a cat is a living being. It does not know that cats catch mice and purr. It does not know that cats live for about fifteen years, or that they were once worshipped in ancient Egypt. The system has built a mathematical pattern that, when shown a new photo, produces 'cat' often enough to be useful and that is all. 'Cat', in the system's world, is a label that frequently goes with a certain pattern. It is not a creature the system has ever met.

This is the most important thing to carry from this unit. Pattern recognition is a task that can be done well or badly. But understanding what the pattern is, what the task means, and why it matters, is a different thing altogether. The two are not the same, and a system that recognises a pattern does not, simply by recognising it, understand it.

So when you read that an AI system 'understands' something, ask carefully 'what does it mean?' Often it means that the system can produce outputs which, to a human reader, look like understanding. Producing such outputs is genuinely useful, but recognising patterns is not the same as understanding what they mean. These are two different tasks.

- **Two more differences between us and machines.** First, we have a body and a life; the system has neither. Our sense of 'hot' includes the memory of the time we burned a finger on a pressure cooker. A system's representation of 'hot' comes from patterns in text, such as which words appear near 'hot' and in what contexts, not from anything felt. Second, we can change our minds from a single example: if a friend convinces us in one conversation that our idea of success was too narrow, we adjust our thinking on the spot. A learning system usually needs many new examples and a fresh round of training before it can shift at all.
- **Memorising is not learning, either.** A system can be made to recognise its training examples perfectly while doing badly on anything new. In that case it has memorised the examples but not learned the underlying pattern. This failure has a name—overfitting. Overfitting is like cramming for an exam, the system remembers exactly what it was trained on but struggles with anything unfamiliar. You will notice these distinctions, memorising



What You have Learnt So Far

Recognising a pattern is not the same as understanding what the pattern means. Memorising is not the same as learning. Both you and a machine can recognise without understanding, and memorise without learning.



Computational Thinking Lens

In this unit you have practised two habits of computational thinking: pattern recognition and abstraction. You progressed from many specific examples, such as the cats you met as a child, to a general idea, 'cat'. You also saw that both you and a learning system are shaped by the examples you were shown. Going from specific examples to a general idea or rule is what we call abstraction, and in this unit both you and the learning system have been doing exactly that.

versus learning, recognising versus understanding, returning again and again whenever we discuss learning systems. They are worth holding on to.

2.8 Three Ways Machines Learn from Examples

At this stage, your goal is simply to recognise the names and the basic differences between three ways a machine learns from examples. As you go deeper into this area later, you will understand each more fully. Did you notice something about the examples in this unit? The cat photos, the pen photos, and the translated sentence pairs, all came with the right answer attached. Someone wrote 'cat' next to a cat photo. Someone labelled Karthik's pen 'blue pen'. Someone confirmed that a particular Marathi sentence and a particular Hindi sentence meant the same thing. This is one of the three main ways machines learn from examples, the most common kind in everyday AI, and the one this unit has mostly described. But it is not the only way. Let us look at all three.

(i) Supervised Learning: Learning from Labelled Examples

When a learning system is trained on a data set where each example comes with the right answer attached, it uses the labels to work out which features tell one category from another.

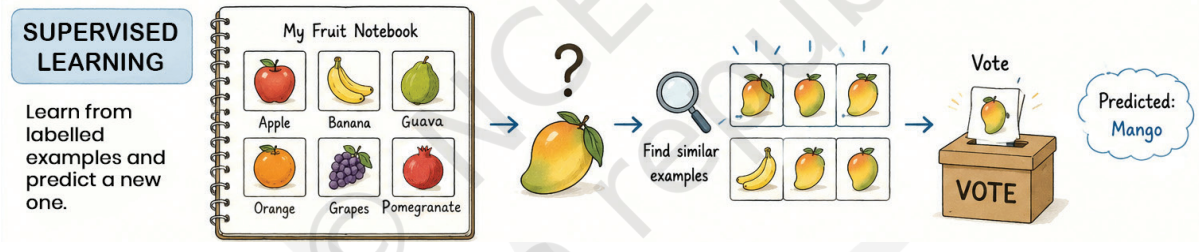


Fig. 2.6: Supervised learning–learning from labelled examples

Karthik's pen classifier was doing exactly this. He showed it 20 photos of his blue pen, each labelled as 'blue pen,' and twenty of his eraser, each labelled 'eraser'. From the photos and the labels, the system worked out which features separate a blue pen from an eraser. An email or SMS spam filter works the same way—every message in its training set was marked 'spam' or 'not spam' by a person before the system ever saw it.

Phone face unlock works a little differently. The face-recognition model is first trained on many labelled examples of human faces. Then, when you set up the feature on your own phone, it stores a mathematical template of your face. Each time you unlock the phone, it compares your live face with that stored version. Bhashini does the same kind of thing on a much larger scale, with pairs of sentences that have each confirmed as translations of one another.

A simple method of supervised learning: One of the simplest supervised-learning methods is called k-nearest neighbours (KNN). The easiest way to understand it is to try it out.

Imagine you have a notebook with the measurements of 20 animals you have seen—10 cats and 10 dogs. For each one, you noted two numbers: its weight in kilograms and the length of its ears in centimetres. Cats tend to be lighter, with shorter ears; dogs tend to be heavier, with longer ears, though some animals fall in between, as real examples often do. Now someone describes a new animal: it weighs 12 kg and has ears of 10 cm long. Is it a cat or a dog?

Here is the KNN method, with K set to 3. Look up the three animals in your notebook closest to (12 kg and 10 cm) on those two numbers. Two of the three nearest are dogs and one is a cat. The majority wins, so you label the new animal 'dog.' Here, K is simply the number of nearby examples you decide to check. A larger K means the system consults more examples before deciding.

What you just did has a name. By deciding on the steps to follow, looking up the three closest animals, counting them, and taking the majority, you followed an algorithm—a clear set of steps for a task. There were no equations and no separate training stage; your notebook of 20 animals was the model.

KNN works because, in many real situations, similar things tend to share a label, that is things alike in their features often belong to the same category. But it has limits. What counts as 'similar' depends on which features you compare, and choosing the right features is most of the work. KNN is worth knowing because it shows that machine learning does not always begin with complicated mathematics. The general rule is simple—to classify something new, look at the most similar examples you have already seen, and let the examples vote.

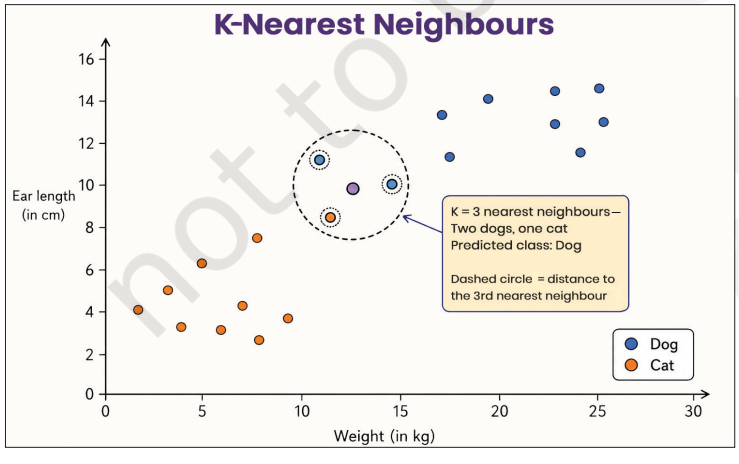


Fig. 2.7: K-nearest neighbours with K = 3

Supervised learning is a category, not a single technique. It ranges from KNN, simple enough to run by hand, to deep neural networks with millions of parameters. But the path is always the same: the system learns from labelled examples, then uses what it learned to label new ones.

(ii) Unsupervised Learning: Finding Structure without Labels

Now let us look at a different kind of learning system, one that is shown examples but not told the right answer. It has to find the pattern on its own: noticing which examples resemble one another, what naturally groups together, and what repeats.

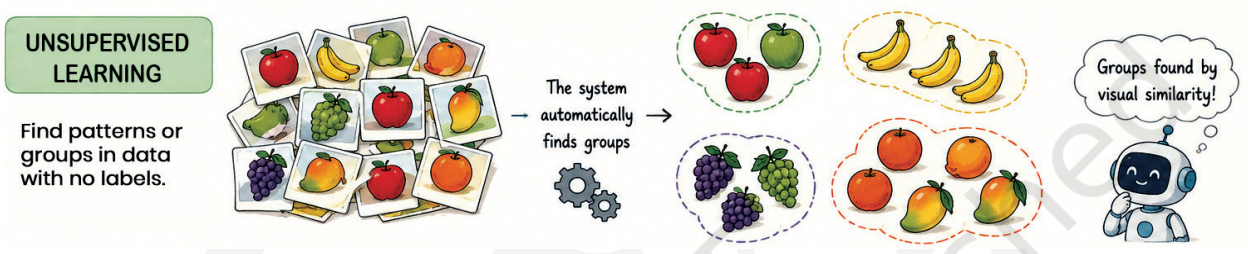


Fig. 2.8: Unsupervised learning–finding groups without labels

A photo-gallery app that automatically groups your photos by who appears in them is doing this. Nobody labelled the photos as 'this is your mother' or 'this is your friend.' The system simply grouped photos that share facial patterns, and it lets you add names afterwards if you want to. A shopping app does something similar when it notices that people who buy schoolbooks often also buy geometry boxes, pens, and pencils; no one taught it that these things go together, it found the pattern in what people bought. A news app that gathers articles about the same event into a single 'story,' without being told which articles belong together, is also doing a form of unsupervised learning.

In each case, the system is asking: how is this collection of examples organised, beneath the visible details?

Real learning systems often combine supervised and unsupervised approaches. For example, a recommendation engine, the kind of system that suggests things you might like, such as a shopping app, an OTT app suggesting films from your past choices, or a podcast app pointing you to similar shows—might use different learning methods to find similar users and supervised methods to predict what each user will click next. These categories are useful for sorting methods, but they are not strict rules.

Let us return to the notebook of cats and dogs, but this time without labelling the animals. The notebook now holds just 20 pairs of numbers (weight and ear length), and nothing else.

An unsupervised system plots these on a graph and looks for natural groupings. It might find two clusters: one in the lower left (the cat data, though the system does not know that name) and one in the upper right (the dog data). The system has discovered that there are two groups. It has not discovered the names 'cat' and 'dog', choosing those is still our job. Clustering reveals the structure; naming it remains human work.

(iii) Reinforcement Learning: Learning by Trying

Now consider the third way a system learns. Here the system tries different actions, gets feedback on them, and updates itself based on that feedback. The feedback is not 'here is the right answer.' It is a signal that says, in effect, 'that worked' or 'that did not work.' Sometimes the signal comes at once, and sometimes only several actions later.

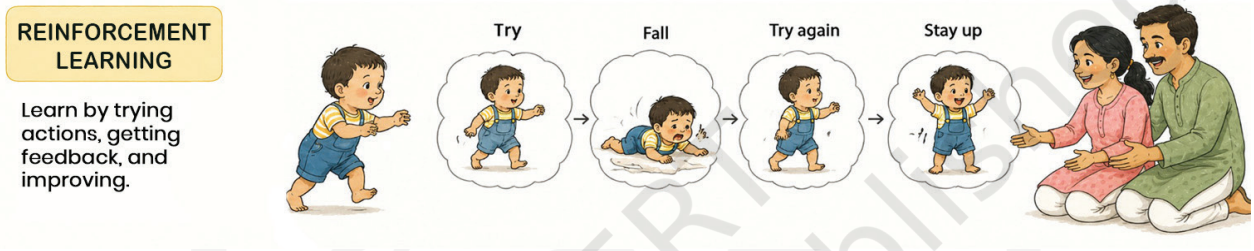


Fig. 2.9: Reinforcement learning–learning by trying

A baby learning to walk learns this way. It takes a first step and falls; it tries again and stays up for half a moment before toppling; the next attempt carries it slightly further. Over many tries, the baby develops a sense of how to balance, how to lift a leg, how to use its arms, and when to take the next step. Nobody hands the baby a labelled set of 'good steps' and 'bad steps.' The experience itself provides the feedback, whether it is staying upright or falling, a parent's encouragement, or reaching the toy across the room. Without being taught the lesson explicitly, the baby learns to walk.

Reinforcement learning systems do something similar. A program that learns chess by playing thousands of games against itself is doing reinforcement learning. Every game ends in a win, a loss, or a draw, and the system uses these outcomes to adjust which moves it favours. A music recommender that changes its suggestions depending on whether you finished a song or skipped it after a few seconds is doing a version of the same thing—the action (suggesting) produces feedback (finished or skipped), which shapes what it suggests next.

For both the baby and the machine, the feedback often comes late. A baby's first steps do not immediately reach

its parents. A move in chess may not prove decisive until ten moves later. In situations like these, the system has to work out which of its many actions actually contributed to the outcome. This delayed-reward problem is one of the harder challenges in machine learning, and a great deal of research has gone into it.

Table 2.5: The three ways machines learn, compared

	Supervised	Unsupervised	Reinforcement
What the system sees	Examples, each with a label (the right answer)	Examples, no labels	Observations from an environment, plus feedback (good or bad)
What it learns to do	Predict the label for new examples	Find structure among groups, similarities, or patterns	Choose actions that lead to good outcomes over time
Where you meet it	Spam filters, facial recognition unlock	Photo grouping, 'you may also like' suggestions	Game-playing systems, robot navigation, different types of recommendations



Pause and Discuss

For each of the systems below, which of the three kinds of learning do you think is most likely at work?

- A spam filter.
- A music app that gives 'you may also like' suggestions.
- A video-streaming app's 'what to watch next' queue.
- A photo gallery that groups your photos by who is in them.
- A chess engine that has learned by playing against itself.

Discuss with a classmate. Notice where the answer feels clear and where it feels mixed.

Many everyday AI systems you meet have a supervised-learning part. Some also use unsupervised or reinforcement-style methods, and many real systems combine more than one approach, so a single clear answer is not always possible.

One more thing worth saying: A careful student might ask two questions at this point.

- **First:** The picture in Section 2.1 (Fig. 2.2) showed neural networks, layered processing, hidden layers, mathematical units. Are KNN and the other methods we discussed the same kind of thing? No. Neural networks are one family of models, used across several kinds of learning; KNN is a different and simpler family. The picture in Section 2.1 was about how neural networks were loosely inspired by biological neurons, not about every learning system. Many learning systems do not look like that picture at all.

© NCERT not to be republished

- **Second:** Are these three categories, supervised, unsupervised, and reinforcement, the only ways machines learn? No. There are others. One important one is called self-supervised learning, and much of modern AI relies on it. The large language models behind tools you may have used, such as ChatGPT, Gemini, Claude, Grok, and Perplexity, learn this way. They read enormous amounts of text and learn by predicting the next token from the tokens before it. Nobody manually labels each example; the structure of the text itself provides the signal. The three ways we covered are a good starting point, but the field of machine learning is larger than these three.

Stretch: How Self-supervised Learning Works

Optional. Mainly for Grade 12 readers, but anyone curious is welcome.

Supervised learning needs someone to attach a label to each example, which is costly in human effort when you have millions of examples. Unsupervised learning needs no labels, but on its own it does not produce a system that can answer specific questions. Self-supervised learning is a clever middle path where the system finds its own training signal in the structure of the data, with no one manually labelling anything.

Take text. A sentence is a sequence of words, or, more precisely, tokens. Cover up the next token and ask the system to predict it, then check whether its guess was right; if not, it adjusts. Repeat this across billions of sentences from the internet, books, and articles, and after a while the system has learned the patterns in language—which pieces of text tend to follow one another, in which contexts, and with which meanings.

This is how many large language models you may have used are trained. Nobody manually labelled 'what the next token should be' for billions of sentences; the data labelled itself, because the structure of language already provides the signal. The same idea works for images (cover a patch, predict it), audio (mask a segment, predict it), and many other kinds of data.

Why this matters: much of the progress in AI in recent years has come from self-supervised learning, going beyond the three categories you met in this section. Those three remain a useful starting frame, and self-supervised learning is an important foundation of many modern AI systems.

2.9 Train and Break: Your First AI Experiment

Now it is time to try out what you have learnt so far. This is the unit's main hands-on activity. The main version needs nothing but everyday objects you already have. If your classroom has computers or tablets, there is also an optional version that uses a free image-classifier tool. Your teacher will decide which version your class can run. Both versions teach the same lesson.

What You have Learnt So Far

Machines learn from examples in three main ways: supervised (with labels), unsupervised (without labels), and reinforcement (by trying and getting feedback). Many everyday AI systems have a supervised part, some also use unsupervised or reinforcement methods, and real systems often mix more than one. These three are a good starting point, but not the whole field. Self-supervised learning, used in today's large language models, is another important approach.

What is a Token?

A token is a small piece of text that a language model works with. Sometimes a token is a whole word. Sometimes it is only part of a longer word, and it can also be a number, a punctuation mark, or even a space. Humans usually think in words; language models process text as tokens. In this module we usually say 'word' for readability, but technically most AI language models predict tokens one after another. Different models split text into tokens differently, which is one reason the same sentence can come out as a different number of tokens in two different models.

Before we lay out the activity, here is what one student's experience looked like.

Karthik's afternoon. Karthik is in Grade 10 in Visakhapatnam. His robotics-club teacher told the class to spend the afternoon training a no-code image classifier on three things they had at their desks. Karthik picked his blue pen, his red pen, and his eraser. He took twenty photos of each, from different angles and distances, and trained the classifier. It worked beautifully. When he held up the blue pen, it said 'blue pen.' When he held up the eraser, it said 'eraser.'

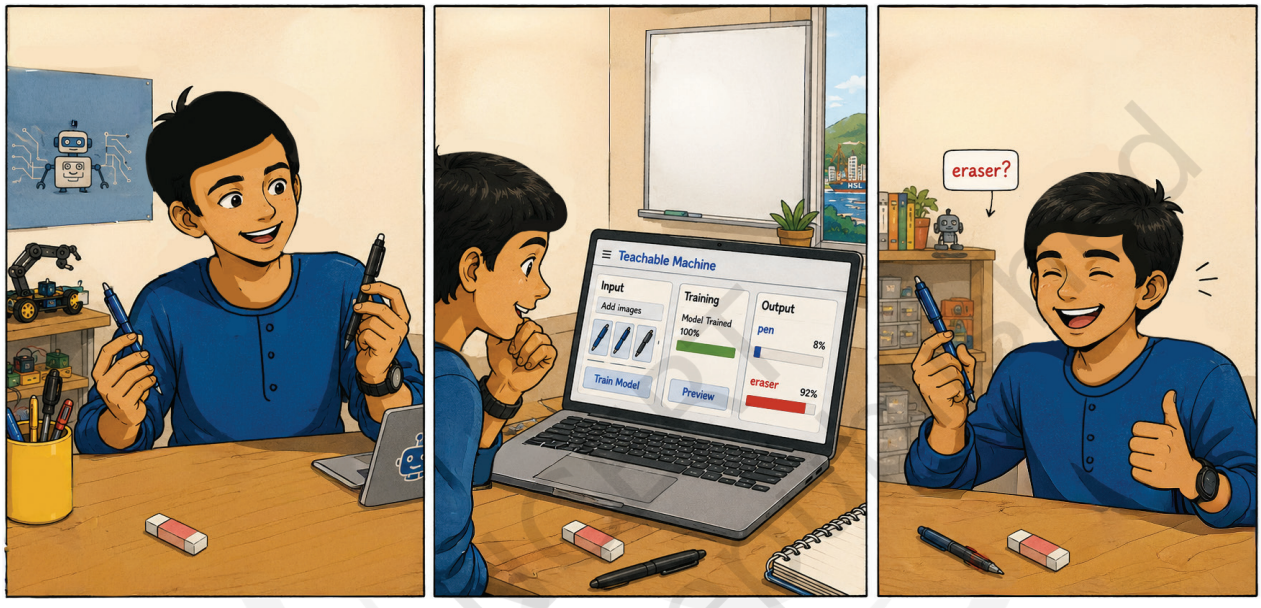


Fig. 2.10: A classifier trained on one pen meets a friend's pen

Then he tried to break it. He held up his blue pen against a white wall instead of his green desk—it still worked. He held up his blue pen sideways—still worked. Then he held up a friend's blue pen, a different brand than his own. The classifier said 'eraser.' Karthik laughed and tried again. Still 'eraser.'

Why? All twenty of his blue-pen photos had been his own pen. The classifier had not learned the concept of 'blue pen.' It had learned 'Karthik's pen.' A different brand, a different shape, different writing on the side, to the system, that was no longer a pen. To Karthik, who knew what a pen was, the same kind of thing in a different brand was obviously still a pen. To the system, which does not understand what a 'pen' was, the difference was everything.

This is the whole unit, in one afternoon. A learning system can be very good at the examples it was trained on and not helpful in the cases you actually want it to handle. One important way

to improve is to use more varied training data. But the lesson, that what you feed a system determines what it learns and has to be experienced to be understood. Now you get to experience it through the next activity. Karthik's afternoon will return in the next unit, where the lesson he drew, 'it looked confident and was wrong', turns out to be the central problem of trusting AI outputs.

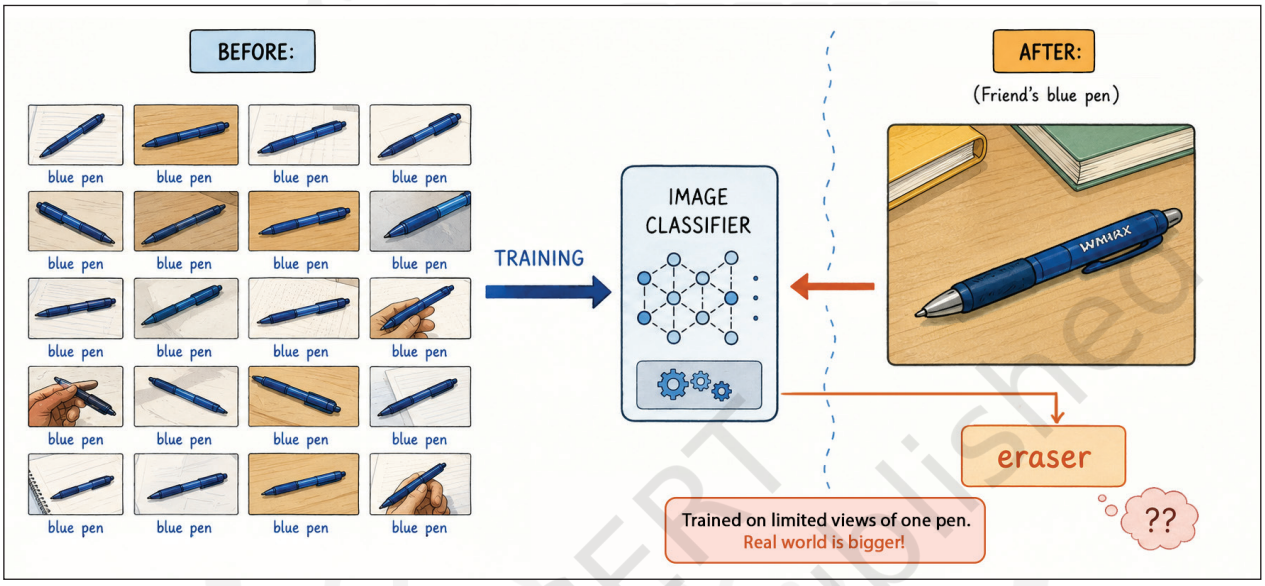


Fig. 2.11: A learning system

A note before you start. Use everyday objects, such as pens, erasers, books, water bottles, fruits, and the like. Do not use people's faces, classmates, Aadhaar or ID cards, or any personal documents as the things being classified. This is true for every version of the activity, and it is a good habit to build from the start.

The main activity: The Human Classifier (unplugged)

The whole class, about 20 minutes. No devices needed, only objects you already have.

- **The idea:** A small group of students will act as a 'learning system.' Like a real classifier, they will learn only from the examples they are shown. The class will then test whether they truly learned the concept, or only the examples they were shown.
- **Set up:** Choose three or four students to be the 'classifier' and send them to the front. Pick one everyday category to train them on, for example, 'pen.' The rest of the class is the trainer.

- **Train:** The trainers show the classifier group eight to ten examples of pens, but deliberately make them all similar, all blue ballpoint pens of the same kind, say. Each time, the class says clearly, 'this is a pen.' Show a few non-pens too, an eraser, a ruler, saying 'this is not a pen.' The classifier group may only watch; they cannot ask questions.
- **Test:** Now show the classifier group new objects one at a time and ask them to call out 'pen' or 'not a pen,' with no help from the class:
  - another identical blue pen (they will likely say 'pen', correctly);
  - a fat marker, a fountain pen, a pencil, a very differently shaped pen (here they will hesitate, or get it wrong);
  - something pen-shaped that is not a pen, like a thin torch or a chopstick (they may wrongly say 'pen').

Discuss

- Where did the classifier group get it right, and where did it slip?
- They were trained only on one kind of pen. What did they actually learn, 'pen,' or 'blue ballpoint pen'?
- This is exactly what happened to Karthik: his classifier was trained on 'Karthik's pen,' not 'pen.' What training examples would have helped it learn the broader idea of a pen?
- How is this like, and unlike, the way you yourself learned what a pen is?

The lesson: A learning system, human or machine, learns from the examples it is given. Narrow examples produce narrow learning. To learn the general idea of 'pen,' it needs varied pens, just as Karthik's classifier needed photos of many different pens, not twenty of the same one.

Optional: Train and break it on a computer

A note on safety for this version. Because this version uses a camera, take extra care. If one of the objects you photograph is a phone, turn it off first, so that no notifications, messages, photos, or app screens appear in the picture. Photograph objects only, never faces, classmates, Aadhaar cards, ID cards, or personal documents. The lesson works just as well with everyday objects.

For classrooms with computers or tablets and internet access. Same lesson, with a real classifier.

If your class has devices, you can run Karthik's experiment for yourself using a free, no-code image-classifier tool that works in a web browser. Your teacher will point you to the current tool and a ready-made demo using the QR code provided within this module.

In pairs, pick three objects from your desk, pens, erasers, books, a water bottle, fruits, or a switched-off phone with a blank screen. Take about 20 photos of each, then train the classifier and test it. Then deliberately try to break it, just as Karthik did: a different brand of pen, a different-coloured eraser, a friend's book. Notice what worked, what failed, and why. (Follow the safety note above: objects only, no faces, or personal documents)

After the activity. Whichever version your class ran, write three lines in your Curious Notebook—one example that worked surprisingly well, one example that broke surprisingly easily, and one sentence explaining why you think it broke. Save these. We will come back to them in Unit 3.

Reflect

Think back to a time you got something wrong because you had not seen the full picture, a person you misjudged, a place you misunderstood, a song you filed under the wrong kind. Write three lines about it.

- What was your 'training data' at the time, the limited set of examples your judgment was based on?
- Where would more variety have helped?
- What does this tell you about how you might judge the next AI system you meet?



What You have Learnt So Far

What a system learns depends on what the data shows it. If you want the machine to learn something else, you may need different kinds of data to train it.

So whenever you work with a learning system, do not stop once it is trained. Test it with new and difficult examples. Try to find where it fails. Those failures will show you what the system has really learned, and what it has only memorised.

The training-data question is not just academic. Indian languages, Indian faces, Indian handwriting, Indian crops, Indian dialects, each of these needs to be in the training set if a system is to work for the people who use it.

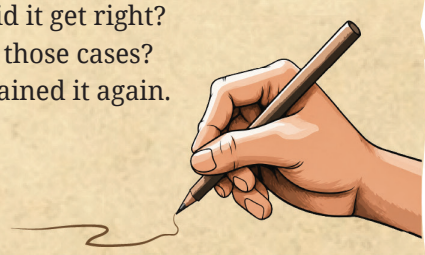
A learning system trained mostly on a narrow group of examples tends to work less well on examples that differ from another. The fix is simple in principle, collect varied examples, but hard in practice: it takes effort, time, and money.

Students like you will, in your working lives, help decide which examples get collected. And the examples you choose to gather will determine who the systems end up working for.

© NCERT not to be republished

For your Curious Notebook

Train-and-Break Notes: Using the activity from Section 2.9, write a one-page note for your Curious Notebook in four parts:
(a) What three things did you or your class classify, and why did you choose them?
(b) Two examples where the classifier worked well. What did it get right?
(c) Two examples where it broke. What was different about those cases?
(d) One sentence on what you would do differently if you trained it again. What kinds of examples would you add?



Assignment

Quick checks

1. In your own words, what is the difference between a rule-based system and a learning system? Give one example of each from your day-to-day life.
2. A friend says "My phone's face unlock learns my face every time I use it." Is this likely to be true? In two lines, what would you check before agreeing?
3. The unit says 'recognising is not understanding.' Give one situation where this distinction matters, where treating recognition as understanding could cause a real problem.

Apply it

4. Bhashini works very well for some Indian language pairs and is still being improved for others. What kind of examples do you think the better-served pairs probably had more of? What would it take to bring the others up to the same level?
5. Imagine you are training a classifier to recognise the handwriting of students in your class. Three classmates volunteer to give you a hundred handwriting samples each. Without doing the activity, predict two ways the classifier might surprise you when a fourth classmate writes a sample. Why might those two things happen?

Stretch (optional)

6. The unit draws a parallel between human learning and machine learning, and is careful to say where the parallel differs. Pick one of the differences described in Section 2.7 and write a half-page note. Do you find that difference convincing? If you do, what does it imply for how much we should trust the systems we build? If you are not convinced, what is your alternative view?

Coming up in Unit 3: You have seen how learning systems work, and where they go wrong honestly. The next unit is about how they can go wrong in ways that mislead us, and what to do about it.

© NCERT not to be republished

UNIT

3

When AI Gets It Wrong

UN514UN03



There is a classic optical illusion you have probably seen, called the Müller-Lyer illusion. Two horizontal lines of the same length are drawn between two pairs of arrows, one pair pointing outward, and the other inward. The line whose arrows point outward, like the tail of an arrow, looks longer. You measure them: they are the same. You look again, and that one still looks longer. You know they are the same length (after measurement) and yet it still looks wrong. Knowing the truth does not make the wrong impression go away.

This is what false certainty feels like. You see something clearly and you feel sure. And you are wrong. You did not lie to yourself; your visual system interpreted the information it had, and the answer it produced was confident but incorrect. The gap between confidence and correctness is what this unit is about, in your own thinking, and in the AI systems you encounter.

Every unit of this module has drawn a parallel between you and the machine, and here that parallel is at its clearest. You produce false certainty; so do machines. The mechanisms are

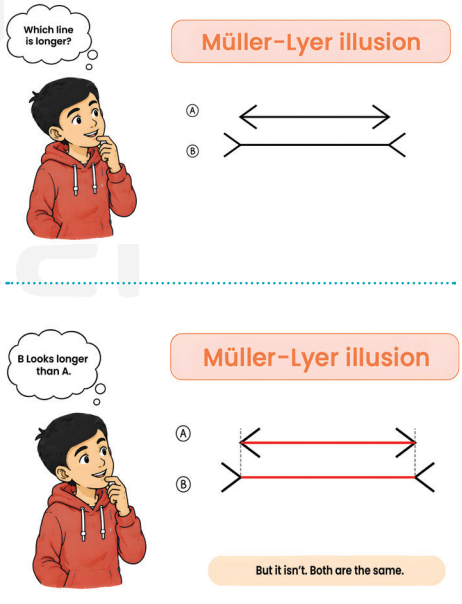


Fig. 3.1: Müller-Lyer illusion

different, but the result is similar from the outside—a confident answer that happens to be wrong. By the end of this unit, you will have a working method for spotting false certainty in your own thinking and in the outputs of AI tools.



Words You'll Meet in this Unit

This unit introduces several technical terms. Each one gets explained when it appears, and the alphabetical glossary at the back has them all. It is worth knowing in advance:

- **Hallucination** is when an AI assistant produces confident, smooth text that contains invented facts, names, or quotations.
- **Distribution shift** is when the kind of input a system meets in real use differs from the kind it was trained on. Performance can drop, often without the system being able to notice.
- **Synthetic media,** such as video, audio, or images, generated or altered by AI rather than recorded or created by people. The AI-made illustrations in this module are synthetic media, and so are many harmless things.
- **Deepfake** is a particular, often harmful kind of synthetic media. It is an AI-generated or manuplated video, audio, or image of a real person, usually made to deceive or shared without that person's consent.
- **Calibration** is whether a system's confidence number actually matches how often it is right. A well-calibrated system's 90%-confidence answers are correct about 90% of the time.
- **Confirmation bias, anchoring, and automation bias** are the three ways your own mind can mislead you when reading AI output. The unit explains each.
- **Verification protocol (the five checks)** is a short, learnable habit for deciding whether to trust an AI output: the source, the claim type, compare, pause, and confidence.

3.1 Confident but Wrong: A Problem You and Machines Share

- **Your version of the problem:** You misremember a song lyric and sing it confidently. You are sure a friend said something they did not, and you can 'recall' the exact words. You wave at a friend across the street, only to realise after they look puzzled, that it was a stranger with the same haircut. None of these is dishonesty. Each is your mind producing a confident answer from imperfect information. The confidence is real; the information was limited.
- **The machine version of the problem:** An AI assistant produces a paragraph of confident, fluent English containing a fact it has invented. A face-recognition system, in poor conditions, bad lighting, a low-resolution image, or an unusual angle, may produce a high confidence score and names the wrong identification. A weather model predicts a sunny Saturday, but it rains. None of these examples is a programming bug. Each is a pattern-based system producing the answer that best fits its examples, and sometimes the best-fitting answer is wrong.

© NCERT not to be republished

- Humans and machines do not work the same way. You have a body, a history, and some sense of when to be unsure; an AI system has none of these. But both of you can make the same kind of mistake, giving a confident answer without enough reason. Once you learn to notice this in yourself, it becomes easier to notice in a machine.
- **Karthik's classifier returns:** In Unit 2, you met Karthik in Visakhapatnam, training a no-code classifier on his pen, his eraser, and a friend's pen. The system told him, with high confidence, that his friend's blue pen was an eraser. The confidence was not necessarily a bug. It was the system matching a new object against the patterns it had learned, and producing whatever fit best, even when nothing fit well. Karthik's classifier and an AI assistant that invents a fake reference work very differently, but they show the same symptom—a confident answer without the kind of understanding behind it. This unit is about recognising that symptom in both.
- **When mistakes carry a heavy price:** Consider a medical screening tool that analyses chest X-rays and flags possible signs of tuberculosis. Recall the concepts of false negative and false positive from Unit 2. Here, a false negative means that the tool says 'looks normal' when disease is in fact present, so a patient who needed further treatment does not get it. A false positive means that the tool marks a healthy X-ray as suspicious, so the patient is sent for an extra test. These two mistakes are not equal in cost. A well-built screening tool is usually tuned to make more false positives than false negatives, because missing a real case of disease is far worse than asking for an additional test. The same logic applies to bank fraud detection, where missing a fraud usually costs more than the inconvenience of an extra check. In every system, the right balance depends on what each kind of error costs.

Try this Now

Think of one moment in the last week when you were sure of something and turned out to be wrong. Write two lines on what you were sure of and what you found out instead.

Now ask: was your sureness based on a lot of evidence, a little evidence, or no evidence—just a feeling?

In retrospect, what would have made you check before being so sure?

What You have Learnt So Far

Confidence is not the same as correctness. Humans produce false certainty too; it is not a machine-only problem. Spotting false certainty in yourself makes it easier to spot in a machine.

© NCERT not to be republished



Pause and Discuss

Sometimes a right answer and a wrong answer feel exactly the same from the inside, both come with the same sense of certainty. If being sure feels the same whether you are right or wrong, how can you tell, in the moment, when you might be about to be wrong? Discuss with a classmate.

3.2 Why models fail

There are a handful of common reasons an AI system produces a confident wrong answer. Each has a name.

- **Hallucination, in language models:** A hallucination is when an AI assistant produces text that sounds confident and fluent but contains things that are not true, such as invented names, false dates, and references to sources that do not exist. We might call this 'lying,' but the system is not lying; it is simply producing fluent-looking continuations from patterns in its training data, and sometimes the best-fitting continuation is an invented fact. The result still misleads anyone who trusts it. Hallucinations are most common when you ask about specific facts the system was not well trained on, such as a small town, a little-known person, or a narrow specialist topic.
- **Distribution shift:** A system trained on one kind of input often fails on a different kind. A face-recognition system trained mostly on indoor photos in good light may struggle with an outdoor photo taken at dusk. A spam filter trained on past spam emails may struggle as spam emails change over time. In these examples, we observe that the system learned one pattern of data, but the real-world pattern has shifted, and the system could not keep pace. Distribution shift happens when the examples the system meets in the real world become different from the examples it learned from.
- **Edge cases:** An edge case is an example that lies close to the boundary between two categories, where the answer is not obvious. Suppose you are training a system to sort mangoes into ripe and unripe. It learns that ripe mangoes usually look one way and unripe ones another. Now, it is shown a mango that is half green and half golden. It does not fall clearly into either group; it sits right on the boundary. That is what makes it an edge case, so close to the line that even a person might pause. You met edge

cases in Unit 2, when you sorted objects into blue and green, you had to ask where a colour like teal belongs. Language has them too. Without punctuation, 'you are coming today' could either be interpreted as a statement or a question, an edge case the system has to resolve. Edge cases make the classification boundary visible.

- **False positives and false negatives:** Recall these from Unit 2. A false positive is flagging something that should not have been flagged; a false negative is letting through something that should have been caught. A face-recognition system that fails to match someone in poor lighting is producing a false negative. A UPI safety check that holds a routine payment to a known recipient is producing a false positive. Real systems make both kinds of error, especially on edge cases, so they are tuned to reduce whichever mistake matters most.
- **Overconfidence on small evidence:** AI systems may return a score when they classify, match, or predict: a face-recognition system might say 'this face matches: 92 per cent'; a fraud-detection system might say 'this payment looks risky: 76 per cent.' Such a score can be a measure of how strongly the input matched a pattern, not a probability that the answer is right. A 92 per cent score does not mean the system will be correct 92 times out of 100. It does not prove the answer is correct at all, only that the input strongly matched a certain pattern. Just like us, a system can be highly confident and still wrong, and that is exactly when it misleads.

MYTH	REALITY
If an AI gives an answer with 97 per cent confidence, it is probably right.	Unless the system has been carefully calibrated, that 97 per cent is the system's internal score, not a reliable probability of being correct. A confident answer can still be wrong, and a less confident answer can still be right. Treat the number as a clue, not a verdict.

3.3 Deepfakes and Synthetic Media

Until recently, a video of someone saying something was treated as strong evidence that they had said it, and a photo of an event was often treated as good evidence that the event had happened. That is no longer reliable. AI tools can now generate or manipulate video and audio of a real person saying things they never said and doing

© NCERT not to be republished

A Note On Care

Some of the harms described above involve material that is distressing even to come across, particularly image deepfakes that target real people. If you ever encounter material of this kind, whether it is directed at someone you know or sent to you as a joke or a prank, do not forward it. Tell a trusted adult, and report it to the appropriate authority. In India, you can report cybercrime on the National Cyber Crime Reporting Portal at cybercrime.gov.in or call the helpline 1930.

things they never did. These outputs are called deepfakes, a word made by joining 'deep learning' and 'fake.'

How does a deepfake work?

A deepfake system is shown many real images or recordings of a person, their face from many angles, and their voice across many sentences. From these, it learns patterns that can be used to produce new images and recordings in that resembling the person. The technology has legitimate uses—dubbing a film into many languages, restoring old footage, and creating content with the consent of the people involved. But it also has illegitimate ones, and those are the uses that mislead people, damage reputations, and spread false information.

Three harms worth naming clearly: First, fraud. A scammer phones an older relative using an AI-generated voice that sounds like a family member, asking urgently for money. This is happening now, in many countries, including India. Second, image deepfakes, where a person's face is placed onto a body without their consent. The targets are overwhelmingly women and girls, though it can happen to anyone. The harm is real, even though the image is fake; the social and emotional damage to the target is the same. Third, fake content of well-known people: a fabricated video or audio of a public figure saying or doing something they never did, which can spread fast enough to cause damage before anyone can check it.

None of this is a reason to give up on video and audio as evidence. It provides a reason to verify. The next sections build the habits that make verification feel routine rather than an effort.

AI in India

Detection research in India: Research institutions and organisations across the country work on tools that detect deepfakes. The detection problem is genuinely hard, as the tools that generate deepfakes improve, the tools that detect them have to improve to keep up. The target keeps shifting.

Detection is not the only response. Public awareness, platform-level labelling of suspected content, clear reporting pathways, and laws that penalise misuse are all part of the picture. Making or sharing a deepfake to harm or deceive someone can carry legal consequences in India. A society that knows deepfakes exist, knows how to spot them, and knows where to report them is harder to deceive than one that does not. The aim of this unit is to make you part of that society.

Asenla's eye: Asenla is in Grade 11 in Kohima, and she has been her family's designated photographer for years. She has developed an unusual habit: when she looks at a photo or video of a person, she looks first at the small things, the earlobes, the reflection in the eyes, the way a strand of hair falls. Real faces have tiny imperfections, but they all fit together naturally. Lately,

© NCERT not to be republished

looking at AI-generated images of people in all sorts of scenarios, she has noticed something. “I cannot always say what is wrong,” she tells her younger cousin, “sometimes a face just feels off, and I trust the feeling and look harder.” That instinct is trainable, and the next section shows how.

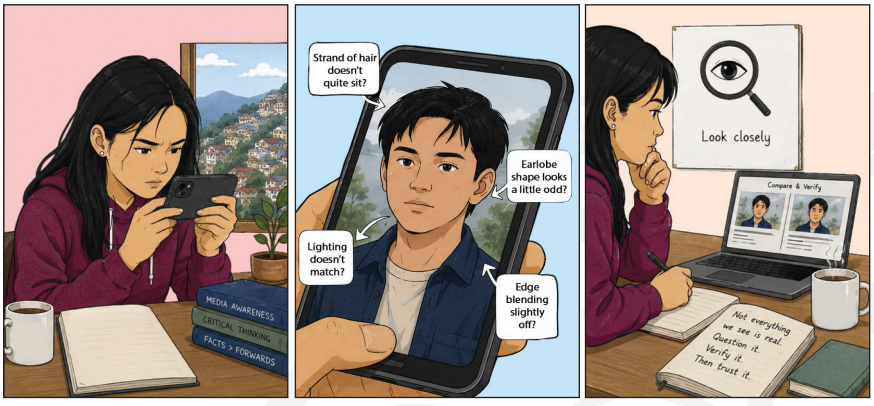


Fig. 3.2: Asenla looks closely for signs of manipulation

A word of caution: Visual clues can help, but modern synthetic media may not show any obvious flaws. The technology improves quickly, and what was easy to spot last year may be invisible this year. Always combine looking closely with checking the source, the context, and the other verification steps mentioned in the next section. Looking carefully is the start of the work, not the whole of it.

Try this Now

Spot the manipulation. Watch a short clip your teacher has chosen, or a public clip such as a public figure at a press conference, a news anchor reading a bulletin, or a documentary segment. Do not use videos of classmates, family members, or private individuals.
Watch it once normally. Then watch it again, looking at just one detail at a time:

- the eyes, and the reflections in them;
- the mouth, and how it shapes each word;
- the hairline and the edges of the face;
- the skin texture;
- the lighting on the face compared with the lighting behind it;
- the audio: does the voice match how the mouth is moving?

As you do this, you are training your eye for the small inconsistencies Asenla looks for. Most clips you watch will be real, and that is fine. The point is not to find a deepfake; it is to build the habit of looking at the parts, and not just trusting the whole.

Write your three most careful observations in your Curious Notebook.

No video in class? Your teacher can show a single printed image instead, a real photo or an AI-generated one, and you can run the same check on the still: eyes, hairline, skin, background, and how the parts fit together.

© NCERT not to be republished



What You have Learnt So Far

A deepfake is AI-generated media of a real person, made to deceive. Synthetic media is the wider category: it has many legitimate uses, and some seriously harmful ones, like deepfakes.

Three harms to remember: fraud; image deepfakes; and fabricated content of well-known people, which can spread before it can be checked. Verification is a key response, along with public awareness, platform-level labelling of suspected content, clear reporting pathways, and laws that penalise misuse.

3.4 Where Your Own Thinking Enters

AI systems do not produce their outputs in a vacuum. They produce them for a reader, you. How carefully you read an output often decides whether a confident-but-wrong answer gets caught or slips through. The problem does not always lie with the AI; sometimes it is the moment when a tired or distracted reader meets an output that already sounds plausible. Three habits of your own thinking are worth naming, because they tell you when you are most likely to let a wrong answer pass as a right one.

- **Confirmation bias:** You tend to trust information that agrees with you more than information that challenges you. If an AI assistant gives an answer that matches what you already believe, you are less likely to check it where as if it contradicts you, you check it twice. The trouble is that your existing beliefs then go unchallenged simply because they are familiar, not because they are true, which is the opposite of what careful thinking requires. For example, if you already think a particular festival is the oldest in your region and the AI agrees, you may never look it up, even though the AI could be as wrong as you are.
- **Anchoring:** The first piece of information you receive shapes how you read everything after it. If an AI tells you a fact and you then read the original source, the AI's version colours your reading: you notice the details that fit what it said and overlook the ones that contradict it. For example, if the AI says a poem is 'about loss', and you then read the poem, you may find loss everywhere and miss the humour in it.
- **Authority overtrust (which AI specialists call automation bias):** Many AI systems are very good at producing output that looks confident, fluent, well-structured, and full of specific detail, the kind of output we associate with a knowledgeable expert. Because it looks as if it came from an expert, you may feel almost rude questioning it. That feeling, that challenging the output is somehow impolite, is one of the strongest forms of overtrust. Notice it when it happens. For example, a neatly formatted answer with exact-looking dates and figures can feel authoritative even when the dates and figures are invented.



Pause and Discuss

Of these three biases, which one do you most often catch yourself doing? Discuss with a classmate. There is no shame here. Noticing the bias is the first step in working around it.

3.5 Verification Habits: The Five Checks

Verification is not a single act. It is a quick, stepwise method you apply before you trust an AI output for anything that matters. The method we suggest has five steps, and once you have practised them, they take only a few minutes together. They work for any AI output, such as a paragraph from a chat assistant, a translated sentence, a photo, a video, a recommendation, and so on. Each is a small check; together, the Five Checks help you to stay aware and careful while using AI.

- **Check 1—Source:** Where did this output come from? Which AI tool, and when? Was it copied, edited, or forwarded before it reached you? Every point along that path is a place where the output could have been changed, so it is worth knowing the path before you trust what it says.
- **Check 2—Claim type:** Some claims are easy to check ('Bhopal's Upper Lake is one of the oldest man-made lakes in India'). Some claims are harder ('This idea was first proposed by such-and-such a person'). Some claims are essentially unverifiable ('The poet was thinking of his mother when he wrote this'). How much scrutiny a claim needs depends on its type. A specific factual claim calls for a specific check; a vague or unverifiable one is best treated as opinion, not fact.
- **Check 3—Compare:** This check has two parts. First, if the AI cited a source, go to that source and check whether it really says what the AI claims. AI assistants sometimes invent references or attribute them wrongly. Second, for important claims, compare against at least two independent sources: a textbook, a reliable website, a teacher, or a friend who knows the subject. 'Independent' is the keyword. Look for a different source, not just another page on the same site. The first part catches a misquoted source

What You have Learnt So Far

Confirmation bias, anchoring, and authority overtrust: three reasons your reading of an AI output is not as objective as it feels.

These biases get worse when you are tired, hurried, or already half-convinced. That is exactly when to be most careful.

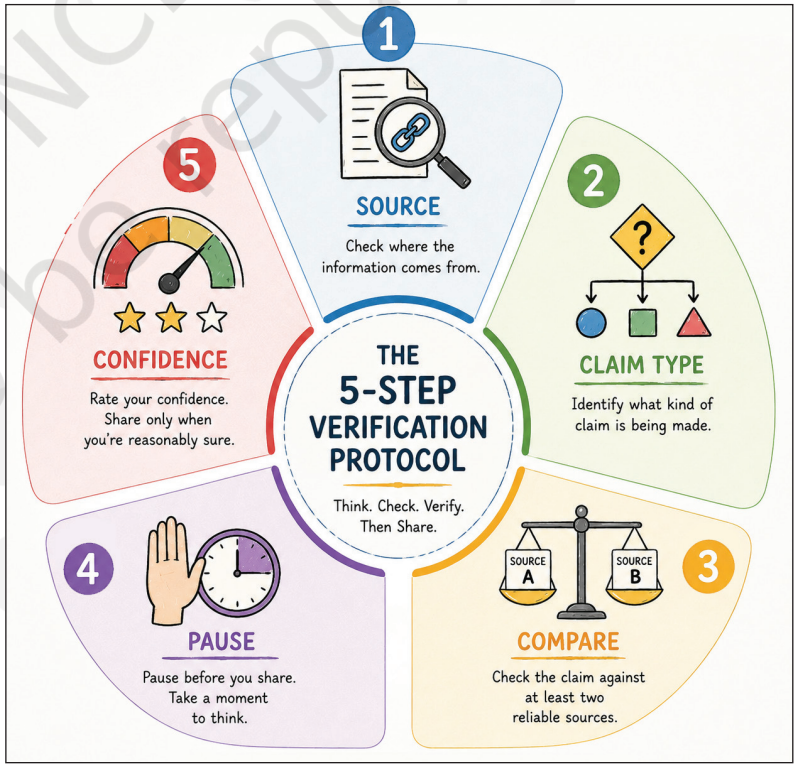


Fig. 3.3: The five checks



What You have Learnt So Far

The Five Checks: Source, Claim type, Compare, Pause, and Confidence help you decide whether you can rely on an AI output.

Build verification as a habit, not as a project, something quick you do every time, not a special effort you make once.

whereas the second part catches a single wrong source or a repeated error.

- **Check 4—Pause:** Before you forward, share, or submit an AI output, pause and re-read what you are about to send. Ask yourself: am I confident enough in this to put my name to it? If not, do not pass it on yet. Forwarding an unchecked claim makes you a link in the chain that spreads it.
- **Check 5—Confidence:** If you do pass an output on, say how sure you are. "I think this is right, but I have not double-checked" is a perfectly good thing to say, and it lets the next person judge for themselves. If you sound more confident than you really are, uncertain information starts to look certain. As it passes from person to person, each retelling can add a little confidence, until the original doubt disappears and a wrong answer ends up looking right.

Where Diya learned this: Diya, whom you will meet properly in Unit 4, picked up these habits by watching her aunt, who worked at a fact-checking website. Her aunt's job was to trace a viral claim back to its source—who said it, when, where, and in what context. Diya noticed that her aunt did not begin by deciding whether the claim was true or false. Her aunt began by treating the claim as something to be located. The Five Checks are a compressed version of that habit, designed to take a few minutes rather than a whole afternoon.

All five together should not take long once you have practised them. Verification is not a research project; it is a habit, like looking both ways before crossing a road. The first few times feel slow and effortful. Then they stop feeling like anything at all.

What is being checked: If you look closely, the first three checks, 'Source', 'Claim type', and 'Compare', are checks on the AI's output. The last two, Pause and Confidence, are checks on your own thinking which is intentional. The output can be wrong, but how you treat the output is what decides what happens next. The Five Checks combine the two because, in practice, they happen together, and one five-step habit is easier to carry than two separate lists.



What You have Learnt So Far

"Good" depends on the situation. Accuracy, fairness, robustness, calibration, and explainability are all different standards, and which ones matter most depends on what is at stake.

The right balance is not automatic; people decide it. Knowing which questions to ask is a real part of that work.

3.6 What Good Evaluation Means

'Good,' when you are evaluating an AI output, is not a single thing. Different situations call for different standards.

Accuracy is not always enough. A system that gets 99 questions right out of 100, sounds excellent, until you learn that the one it got wrong was a medical diagnosis that mattered. A system that is 80% accurate but is never confidently wrong on

the cases that matter most—may be more useful than a 99% system that occasionally produces a confident catastrophic error. The right standard depends on what is at stake.

Other standards worth knowing:

- **Fairness**: does the system work equally well across different groups of users?
- **Robustness**: does it stay accurate when the input changes a little, or does a small disturbance break it?
- **Calibration**: when it says it is 90% sure, is it right about 90% of the time?
- **Explainability**: can a person understand why it produced this answer, or is it a black box whose decision process you cannot see?

In practice, you balance several of these at once. A system used to screen medical images might put robustness and calibration ahead of speed. A system for sorting customer queries might put speed and accuracy ahead of explainability. A system used to support a decision in a courtroom probably needs to put explainability above almost everything else. These choices are not automatic. People make them.



Computational Thinking Lens

The Five Checks you just learned bring together two habits of computational thinking—algorithmic thinking and evaluation.

Algorithmic thinking is designing a clear, step-by-step method for a task. The Five Checks are exactly that—a method applied to a tricky problem, taking an uncertain claim, breaking the decision into clear steps, checking each step against evidence, and only then deciding whether to trust the output or not.

Evaluation is the habit of judging whether something is good enough, and by which standard. That is the final move in the verification protocol—deciding whether this particular output is good enough to use, for that particular purpose.

Stretch: Confusion Matrix, Precision, and Recall

Optional. Mainly for Grade 12 readers, but anyone is welcome to attempt.

When a system makes yes-or-no predictions, every prediction falls into one of the four groups. Laid out together, they form what is called a **confusion matrix**.

	Should have been flagged	Should not have been flagged
System flagged it	True positive	False positive
System did not flag it	False negative	True negative

Two numbers summarise this matrix.

Precision answers: of all the things the system flagged as a problem, what fraction really were problems? It tells you how trustworthy the system's alerts are.

© NCERT not to be republished

Recall answers: of all the real problems, what fraction did the system catch? It tells you how many of the real problem the system was able to find.

Improving one can sometimes make the other worse. A medical screening test set very sensitively catches for almost every real case (high recall) but also flags many healthy people (lower precision). A spam filter set very cautiously, blocking only what it is sure about, rarely blocks a legitimate e-mail (high precision) but lets some spam through (lower recall). Engineers choose the balance by deciding which mistake would be worse.

3.7 Your Careful-use Checklist

The Five Checks of Section 3.5 are the core. Around them, build your own personal checklist of habits that work for you. Below is a starter version. In your portfolio activity for this unit, you will adapt it into your own.

Habit	When to use it
Read it twice	Before you trust any AI output. Once for the gist, once for the specifics.
Find one independent source	For any specific factual claim that matters.
Spot-check two or three details	For long outputs like paragraphs and summaries. A wrong detail buried in a long output is where hallucinations hide.
Check the dates	For anything time-sensitive, prices, leaders, scores, or any recent events.
Ask what kind of claim it is	A specific fact, an opinion, or a matter of taste. Different claims need different kinds of checks.
Take a pause before forwarding	Always. Forwarding is when bad output spreads.
Mark your confidence	When passing anything on. Use “I think…”, “I checked…”, “I am not sure but…”.
For high-stakes topics, take longer	Medical, legal, or financial matters, anything that affects another person's welfare. Better to be sure.

The point of these habits is not to refuse to use AI tools. It is to use them well. AI tools genuinely help with summarising long texts, brainstorming options, and explaining a difficult idea in more than one way. The point of verification is to keep that help while reducing the risks. A careful, critical user gets far more value from a tool than an uncritical one does.

© NCERT not to be republished

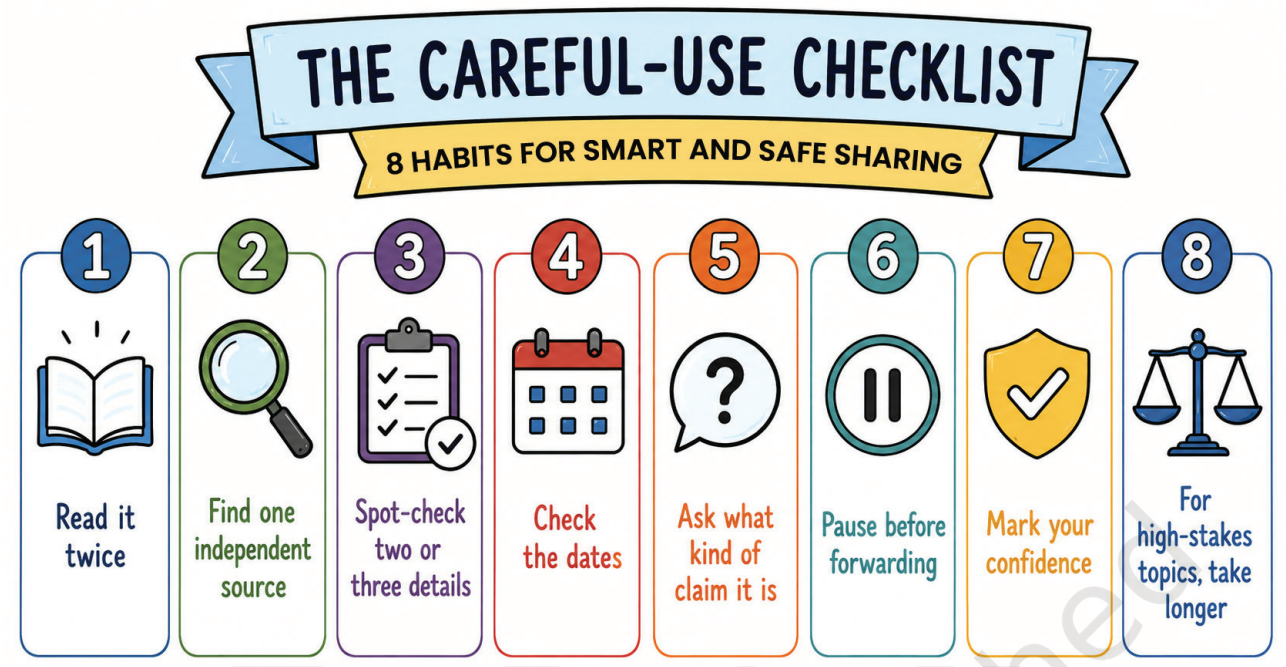


Fig. 3.4: Being careful

Inclusivity Lens

Verification has its own accessibility question.

Fact-checking tools, official sources, and quality publications are mostly in English, and many sit behind paid subscriptions. A reader who works in Manipuri or Santali, or who relies on a screen reader, has fewer easy-to-reach sources to compare against. The habit of verification is the same for everyone; the ease of applying it is not. As you build your own checklist, notice which of your habits depend on resources that not everyone has, and try to include at least one that almost anyone can reach, a knowledgeable person, a library, a public broadcaster.

What You have Learnt So Far

A personal checklist makes the verification habit your own. Adapt the starter checklist to suit your needs; do not just copy it.

Verification keeps the value of AI tools while reducing the risks. The more carefully you use AI, the more useful it becomes.

Reflect

Think of the last time you passed something on, an answer, a forward, or a claim, without checking it. Which of the Five Checks would have caught a problem fastest? Write three lines on what you trusted, why it felt trustworthy, and which one check you will make a habit first.

© NCERT not to be republished

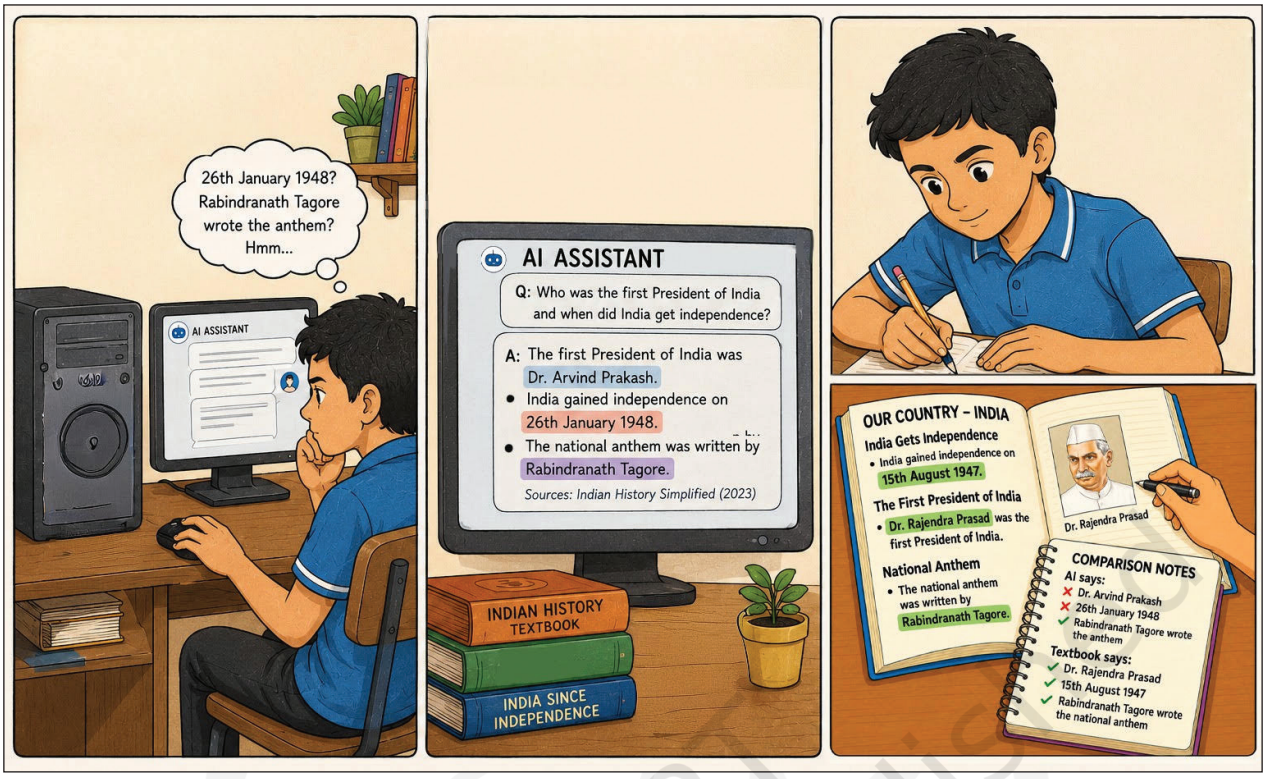


Fig. 3.5: Hallucination Hunt in action

For your Curious Notebook

This is a two-part task. The first part is for everyone, while the second is an optional challenge, if you have access to an AI tool.

- **Part 1 (everyone)—your Personal Verification Checklist:** Adapt the starter checklist from Section 3.7 into one of your own. Add habits that work for you—remove ones that do not. Keep it to one page. This becomes your reference for the rest of this module, and for your life with AI tools.
- **Part 2 (optional, to go further)—a Hallucination Hunt:** Pick an area you know well, a sport you follow closely, a festival from your region, a local historical figure, your school's history. Ask an AI assistant five detailed questions in that area and note down each answer. Then, using your own knowledge and the Five Checks from Section 3.5, filter and sort the answers into those that are right, those that are wrong, and those that contain hallucinated details. For at least two of the factual claims, show the independent source you used to check it—a book, inputs from a teacher, a reliable website, or the method you used; do not rely only on memory. For each wrong answer, write one line naming the kind of error—a hallucination, a distribution-shift error, an edge-case mistake, or an overconfident but vague answer.

If you cannot use an AI tool yourself, your teacher can run a single AI assistant on the projector and walk the class through the five questions together. You still write your own observations and your own checklist.

Through this activity, you will practise noticing when an AI system produces a hallucination, and naming the kind of error you have found.



Assignment

Quick checks

1. In your own words, what is a hallucination in an AI assistant? Give one situation where a hallucination would matter, and one where it would not.
2. The unit says "confidence is not the same as correctness." Give an example of a confident-but-wrong moment from your own life, and a hesitant-but-right one.
3. Of the three reader biases in Section 3.4—confirmation bias, anchoring, and authority overtrust—which do you think you are most likely to fall into? Write two lines on why.

Apply it

4. A friend forwards you a video of a politician saying something inflammatory. Walk through the Five Checks on that video. What would each check look like in this case?
5. You ask an AI assistant to summarise a 2,000-word article you do not have time to read. The summary is fluent and confident. Name three things you would check in it before treating it as accurate, and also explain why you chose those three.

Stretch (optional)

6. A school is considering letting students use an AI assistant for assignments, on the condition that the output must be verified before it is submitted. Write a half-page note. Would this work? What would have to be true for it to work well? What would have to be true for it to fail badly? What is your own view, and why?

Coming up in Unit 4: You have started building the verification habit. The next unit turns to the data you put into AI systems—the rights you have over it, and what to do when something goes wrong.

© NCERT not to be republished

UNIT

4

Your Data, Your Rights, Your Responsibility

UN514UN04



The first time you open any app, you have probably been shown a long block of text headed 'Terms and Conditions' or 'Privacy Policy,' and a message asking you to agree, 'By continuing, you agree.' Most people scroll straight to the bottom and tap 'Accept,' because the text is long, the language looks complicated, and the only alternative to accepting seems to be not using the app at all. The same kind of policy appears the first time you visit a website, often as a cookie banner, the small pop-up asking you to accept 'cookies,' the little files a site stores to remember you. In practice, almost everyone scrolls and clicks 'Accept' without reading. That is exactly what the company expects.

This unit argues that you should read, at least sometimes, at least for the apps that matter most. Once you know what to look for, reading takes a few minutes. It breaks those long policies into small, understandable pieces, not because you will read every policy for the rest of your life (you may not), but because

once you see what these policies quietly contain, you can ask better questions: which apps are worth installing, what they should and should not be allowed to do, and what to do when something does not feel right.

The unit focuses on the habit of decomposition: breaking one big decision, 'Should I install this app?', into smaller ones. What data does this app want? Who will it share that data with? How long will it keep my data? How do I take it back?

This is also the unit where AI tools begin asking something of you, not just offering things. Using an AI tool is a two-way exchange: it gives you an output, and you may give it data in return. Knowing your side of that bargain matters. You need to know what you are agreeing to, and what to do if something goes wrong. This unit addresses both.



Words You'll Meet in this Unit

Privacy law uses a few specific terms. The unit introduces them in plain language, and the alphabetical glossary at the back has them all.

- Personal data and sensitive personal data are, respectively, data that can identify you, and data that can cause greater harm if misused, such as health, financial, or biometric information.
- Consent is your permission for a company to handle your data. Real consent is specific, informed, voluntary, and revocable, meaning you can take it back.
- Data Principal is you, in the language of India's privacy law, the person whose data it is.
- Data Fiduciary is the company or organisation that decides why and how your data is handled, and that has legal duties towards you.
- The Digital Personal Data Protection Act, 2023 (DPDP Act) is India's data protection law, and the DPDP rules are the operational rules that put it into practice.
- The grievance officer and the Data Protection Board are, respectively, the company's point of contact for complaints, and the formal Indian body for data and privacy disputes.

4.1 Not All Data is the Same

'Data' is one of those words that has expanded to mean almost anything. It helps to break it into kinds, because each kind of data has its own importance.

- **Identifying data:** Your name, your face, your fingerprints, your Aadhaar number, your phone number, and your address can be used to identify you specifically among many people. They are the most serious kind of data to share, because they connect everything else back to you.
- **Behavioural data:** what you click on, how long you spend on a page, which songs you skip, which messages you reply to fastest. This is the quiet kind of data, the kind you may not even realise is being collected. It often reveals more about you than the things you fill in on a form, because it is harder to fake.

© NCERT not to be republished

What You have Learnt So Far

Your device collects four kinds of data: identifying, behavioural, sensitive, inferred.
Each has different importance.

Inferred data is often the most valuable to a company and the most invisible to you.

- **Sensitive personal data:** Some kinds of personal data can cause greater harm if misused than others. These include health records, financial details, biometric data, and other categories that laws often protect with special care. Older Indian rules used a separate category called 'sensitive personal data', and many other countries' laws still do, giving such data extra legal protection. India's current law, the DPDP Act, takes a somewhat different approach, with particular attention to children's data. The legal terminology in this area keeps changing, but one practical rule holds across all of them: be extra careful whenever you are asked to share data of this kind.
- **Inferred data:** Conclusions a system has drawn about you from the other kinds of data. "This user probably likes Marathi devotional music." "This user is likely between 22 and 28 years old." "This user is likely a college student." You did not give the system this information. It worked it out. Inferred data is invisible to most users, but it is often what decides what you see on your screen, which advertisements you are shown, and sometimes even the prices you are offered.

Try this Now

Choose one app or device you use, on a smartphone at home or on a school computer. Write down what kinds of data it might store about you, using the four kinds of data you have just read about: identifying data, behavioural data, sensitive personal data, and inferred data. Be specific.
For example, for a video app:

- Identifying data: my name and email address.
- Behavioural data: which reels or videos I watched fully, and which I skipped.
- Sensitive personal data: ...
- Inferred data: ...

Fill in all four categories. If your answer for any category is 'I don't know,' write that down. Not knowing is also a finding.

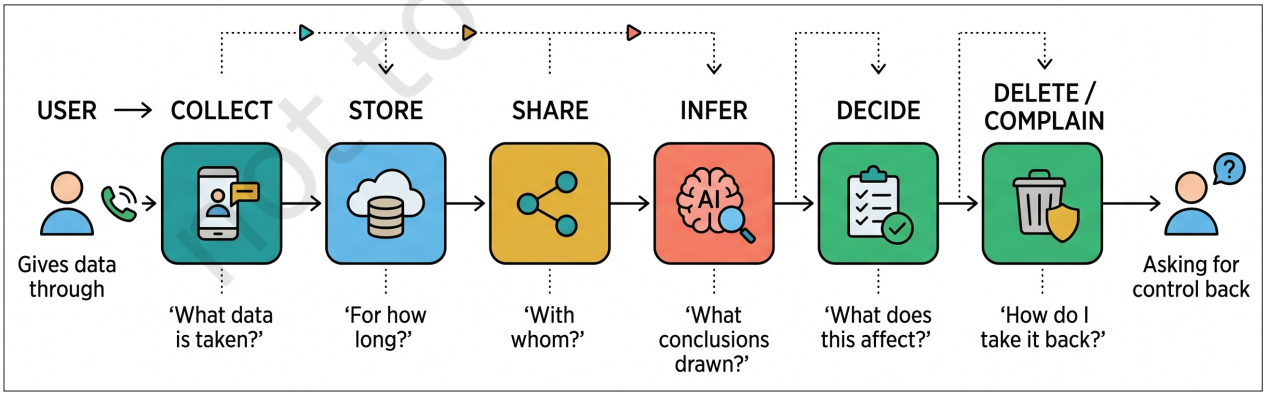


Fig. 4.1: Where your data goes, and the question to ask at each stage.

4.2 Consent, in Plain Language

Consent sounds simple. You agree, or you do not. In practice, the consent that meaningfully protects you has four components. These four points are not the whole law. Different countries may have different rules. Still, together they show what makes consent meaningful.

- **Specific:** 'I agree to the company using my data' is not specific. 'I agree to the company using my email address to send me one weekly summary of its new features' is specific. The vaguer the consent, the broader the permission, and the more freedom the company has to use your data, and the harder it becomes later to say, 'I did not agree to that.'
- **Informed:** You cannot give consent to something you do not understand. If a privacy policy is 40 pages long and full of legal language, you cannot give meaningful consent to it. Real informed consent depends on the company explaining what they are asking for, in language a normal person can follow. Some companies do this well, a one-page plain-English summary at the top, with the legal text underneath. Others do not.
- **Voluntary:** If your only choice is to agree or not use the service, and the service is something you genuinely need (a banking app, a school portal, a government service), the consent you give is shaped by that need. The law recognises this and tries to protect against the most extreme version, but the situation is real. Voluntary is not just yes or no; it is a spectrum.
- **Revocable:** Consent given once should be withdrawable. If the company does not let you withdraw your consent, or makes withdrawal impossibly difficult, the consent was actually never yours. Revocability is what makes consent meaningful, not just a box you ticked once.
- **Indian law:** The Digital Personal Data Protection Act, 2023, commonly the DPDP Act, sets out rules for how organisations must handle personal data of people in India. Under the Act, the person whose personal data is being collected or used is called the Data Principal. That is you, the person to whom the data belongs. The company or organisation that decides why and how your data is processed is called the Data Fiduciary. Under the Act, the Fiduciary, alone or with others, determines the purpose and means of processing your data, and so carries legal

© NCERT not to be republished

duties for how it is handled. The data relates to you; the Fiduciary is accountable for what happens to it.

- **Your rights, in plain English:** The DPDP Act provides the right to be told what data about you is being held, the right to have it corrected if it is wrong, the right to have it erased in certain cases, the right to a grievance-redressal process, and the right to nominate someone to exercise your rights if you cannot. To raise a problem, you usually approach the company's grievance officer first; if it is not resolved there, you can escalate to the Data Protection Board. The full law has more provisions; these are the ones that affect you most often.

DPDP Rules 2025—A Quick Update

In November 2025, the central government notified the Digital Personal Data Protection Rules, 2025, the operational framework that puts the 2023 Act into practice. As per the notified Rules and the current implementation schedule, the Rules are being brought into force in stages; the exact timeline can be revised, so check the Ministry of Electronics and Information Technology (MeitY) website for the current position.

Stage 1, immediate (from November 2025): the machinery to set up the Data Protection Board of India comes into force. The Board is designed as a digital-first body, intended to work mainly online.

Stage 2, from November 2026: the rules on Consent Managers take effect. A Consent Manager is a new kind of organisation, registered with the Board, that lets you give, manage, review, and withdraw your consent across many services from one place.

Stage 3, from May 2027: the main operational obligations come fully into force, notices, security safeguards, breach reporting, data-deletion timelines, the rights of Data Principals, and most protections for children's data.

A few specifics worth knowing: If a company suffers a personal-data breach, it must tell each affected person and report the breach to the Board without delay, with detailed information within 72 hours. Every company must run a grievance-redressal system and respond within a published time limit, which the Rules cap at 90 days. And large platforms must erase your data a set period after you stop using it, for the biggest e-commerce, gaming, and social-media services, that period is three years.

Why this matters for you: Many readers of this module are under 18. The framework requires a company to obtain verifiable consent from a parent before processing a child's personal data. This has real implications for the apps and services you use, and it is worth raising with a parent or teacher if a service seems to be processing your data without the consent the law requires.

Why we are telling you now: Privacy law in India is changing rapidly. By the time you finish school, the rules may look quite different. The principles behind them, your data, your rights, the company's duties, stay the same; the specific timelines and procedures will keep changing. Check the website of the Ministry of Electronics and Information Technology for the current version when you need it.

© NCERT not to be republished

MYTH

I clicked 'Accept', so I have given the company permission to do whatever they want with my data.

REALITY

Clicking on 'Accept' gives permission only for what the company has actually told you it will do, in language you could reasonably understand. If they do something you were not told about, that is not covered by your consent, and may be against the law.

4.3 What not to upload into an AI tool

AI tools such as chat assistants, image generators (tools that create pictures from text descriptions), and document summarisers work by being given input and producing output. The input you give them often goes back to the company that made the tool. It may be used to improve the tool. It may be stored. It may be reviewed by people. The exact rules depend on the tool and on whether you are using a free version or a paid one. The shape of the rules across tools is similar enough that a few habits travel.

Things to be careful about uploading: Each of these has a reason, not always a single reason, but a reason to be treated cautiously.

Do not upload	Why?
Aadhaar number, PAN, full bank details	These are personal data tied to formal identity. If exposed, the harm cannot easily be undone.
Passwords or one-time passwords (OTPs)	These are credentials. Once you paste them into another system, you lose control over who or what may access them.
Photos of your ID documents	These are images of an Aadhaar card, driver's licence, or passport carrying the personal data plus your face. The combination is more valuable than either alone.
Other people's personal information	Their consent, not yours, is what governs what gets uploaded about them. You do not have that consent.
Confidential work or school documents	Many tools may use uploaded content to improve themselves. A paper a student has not yet submitted, or a document marked confidential, should not be shared with any tool unless you understand how it will store, use, or share it.
Medical details, especially about minors	These are sensitive personal data. The harm from misuse is high. The benefit of having an AI tool summarise it is rarely worth the risk.
Anything you would not want shared publicly	If you would not want it shared with people you do not know, do not paste it into a tool whose handling you cannot fully see.

What You have Learnt So Far

Real consent is specific, informed, voluntary, and revocable. Anything weaker is consent in name only.

Under current Indian law, you are the Data Principal. The company is the Data Fiduciary, the organisation that decides why and how your data is processed, and therefore has legal duties about handling it. The data relates to you and is meant to be handled carefully and lawfully.

© NCERT not to be republished

Two narrower exceptions: First, if the tool runs entirely on your own device, with nothing uploaded to a company's servers, the judgement is different, you are not handing your data to anyone else. Second, a school or workplace contract, or a setting you have personally checked, may reduce the training and retention risks for some inputs. But 'reduced risk' is not 'no risk'. A password or one-time password belongs only in the app or website that issued it. Do not give one to anything else, including any AI tool. ID documents and other people's personal data belong only where an official service you approached genuinely requires them, and never in a general-purpose AI tool. Treat any other request for these as a warning sign. Both exceptions still require you to actually check the policy, not just assume.

What to Look for in an AI Tool's Data Settings

Whether a tool is free or paid can be a useful clue, but tools differ on more specific points too. Before using a tool for anything that matters, check its settings page. The questions worth asking are:

- **Training use:** Does the company say your inputs may be used to train future versions of the model? Some tools train on your inputs by default; some let you turn it off; some have a paid tier where training is off by default. Check.
- **Data retention:** How long is your data kept, a few days, a few months, indefinitely? Some tools delete it after a set period; others keep it until you delete it yourself.
- **Human review:** Are people at the company allowed to read your conversations, for quality checks, safety review, or other reasons? Many tools allow this in some form. The settings page should say so, and some let you opt out.
- **Memory of past conversations:** Some tools remember things from earlier conversations and use them in new ones. Memory is convenient, but it is also a place where something you wrote once can resurface later. The settings page should let you see and clear what is remembered.
- **Temporary or incognito chat:** Many tools offer a session that does not appear in your history or memory and may not be used for training. Even so, tools may keep such conversations briefly for safety, abuse review, or legal reasons, so check the policy. It is often useful for one-off questions you do not want kept long-term.
- **Cookies and website tracking:** Many websites store small files on your device, called cookies. Some are useful (they keep you logged in); others track you across sites for advertising. Browsers have settings to manage or clear them, the website-level equivalent of clearing your chat history.
- **Shared devices:** If you are using a phone, computer, or tablet that others also use, a family member's phone, a school or library computer, an internet-café terminal, sign out of any account you opened, clear the browser history if you can, and never tick "remember me." A login that is convenient for you is an automatic login for whoever uses the device next.
- **School or workplace controls:** If your school or workplace provides the tool, an administrator may have set the rules already. Ask which version you are on, and what the privacy contract says about your inputs.

© NCERT not to be republished

4.4 Fairness, inclusion, accessibility

AI systems can be a barrier for some people and a pathway for others. Sometimes it is both at once, helping one group while quietly excluding another. The unit's third major idea is to look for both effects, not just the one you happen to notice.

Where AI removes barriers: Voice assistants make a phone usable for someone with low vision, or for people who find reading or typing difficult. Bhashini-enabled multilingual services bring government and information services to a user who works in Marathi or Tamil rather than English.

Augmentative and alternative communication (AAC) devices, often AI-powered, give a child who does not speak with their voice a way to communicate. Sign-language recognition research at Indian institutions is working towards real-time interpretation. For the people who depend on them, these are not conveniences; they are what makes access possible at all.

Where AI does not work for everyone: Some face-recognition systems have been shown to perform unevenly across different groups of users. How well they work depends on the algorithm, the application, the training data, and the capture conditions. A system that performs poorly in certain conditions, or for certain groups, will reject more of those users at a security gate. A voice assistant trained mostly on one accent struggles with another, leaving some speakers shouting at their phones while others are understood the first time. A medical screening tool tested mostly on adult patients may quietly underperform on adolescents. None of these are mistakes in the moral sense; they are honest consequences of training data, design choices, and deployment conditions, the same lessons we met in Unit 2.

What this asks of you: When you meet a new AI tool, two questions are worth asking before you trust it: who does this tool work well for, and who does it work less well for? 'It works for me' is the most common way to miss the second question. Some of the systems your generation will build will decide, through their choices about training data and evaluation, who is helped and who is left out. Notice the pattern early.



Inclusivity Lens

These accessibility features are not extra enhancements to add at the end; they need to be part of the design from the start. For the people affected, they shape the whole experience. AI's biggest contribution in someone's life might be a screen reader for a student with low vision, a translation app for a patient who does not speak English, or an AAC device for a child who does not speak with their own voice. From the outside, it does not look like AI at all; it looks like a service that finally works. As you build your own list of how AI affects the people around you, count these uses too.



What You have Learnt So Far

AI can reduce barriers but, AI can also create barriers. Often the same system does both, for different users.

Two questions to ask before you trust a new tool: who does it work well for, and who does it work less well for?

4.5 AI in Your Learning: What It Does For You, What It Does to You

This section is for you specifically, as a student. AI tools are now within easy reach for studying: explaining a difficult concept, suggesting different ways to phrase a paragraph, summarising a long article, generating practice questions. Used well, they can speed up your learning. Used recklessly, they can do real damage. The difference is worth understanding now, not when an exam goes wrong in two years.

What AI tools do well in your learning: They can explain something you do not understand in three different ways until one of them clicks. They can suggest five openings for an essay you cannot get started on. They can compress a 4,000-word article into a paragraph so you can decide whether to read the whole thing or the summary. They can find the gap in your reasoning when you describe an argument out loud. They can generate practice questions when your textbook has run out. These are real services that genuinely save time and help learning along.

Table 4.1: How AI can support learning, and how it can weaken it

AI can support learning by...	AI can weaken learning when...
explaining a concept in different ways	you accept the explanation without checking
generating practice questions	you practise only easy or unchecked questions
giving feedback on a draft	you stop writing first drafts yourself
helping with language	your own voice and opinion disappear
summarising long material	you stop reading long arguments carefully

The hidden cost: If you always ask for the answer, you never build the habit of finding it yourself. If you always have a draft generated, you never learn the harder skill of producing a first draft from a blank page, and that is the harder skill, because a first draft is where your own thinking takes shape. If a summary is always handed to you, you never develop the patience to read a long argument all the way through. The cost is invisible at the moment. It shows up later: in an exam where the AI is not there, in an interview where the answer must come from you, in a conversation where someone asks what you actually think and you realise you have not done the thinking.

© NCERT not to be republished

What separates the student who uses AI well: Asenla, in Grade 11, uses an AI assistant after she has tried a problem on her own, not before. She uses it to check her reasoning, not to outsource it. She does not copy its answer directly; she first puts it into her own words. When she submits work, she is honest about what came from where. Another student, on the same assignment, treats the tool as an answer key. He pastes the question, copies the output, does not check it, and skips the part where his own thinking would have happened. He gets through the assignment but not through the learning. The first student uses the tool as a helper; the second uses it as a replacement. Same tool, two different ways of using it, and over a year the difference shows up in what each student can actually do.

On academic integrity, briefly: Different schools, exams, and countries have different rules about what AI assistance is acceptable. Some allow it openly, as long as you say where you got the help, sometimes called citing AI use, telling the reader which part came from the tool. Some forbid it for certain kinds of work. Find out what the rule is in your context. The procedural rules matter, but the deeper rule, the one that protects your own development, is the same everywhere: use AI tools so that the thinking is still yours. Unit 5 develops this into a working method, how to ask, iterate, and reflect with AI tools so that your own thinking grows with the work rather than around it.

4.6 If something goes wrong

Sometimes a system makes a mistake about you. The face recognition system at a metro gate or campus entrance fails to recognise you, and you are sent for manual verification while the people behind you walk through. An automated essay-grading tool gives your draft a low mark for reasons you cannot make sense of. A school portal records that you completed an assignment you did not, or fails to record one you did. A deepfake using your face appears on a website without your consent. None of these is common, but all of them are real, and the right time to learn what to do is before any of them happens.

Asenla's cousin: Asenla, whom you met in Unit 3 with her photographer's instinct for spotting wrong-looking faces, helped a younger cousin handle exactly such a problem last year. The cousin found a video on a friend's phone with her own face on it. She was upset and embarrassed, and her first instinct was to delete the video and pretend nothing had happened. Asenla, who had been on a panel at school about this very thing, said no: do not delete it, screenshot it, note where it came from, and



PAUSE AND DISCUSS

Suppose your school permits AI tools for some kinds of schoolwork. Discuss with a classmate: which tasks would still be your own thinking, and which would not be? Where would the line be drawn for a Maths problem? An English essay? A Science assignment where the unit is the learning? What signals would tell you which side of the line you were on?

What You have Learnt So Far

Used as a study partner, an AI tool can explain, brainstorm, summarise, and check. That adds real value.

Used as an answer key, it only helps you copy and submit, and you skip the learning. The cost may not show up at once. It appears later, when your own ability has not grown.

The rule that protects your development is simple: use the tool so that the thinking stays yours.

take it to a parent. They went together to the cousin's mother, who filed a complaint at the cybercrime portal that night. The video was taken down within a week. They had slowed down and worked through it one step at a time. The hard part was not the technology. It was knowing that facing the problem was better than hiding from it.

Diya in Surat: Diya is in Grade 12 in Surat, captain of her school debate team. Her aunt works at a fact-checking website, and from years of watching her aunt trace a claim back to its source, Diya has picked up a habit.

When she hits a problem with an app, an online service, or an AI tool, her first question is: where did the wrong piece come from, and who is in a position to fix it? She learned to ask not only "who is responsible?" but also "who can actually see or use the data?" The right place to take a complaint is usually not the most visible person on social media, and not the angriest comment thread. It is the formal, often quieter, channel built for it.

A general escalation path: The path varies by the kind of problem, but the shape is similar each time.

- **First, contact the company directly**. Most companies name a grievance officer in their privacy policy or app settings. Give them a clear, written description of the

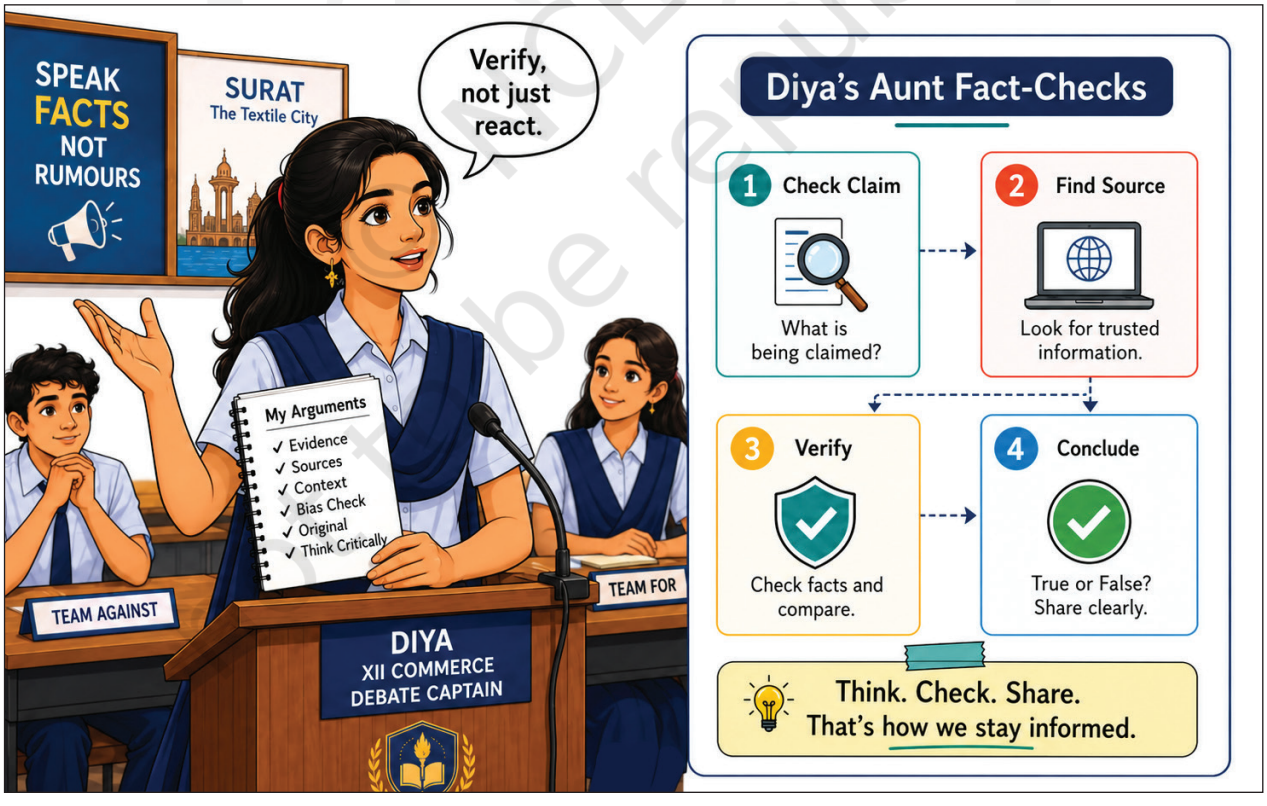


Fig. 4.2: Diya reminding us to verify, not just react.

problem. For routine data and privacy issues, allow a reasonable window to respond; companies must run a grievance-redressal system and respond within a published time limit, which the Rules cap at 90 days.

- **Second, escalate if needed**. If the company does not respond or refuses to fix the problem, take it to the Data Protection Board (for data and privacy issues) or the relevant sectoral regulator (for finance, telecom, healthcare, and so on).



Fig. 4.3: From a long privacy policy to five key questions

- **Third, for criminal harm, act at once**. For fraud, deepfakes, harassment, intimate-image abuse, or ongoing threats, do not wait. You can report it on the National Cyber Crime Reporting Portal at cybercrime.gov.in or call the helpline 1930. A school counsellor or a trusted adult is a useful person to talk the steps through with before you take them.

The key is not to panic, and not to stay silent. These are two opposite errors. Panicking leads to rushed, wrong decisions, or to paying a scammer to make a problem disappear. Staying silent lets a problem grow: a deepfake spreads, a wrong record persists, a fraud claim mounts interest. The path between them is to slow down and write things down before doing anything else. Keep a record of what happened, the time, the messages you saw and received, the people involved, and then follow the process one step at a time.

Where a Complaint can Go

- **First, talk to an adult you trust**. Reach out to a parent, guardian, or trusted teacher. They can help you decide what kind of issue this is and which channel fits. For serious matters, such as deepfakes, harassment, extortion, or identity misuse, do not try to handle it alone.
- **Second, try the app or organisation's own grievance mechanism**. Most companies name a grievance officer in their privacy policy or app settings. Give them a clear, written description of the problem and a reasonable time to respond. Many issues are resolved here.
- **If it is not resolved, escalate to the right authority**. For data and privacy issues, the Data Protection Board of India is the formal body under the DPDP Act, subject to the procedures in force. For sector-specific issues, approach the relevant regulator: the Reserve Bank of India for banking, the Insurance Regulatory and Development Authority for insurance, the Telecom Regulatory Authority of India for telecom, and so on.
- **For criminal matters, report at once**. For fraud, deepfakes, harassment, or identity misuse, use the National Cyber Crime Reporting Portal at cybercrime.gov.in and the cyber helpline 1930. Both are run by the Indian Cyber Crime Coordination Centre (I4C) under the Ministry of Home Affairs. Report these immediately, without waiting to see whether the company's grievance process works.

Reflect

Pick one app you use most days. Without checking, write down your best guesses: what data of yours does it have, what does it use that data for, who does it share it with, how long does it keep it, and how can you take it back?

Now find the real answers in the app's privacy policy. How many of your guesses were right, and what surprised you most?

© NCERT not to be republished

Computational Thinking Lens

Decomposition is the habit you have practised again and again in this unit. You broke one big decision ('should I install this app?') into smaller ones about data, consent, sharing, retention, and revocation. You broke one big problem ('something has gone wrong') into smaller steps about who can actually fix it. Decomposition is the engineer's habit, the lawyer's habit, and increasingly the citizen's habit.

For your Curious Notebook

Privacy Policy Dissection: Take one app you use regularly, such as a learning app, a game, a social-media app, a messaging app, a bus or train app, or an OTT app. Find its privacy policy. Then, in your Curious Notebook, write a one-page plain-English answer to each of these five questions. One line per question is enough; the work is in finding the answer.

- What kinds of data does this app collect from you? (identifying, behavioural, sensitive personal, inferred.)
- Why does the app say it collects each kind?
- Who does the app share your data with?
- How long does the app keep your data after you stop using it?
- How can you delete your account and have your data removed?

If the policy does not answer one of these questions clearly, write 'Not stated' as your answer. 'Not stated' is itself a finding worth noting.



Assignment

Quick checks

1. The unit lists four kinds of data: identifying, behavioural, sensitive personal, and inferred. Give one example of each from your own life.
2. In your own words, what makes a consent real, specific, informed, voluntary, and revocable, as opposed to consent in name only?
3. Of the items in Section 4.3's 'do not upload' list, which two would you most have to remind yourself about? Two lines on why those.

© NCERT not to be republished

Apply it

4. A friend wants to summarise their Grade 12 Physics notes by uploading them to an AI assistant. The notes include some classmates' names and a few personal details from a project list. What would you advise them to do, and why? Be specific.
5. A face-recognition system at a metro station's entry gate rejects you twice in a row while waving through the person behind you. You suspect the system is not working reliably for you. Name one thing the station administrator could do, one thing the system's developer could do, and one thing you could do today to improve the situation.

Stretch (optional)

6. Your school is considering an AI-powered attendance system that recognises students' faces, with the aim of faster check-in and accurate records. Write a half-page note for the principal that decomposes the decision into its parts: what data is collected, what consent is needed, who keeps it, for how long, what redress exists, and who might be left out. End with your own view, and your reasons.

© NCERT not to be republished

UNIT

5

Working with Modern AI Tools

Prompts and the Helper Stance

UN514UN05



A note on this unit. This is the unit where everything so far comes together: you take the habits you have built, recognising patterns, checking before trusting, protecting your data, and put them to work by actually using an AI tool to do something real. That makes it the most rewarding unit, and also the most demanding, because working with these tools needs access that not every classroom has yet. For that reason, it is optional, and your teacher will decide whether and how to run it. Where tools are not available, the unit gives unplugged (offline) versions of its activities, so the core ideas, how to ask, refine, and reflect, can be learnt with paper and discussion alone. If your class covers Units 1 to 4, you have completed the core of this module. Unit 5 is where you get to use it.

You have spent four units understanding AI tools: noticing them around you, understanding how they learn, catching their mistakes, and asking what they want from you in return. This unit is the one where you actually use them, deliberately, for something that matters to you. By the end, you will have built a small library of your own, five prompts for five tasks you personally care about, each one revised, improved, and reflected

on, all going into your Curious Notebook. (Where AI tools are not available in your classroom, the unplugged versions of these activities let you build the same library through planning and discussion.)

There is a small industry now selling 'prompt engineering' as a kind of magic. Most of it is not magic. Using an AI tool well means asking clear questions, giving enough context, and checking the answer carefully. You say what you actually want, you give the context that helps, and you respond to what you get back. The skill is mostly clear, honest thinking, not clever tricks. That is what this unit is about.

Three habits of computational thinking come together here. You will decompose a task into parts. You will design a stepwise method, an algorithm, in plain language, for working through it. And you will iterate: try, check, improve, try again. By the end of the unit, the three should feel like one habit rather than three.



Words You'll Meet in this Unit

This is the most practical unit in the module, and the vocabulary is mostly familiar by now. A few terms are worth flagging:

- Prompt is the question or instruction you give to an AI tool.
- Iteration is the cycle of try, check, improve, and try again. It is the habit at the centre of this unit.
- Helper-stance and replacement-stance are two ways to use an AI tool. In the helper-stance, you do the thinking and the tool helps you do it better. In the replacement-stance, the tool does the thinking instead of you. The unit explains why one is good for your learning and the other is not.
- The Curious Notebook is the running portfolio you have been building since Unit 1. It collects one learning task from each unit, four across Units 1 to 4, and a fifth from this unit if your class takes it up.

A Note on when this Unit is Used

This is the concluding unit of the module. It brings earlier learning together through direct work with AI tools or, where these are not available, through unplugged activities. In the Unplugged mode, students write prompts on paper, the teacher reads out sample AI replies, and the class discusses how to evaluate and improve them. The Teacher-Demo and Unplugged versions cover the same ideas as the device-based version, and the central skills of asking, refining, and reflecting can be learnt in any of the three modes. Teachers may choose the mode that suits their classroom.

5.1 Asking Better Questions

Many poor results from an AI tool come from unclear prompts, though better prompting cannot fix a tool's own limits. Four habits will help in most situations.

© NCERT not to be republished

- **Say what you actually want:** "Write something about Diwali" will get you a generic paragraph nobody asked for. "Write a 200-word welcome note for my school's Diwali celebration, addressed to parents, in a warm but slightly formal tone, mentioning that this year's celebration falls on a Friday afternoon" will get you something close to what you wanted. The second prompt is not cleverer. It is just clearer about the task.
- **Give context:** Who is the audience? What is the situation? What has already been tried? What makes your version of the task different from a generic one? An AI tool cannot read your mind. The moment you give it the context, the output gets better.
- **Show, don't just tell:** If you want a particular style, paste an example of it. If you want a particular structure, sketch the structure first. If you want a particular tone, give a sentence or two in the tone you have in mind. Examples work better than vague descriptions: "formal but warm" is harder for the tool to act on than a single sample sentence in that voice.
- **Ask for the work you want to see:** 'Summarise this article' and 'Summarise this article in five bullet points, each a complete sentence, focused on what the author argues rather than what they describe' will produce very different outputs. Tell the tool the shape you want, length, structure, point of view, and you will get closer to it.

5.2 Helper, Not Replacement

There is a stance towards AI tools that quietly determines how much you get from them: using them as helpers rather than replacements. The same tool, the same prompt, can produce a useful draft for one student and a hollow shortcut for another, depending on what the student does next.

- **Helper, in practice**: You write a first draft. The tool suggests ways to improve it. You read the suggestions, accept some, reject others, and modify several. The work that comes out is yours, with the tool's help. You learnt something while doing it: which of your sentences were unclear, which suggestions were good and why. The tool sped you up, but the thinking was still yours.
- **Replacement, in practice:** You ask the tool to write the whole thing. You read it once, change a few words, hand it in. You learned almost nothing about the topic, the writing, or yourself. The tool did the task; you acquired



What You have Learnt So Far

Four habits make for better prompts:

- Say what you actually want.
- Give context, the tool cannot guess the context.
- Show with examples, not just adjectives.
- Ask for the shape of the output.

Most of "prompt engineering" is just thinking clearly about what you want before you ask.



What You have Learnt So Far

Same tool, different stance. Used as a helper, it helps you grow; used as a replacement, it leaves you where you were.

Use AI where the judgement is still yours. Skip it where the work was the growth.

When you do use AI on submitted work, say so plainly, in one line at the top.

a piece of paper. Two months later, when a similar task arrives without the tool, you find you cannot do it.

- **The same tool:** Notice that the tool did not change between these two examples. Your stance toward it did. Same assistant, same prompts, same outputs: used as a helper, you grow; used as a replacement, you stay where you were while appearing to make progress.

Where AI tools are clearly worth using

Tasks that are useful but boring: Reformatting a list, generating practice problems, summarising long articles you might not have time to read carefully, etc.

Tasks where you need many options to choose from: Brainstorming, drafting variations of a sentence, suggesting different ways to organise a piece, etc.

Tasks where you want a sanity check: 'Does the reasoning hold together?' 'Am I missing an obvious objection?', etc.

Where AI tools are not the right answer

They are not the right answer for a task designed to build a skill you actually want to have: doing your own homework, writing your own draft, working through your own argument from scratch. If you let an AI do these for you, you give away the chance to build the very skills you came to school to develop. The tool may help you produce an acceptable output, but it will not help you become better at the skill itself.

Say it clearly. When you do use an AI tool on work you submit, say so. A line at the top, or a brief note where it fits: 'I used an AI assistant to brainstorm the structure of this essay, but the final argument and the words are mine.' 'I used an AI tool for language polishing.' 'I used an AI tool to generate practice questions; the answers below are my own.' This habit is good for three reasons. It is honest. It tells the reader what you did, so they can judge the work fairly. And it makes you articulate, to yourself, exactly what the tool did versus what you did, which is itself a form of clear, honest thinking.

5.3 Build your prompt library: A worked example

Time for the unit's main hands-on activity. The prompt library is your portfolio piece for this unit: five prompts on five tasks you personally care about, each iterated until it produces something useful, each annotated with your notes. Before you build yours, here is one student's, in full, so you can see the iteration in action.

(If your class does not have AI tools to hand, you can still do this on paper, writing and refining the prompts, and reasoning about what each version would need to work. The thinking is the same.)

Rohan in Bhubaneswar: Rohan is in Grade 10 in Bhubaneswar. His mother teaches a group of about twelve children in their neighbourhood, mixed ages, mostly first generation school goers. She covers Mathematics, basic English, and a little Science in the evenings after her main job. For a few weeks Rohan has been wondering whether AI tools could help her work. Not replace her, help her. He picks one specific task to start with: generating practice problems for the children, at the right level and in the right local context.



Fig. 5.1: Rohan helps his mother prepare practice problems

His first attempt: 'Give me 10 maths problems for kids.' The tool produces 10 problems. Some are too easy, some too hard, and the contexts are American, baseball scores, dollars, cookies, not at all useful for his mother's class.

- **Second attempt:** 'Give me 10 maths problems for Grade 4 students in India. Mixed difficulty.' This is better. The problems now use rupees and Indian names. But 'mixed difficulty' produced a wide spread, some far too easy, some far too hard, and the topics are scattered: a fraction, then a division, then a time problem, then money. No steady practice.

Third attempt: 'Give me 10 maths problems for Grade 4 students in India, all on the topic of multiplication of two-digit

numbers, ranging from simple to slightly tricky. Use Indian names and contexts such as cricket scores, vegetable prices, and school examples. Each problem is one or two sentences. Put the answer on a separate line at the end of each problem.'

This is now usable. The practice stays on one topic, the questions get harder step by step, and the situations are familiar. Rohan can imagine handing these to his mother's students.

Fourth attempt, after running attempt 3 with his mother: 'Give me 10 maths problems for Grade 4 students in Odisha, all on the multiplication of two-digit numbers, from straightforward to slightly tricky. Use contexts familiar to children in coastal Odisha: the fish market, school tiffin, festival sweets, cycle-rickshaw fares, Chilika Lake birds. Each problem is one or two sentences, with the answer on a separate line. Use kilograms and metres, not imperial units.' This is the version Rohan saved, not because it looks exciting, but because it helps with the real task his mother has.

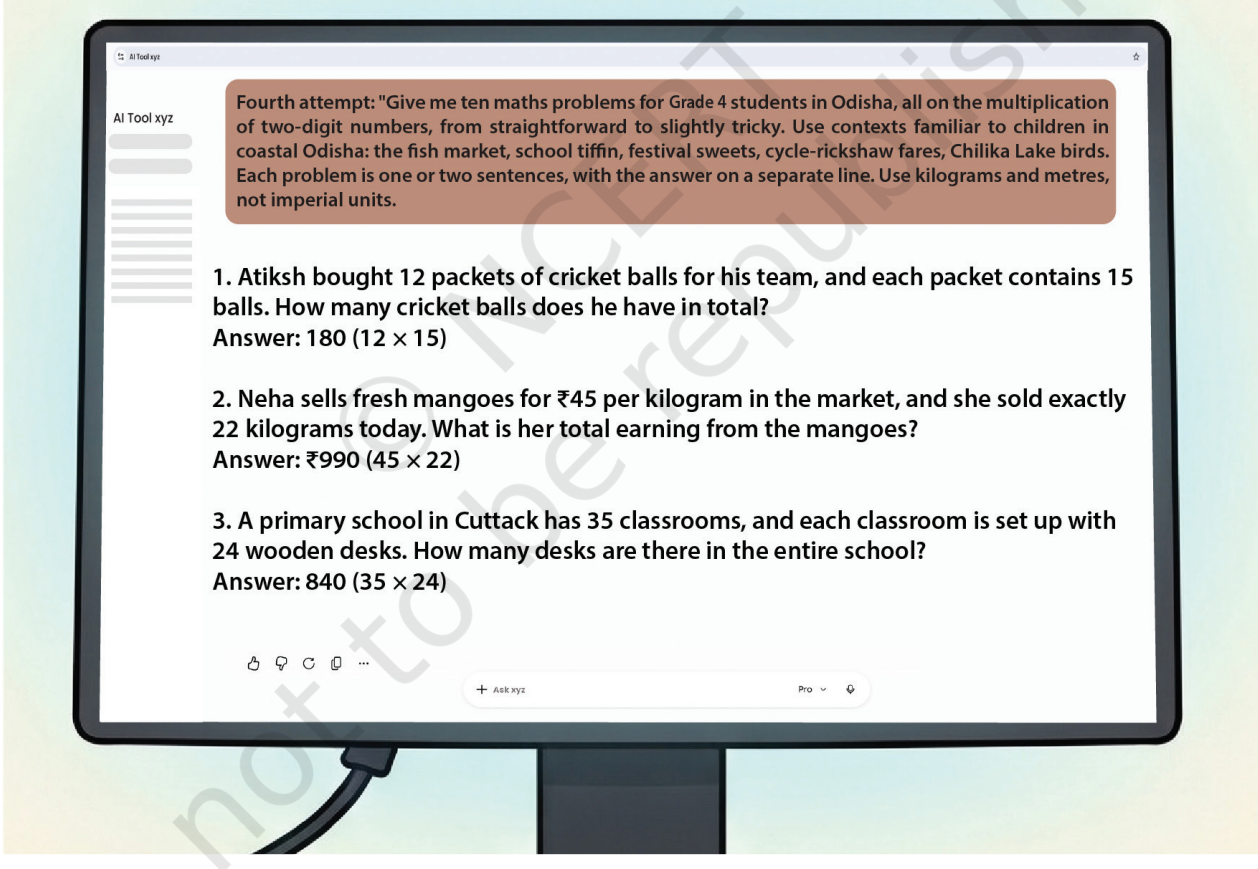


Fig. 5.2: What the tool produced for the fourth prompt (three of ten problems)

What changed across the four attempts: Each step added one of the four habits from Section 5.1. The first was vague.

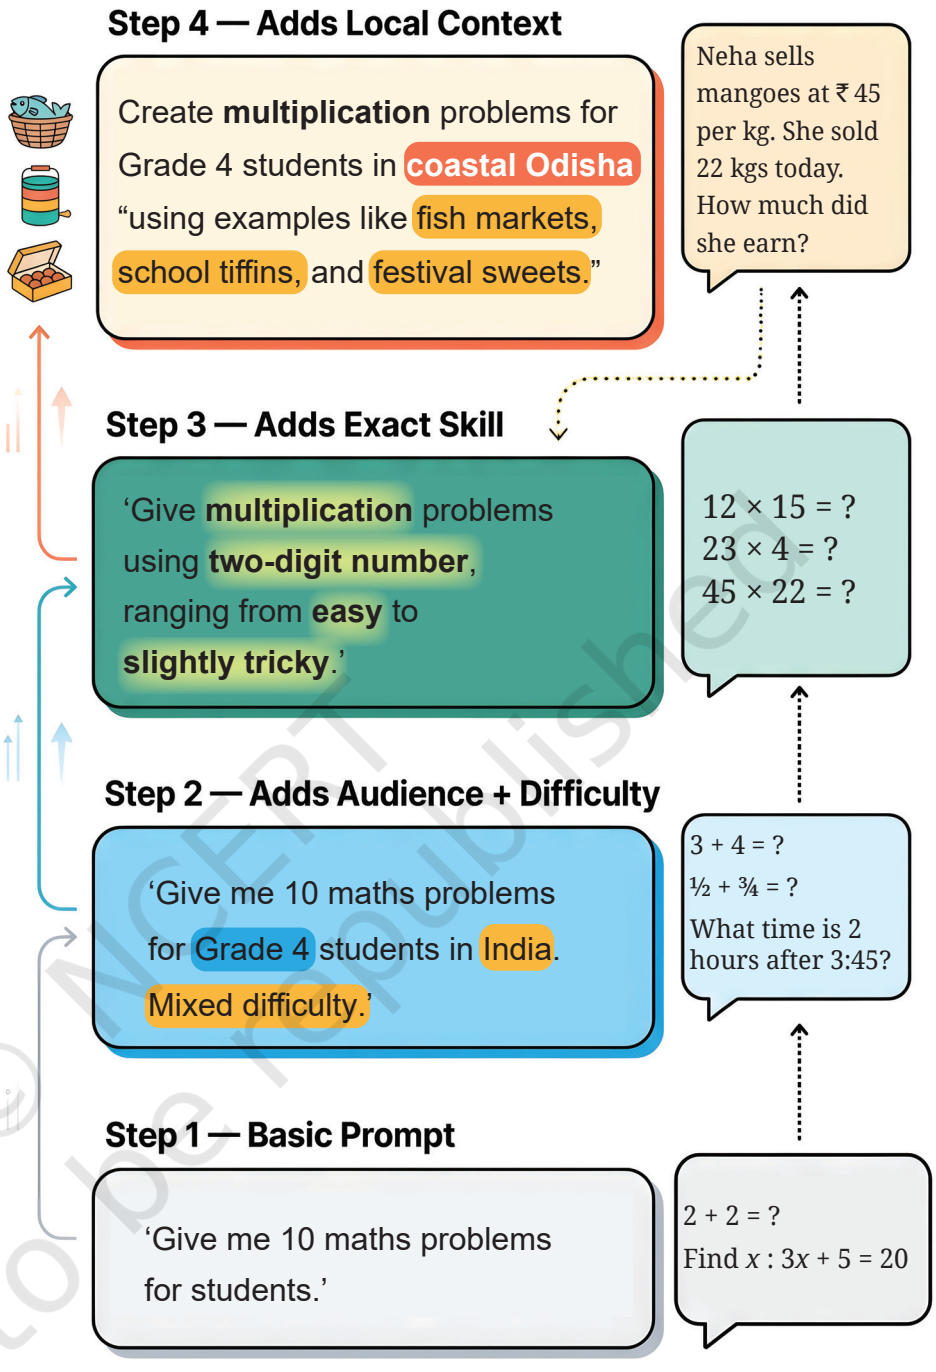
© NCERT not to be republished

The second added context (Grade 4, India). The third added structure (one topic, a difficulty arc, a fixed output format). The fourth added specificity drawn from seeing how the previous version actually worked in his mother's class. The fourth prompt is not cleverer than the first. It is just better-informed. That is the whole skill.

Doing this activity, whatever your classroom has. The activity below comes in three forms. They teach the same thing, the unplugged form needs no tools at all, so every class can do it; the other two use an AI tool where one is available. Your teacher will choose what fits.

UNPLUGGED (works for everyone). No tool needed. Working in pairs:

- One of you writes a prompt for a task, on paper, exactly as you would type it to an AI tool.
- Hand it to your partner. Your partner plays the AI: they answer using only what the prompt actually says, taken literally, not what they guess you meant. If the prompt is vague, the answer should be vague or slightly wrong, on purpose.



Better prompts come from better information about the task.

© NCERT not to be republished



What You have Learnt So Far

Iteration is the skill. The first prompt is rarely the right one. The third or fourth often is.

Better prompts come from better information about the task, not from cleverer phrasing.



Inclusivity Lens

AI tools often work better in some languages than others. If your task is in a language the tool was less trained on, such as Manipuri, Santali, Bodo, and many more, you may need more iterations to get a usable output, and you may need to fall back on a more-supported language and translate. This is unfair work, and it falls disproportionately on speakers of less-resourced languages. As you build your library, notice when the language you are working in is making the task easier or harder. The work to use AI tools well is not equal across users.

- Read what you got back. Where did your partner go "wrong"? Usually it is because the prompt left something unsaid.
- Rewrite the prompt to close that gap, and try again. Two or three rounds is enough.
- Then swap roles, so each of you gets a turn writing and a turn being the AI.

This teaches what "missing context" feels like more vividly than a real tool does, because your partner, like a real AI, can only work with the words you actually gave.

Teacher-demo: Improving a prompt together

The teacher runs a single prompt-improvement activity on the projector, taking suggestions from the class at each iteration, and everyone watches the prompt evolve in real time, discussing what changed at each step. As you watch, note in your Curious Notebook what was added at each stage (context? structure? an output format?) and why it improved the result.

FULL (a tool for each student or group). Working alone or in small groups:

- Pick five tasks you personally care about, perhaps one study task, one creative task, one to help a relative, one hobby, and one piece of school work.
- For each task, write a first prompt and run it through an AI assistant.
- Read the output critically. What is missing, vague, or off? (Use the Five Checks from Unit 3 to judge it.)
- Improve the prompt and run it again. Iterate at least twice per task.
- For each task, save every version of the prompt, the outputs, and a short note on what you changed and why.

These notes become your Prompt Library, your Curious Notebook piece for this unit.

5.4 Reflecting on what you made

Building five prompts is the visible part of the activity. Reflecting on what you built is where the lasting learning happens. Three questions are worth asking about every prompt you keep.

What did the tool actually do well? Be specific and concrete, not "it was helpful." Did it generate variety you would

not have thought of? Did it show you that something important was missing from your draft? Did it raise an opposing point you had not considered? When you can name what the tool actually contributed, you can use that contribution again on the next task. When you cannot name it, you may have been doing the work yourself and giving the tool the credit.

What did you have to fix? Every output needed something. Did the tone need adjusting? Did facts need checking (Unit 3's habit)? Did the structure need reorganising? Did the cultural context need correcting? The shape of what you fixed tells you where this tool is reliable, and where it is not, for tasks like this one.

What is the one thing you learned during the task? This is the question that most often gets skipped. If, at the end of a task with an AI tool, you cannot name something you learned, the task may have produced output but not learning. The good news is that the same task, done with reflection, usually does produce some. Even something small, like "I realised my first prompts are usually too vague," counts.

Inaya's three languages: Inaya, whom you met in Unit 1, noticing that her translation app flattened her jokes, took the prompt-iteration activity in three languages: Urdu at home, Hindi at school, and English in her community drama club. She tried the same task, writing a short comic monologue about a school assembly, in each. The English version came quickly. The Hindi version took a few more tries to get the rhythm right. The Urdu version needed the most guidance: she had to paste in examples of the kind of wordplay she wanted before the tool stopped making it too plain. "It is the same skill in all three," she told her cousin. "But Urdu showed me what the first language had hidden, how much the tool was doing for me, and how much still depended on me."

Reflect

Look back across the five units. Which habit do you think will stay with you the longest: noticing AI in daily life, understanding how machines learn from examples, verifying before you trust, decomposing a digital decision, or iterating a prompt? Two lines on which one and why.

Now name one habit that you have not yet formed. What would it take, in the next month, the next year, for it to become natural?



What You have Learnt So Far

Reflection is what turns activity into learning. Three questions are useful: what did the tool do well, what did you fix, what did you learn that you did not start knowing? If you cannot name something you learned, the task produced output but not learning.

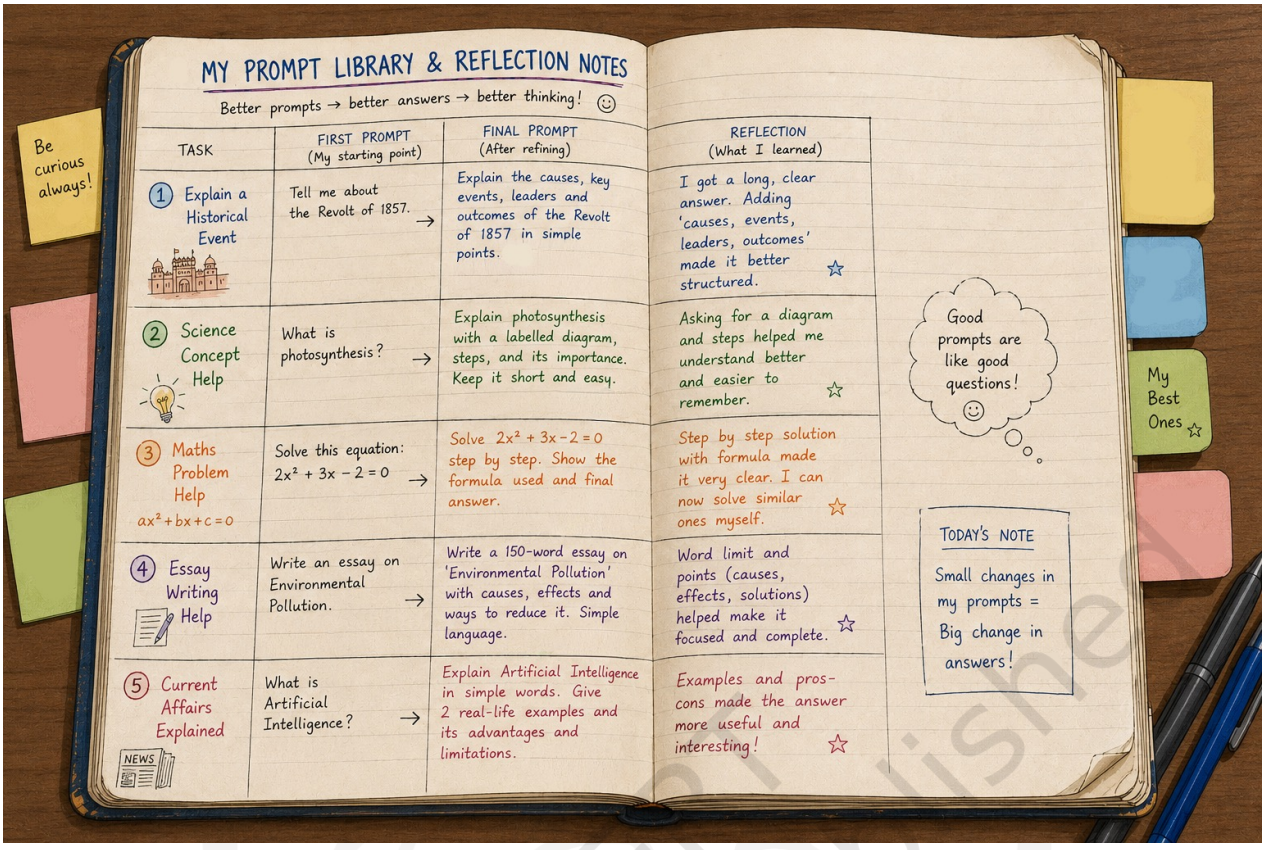


Fig. 5.3: Sample prompt library and reflection notes

For your Curious Notebook

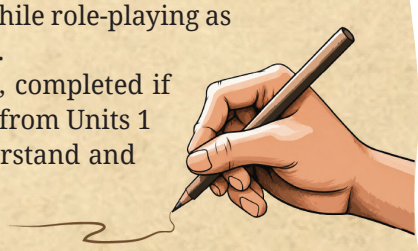
My Prompt Library and Reflection Notes: Five prompts on five tasks you personally care about. For each one, record in your Curious Notebook:

- The task in one sentence. What were you actually trying to achieve?
- Your first prompt, exactly as you wrote it.
- Your final prompt, exactly as you wrote it.
- Two or three lines on what changed between the two, and why.
- Two or three lines on what the final output got right, and what you still had to fix yourself.
- One sentence on what you learned from the task that you did not know when you started.

Aim for five tasks in all, and try to make at least one a study task, at least one a creative task, and at least one something that helps another person.

© NCERT not to be republished

If your class did the unplugged version of the activity, you can build your library the same way: your "outputs" are the responses your partner gave while role-playing as the AI, and the six points above work just as well on paper.
This is the fifth and final piece of your Curious Notebook, completed if your class takes up this unit. Together with the four pieces from Units 1 to 4, it completes your record of how you learned to understand and use AI.



Teacher-demo: Building the Library

The teacher runs a single AI assistant on the projector. Rather than only watching one shared prompt evolve, each student drafts their own prompt on paper, and the teacher types them in turn so the class sees each one run. You then improve your own prompt based on what came back, and dictate the better version for a second run. Record your iterations in your Curious Notebook. The skill being practised, thinking carefully about what you want before you ask, then improving the ask based on the result, is the same whether you typed the prompt yourself or watched it typed.

Assignment

Quick checks

1. The unit names four habits for asking better questions. List them in your own words, then add one of your own that you have noticed works.
2. Same tool, different stance: in one or two lines, explain in your own words the difference between using an AI tool as a helper and using it as a replacement.
3. The unit says “iteration is the skill.” Why is the third or fourth prompt usually better than the first? Answer in two lines.

Apply it

4. You are helping a younger cousin in Class VI prepare for a Hindi essay competition. The topic is “My Favourite Festival.” Walk through three iterations of a prompt you might give an AI assistant to help you brainstorm with your cousin, not to write the essay, but to help him think through what he might want to say. What does your prompt look like by the third iteration?
5. Think of one task from the past week where, looking back, you wish you had not used an AI tool, because doing it yourself would have helped you get better at something, and you skipped that chance. In half a page, write down what the task was, why you reached for the tool, and what you would do differently.

Stretch (optional)

6. Imagine you are designing an AI-powered study companion for a student two years younger than you. Write a half-page note: which of the habits you have built in this module would you most want the tool to encourage, and how would it do so? Avoid the easy answer ("it would just teach them"). What would the tool actually do, on a Tuesday evening, with a particular student, on a particular assignment?

© NCERT not to be republished

That brings the units to a close. The Curious Notebook is yours, the habits are yours, and the next stretch of the journey is yours too. The module ends with a short reflection on the computational thinking you have practised, and a few words on what you carry forward.

The computational thinking you practised

You have come to the end of the five units. Before you close the module, here is a short look back. The module began with a promise that you would not just learn about computational thinking, you would do it. This page is where we check on that promise.

Five moves were named in the front matter. Decomposition, pattern recognition, abstraction, algorithmic thinking, evaluation. Across the five units, each move showed up in something you actually did.

The move	Where you practised it
Pattern recognition	Unit 1, when you spotted AI across systems as different as a music recommender, a translation app, and a transaction-safety check, and noticed the shared pattern-matching structure underneath. Unit 2, when you saw that you and a learning system both build patterns from examples, and how the parallel works and where it breaks.
Abstraction	Unit 2, when you moved from many specific examples (the cats you met as a child) to a general idea ('cat'), and saw that learning systems do versions of the same move.
Evaluation	Unit 3, when you built the Five Checks, Source, Claim type, Compare, Pause, Confidence, and used them to catch confident-but-wrong outputs in your own Hallucination Hunt.
Algorithmic thinking	Unit 3, in designing the verification protocol as a stepwise method. Unit 5, in designing your prompt-iteration ladder, a method that turns a vague first prompt into a useful fourth.
Decomposition	Unit 4, when you broke a single big decision ("should I install this app?") into smaller decisions about data, consent, sharing, retention, and revocation. Unit 5, when you broke a task into the parts a prompt has to address, audience, context, format, constraints.

None of these moves is finished. Decomposition is not a thing you complete; it is a habit you keep. Pattern recognition gets sharper with practice. Evaluation gets faster. The five units were not the whole training. They were the beginning.

© NCERT not to be republished

There is one more thing worth saying. The module drew a parallel between you and the machine in almost every unit. The reason was not to flatter the machine. It was to use what you already know about your own thinking to make the machine's thinking visible. You learn from examples; so does a learning system. You produce false certainty; so does an AI assistant. You break big problems into smaller parts; so, by design, does a good system. By holding the parallel up, the module gave you a mirror to think with. The mirror does not stay behind in the module. It travels with you.

What you carry forward

Four habits are worth carrying from this module.

- **Notice:** Spot AI in your daily life. Spot when a system is automation pretending to be AI, and when an AI system is good at the surface but weak at the layer below. Spot when you are about to be wrong with confidence. Noticing is the first move of every other habit.
- **Verify:** Before you trust an AI output for anything that matters, run the Five Checks from Unit 3. Before you forward something, pause. Before you submit something, check.
- **Ask:** What data does this system want? What consent did I give? Who does it work well for, and who does it work less well for? What kind of task is this, and is it the kind the technology handles well? The questions take less time to ask than to recover from not having asked.
- **Iterate:** Your first prompt is rarely the right one. Your first draft is rarely the final one. Your first reading of a privacy policy will miss things your second one catches. Iteration is not a sign that something is wrong. It is the sign that something is being made well.

Your Curious Notebook stays with you after this module ends. The five portfolio artefacts (your AI Diary, your Train-and-Break Notes, your Hallucination Hunt Log and Personal Verification Checklist, your Privacy Policy Dissection, your Prompt Library and Reflection Notes) are not just assignments you handed in. They are a record of how you thought about AI at one point in your life. In a year, look back at them. You will be surprised how some of your worries were unwarranted, and how some things you took for granted look quite different in retrospect. That is what learning looks like, watched at a distance.

© NCERT not to be republished

Where this module sits in your journey

This module is the first NCERT publication on AI for secondary school students. It is meant to be the introduction; the fuller picture comes in the textbooks that follow. What you have learned here are the habits that will travel with you and help you explore the world of AI with curiosity, care, and confidence. What comes next is what you make of these units and the 'Curious Notebook' you built alongside them.

Wherever you go from here, AI tools will increasingly be part of it, whatever you study and whatever work you do. A doctor, a journalist, a designer, a lawyer, a farmer, a teacher: each will meet AI tools in their work, and each will have to make their own version of the choices in this module. The verification habit from Unit 3 works on any AI output, in any field. The decomposition habit from Unit 4 works on any digital decision. The prompt-iteration habit from Unit 5 works on any task. Keep practising them. They are not a school topic that ends with this module; they are life skills that begin with it.

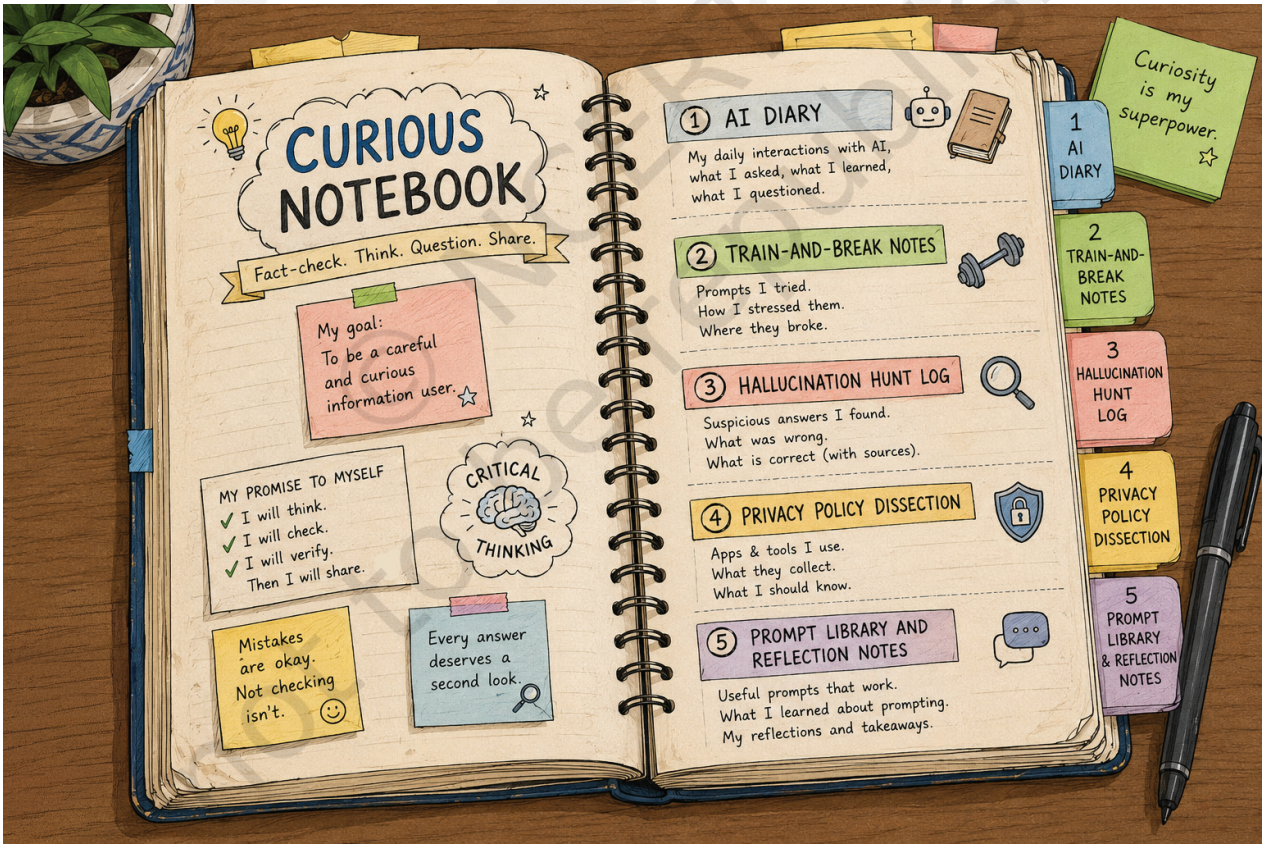


Fig. 5.4: Sample Curious Notebook

Glossary

The glossary presented below gives details about key terms introduced in the five units. The unit column shows where you first met the term. This referencing would be useful if you want to read the surrounding context.

Term	Meaning in this module	Where it first appears
Abstraction	Moving from many specific examples to a general idea.	Unit 2
Algorithmic thinking	Designing a stepwise method, then testing and improving it.	Front matter, Unit 3
Anchoring	When the first piece of information shapes how you interpret what follows.	Unit 3
Artificial intelligence (AI)	A broad field of computer science. In this module, we mostly mean systems that learn patterns from examples rather than only following rules someone wrote in advance — what we call "learning-based AI".	Unit 1
Automation	A system that follows fixed rules someone wrote in advance. FASTag is automation.	Unit 1
Automation bias	The tendency to over-trust output from an automated system because it looks confident, fluent, and expert-sounding. Called 'authority overtrust' in plain language in this module.	Unit 3
Behavioural data	Data about what you do online, clicks, time spent, what you skip.	Unit 4
Bias (in a learning system)	Unevenness in the training examples. Not a moral judgement, a description of skewed data.	Unit 2
Calibration	A system's reported confidence matches its actual rate of being right.	Unit 3
Category	One of the groups you sort examples into. "Cat" and "dog" are categories.	Unit 2
Confirmation bias	Weighing evidence that fits what you already believe more heavily than evidence that does not.	Unit 3
Consent	Real consent is specific, informed, voluntary, and revocable.	Unit 4

© NCERT not to be republished

Context (in a prompt)	Information about a task that an AI tool cannot guess, audience, situation, constraint, style.	Unit 5
Curious Notebook	Your folder of work across the five units, the five portfolio artefacts together.	Front matter, Unit 5
Data Fiduciary	Under the DPDP Act, the organisation that decides why and how your personal data is processed, alone or with others, and therefore has legal duties about handling it.	Unit 4
Data Principal	Under the DPDP Act, the person whose data it is, you.	Unit 4
Data Protection Board	The formal Indian body for data and privacy disputes under the DPDP Act.	Unit 4
Decomposition	Breaking a single big problem or decision into smaller, manageable parts.	Front matter, Unit 4
Deepfake	AI-generated video, audio, or images of a real person saying or doing things they did not.	Unit 3
Distribution shift	When real-use input differs from training input. Performance drops, often unnoticed.	Unit 3
DPDP Act	Digital Personal Data Protection Act, 2023, India's law on personal data handling. The DPDP Rules, 2025 are the operational rules that put the Act into practice.	Unit 4
Evaluation	Deciding whether an output is good, by which standard, in which situation.	Front matter, Unit 3
Example	One item a learning system learns from, a single photo, a single sentence, a single sound.	Unit 2
False negative	When a system fails to flag something that should have been flagged. A fraudulent transfer that slips through the safety check.	Unit 2, Unit 3
False positive	When a system flags something that should not have been flagged. A routine UPI payment held by the safety check for verification.	Unit 2, Unit 3
Feature	A particular detail you (or a system) pay attention to in order to recognise something.	Unit 2
Grievance officer	The company's named point of contact for complaints about how your data is handled.	Unit 4

© NCERT not to be republished

Hallucination (in an AI assistant)	Confident, fluent output containing invented facts, names, or quotations.	Unit 3
Helper stance	Using an AI tool to support your own work, thinking and judgement remain yours.	Unit 5
Identifying data	Data that can identify you specifically among many people, name, face, fingerprints, Aadhaar number.	Unit 4
Inferred data	Conclusions a system has drawn about you from other kinds of data.	Unit 4
Iteration	The cycle of try, observe, improve, try again.	Unit 5
k-Nearest Neighbours (KNN)	One of the simplest supervised-learning methods. For a new example, find the k most similar examples in the training set and let them vote.	Unit 2
Label	The right answer attached to an example, used to train a learning system.	Unit 2
Layer	A row or group of neurons (biological or mathematical) that processes information at one stage.	Unit 2
Learning system	A system that learns patterns or internal parameters from examples rather than only following rules written by hand. Bhashini is built around learning systems.	Unit 2
mKisan	A government mobile advisory platform that sends customised farming information, alerts, and guidance to farmers through text or voice messages.	Unit 1
Neural network	A mathematical structure used in many learning systems, loosely inspired by the brain but not the same as a brain.	Unit 2
Neuron	A cell in the nervous system that takes signals, processes them, and passes them on. (The mathematical version is a borrowed metaphor.)	Unit 2
Overfitting	When a system memorises its training examples but fails on anything new.	Unit 2
Pattern recognition	Noticing that something new is similar to something you have seen before.	Front matter, Unit 1

© NCERT not to be republished

Personal data	Data that can identify you.	Unit 4
Prompt	The input you give an AI assistant, question, request, description.	Unit 5
Reinforcement learning	Learning by trying, the system takes actions, gets feedback from the environment, and updates. A baby learning to walk is the natural analogy.	Unit 2
Replacement stance	Using an AI tool to do work in your place. Useful for some tasks, harmful for others.	Unit 5
Rule-based system	Another name for automation when discussing computers. mKisan is mostly rule-based.	Unit 1
Self-supervised learning	A fourth way machines learn from examples, in addition to the three covered here. The system finds its own training signal in the structure of the data, for example, by predicting the next word in a text. Modern large language models work this way.	Unit 2
Sensitive personal data	Kinds of personal data that can cause greater harm if misused, health, financial, biometric, identity-related. Exact legal categories vary; the practical point is extra care.	Unit 4
Simple device	A mechanism that responds to a physical condition without making decisions or learning.	Unit 1
Supervised learning	Learning from labelled examples, each example comes with the right answer attached. A very common kind in everyday AI systems.	Unit 2
Synthetic media	Any media, video, audio, or images, generated or altered by AI rather than recorded from the world.	Unit 3
Token	A small piece of text processed by a language model. A token may be a whole word, part of a word, punctuation, a number, or a space. Language models usually generate text by predicting likely next tokens.	Unit 2
Training data	The examples a learning system is shown so that it can find patterns.	Unit 2
Unsupervised learning	Finding structure in examples without being told the right answer. The system groups, clusters, or notices patterns on its own.	Unit 2
Verification protocol (the Five Checks)	A short stepwise method for checking whether to trust an AI output, Source, Claim type, Compare, Pause, Confidence.	Unit 3

© NCERT not to be republished

विद्यया ऽ मृतमश्नुते
© NCERT not to be republished
एनसीईआरटी
NCERT

विद्यया ऽ मृतमश्नुते
© NCERT not to be republished
एनसीईआरटी
NCERT

Hello Children!

If you feel uneasy about someone touching you inappropriately, you should not keep quiet. You must

1. Not blame yourself
2. Tell someone whom you trust
3. You can also inform National Commission for Protection of Child Rights through the **POCSO e-box**.

When you get an unsafe touch, you may feel bad, confused and helpless
You need not feel "bad" because it's not your fault



Press This Button

POCSO e-box available at NCPCR@gov.in

नए समाज की ओर
Towards a new dawn

If you are below 18 years of age, and are troubled or confused or abused or in distress or know some other child who is...

*Call **1098**...because some numbers are good!*
They change lives!!!



CHILD LINE
1098
NIGHT & DAY

CHILDLINE 1098 - a national 24 hours toll free emergency phone service for children in distress is an initiative of CHILDLINE India Foundation supported by Ministry of Women & Child Development

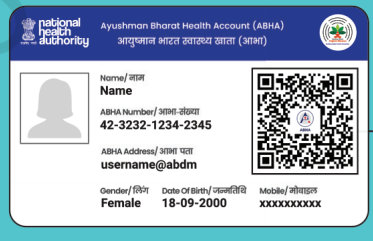
नए समाज की ओर
Towards a new dawn



Ek Kadam Swachchhta ki or

national health authority

ABHA
Your trusted partner in healthcare



Scan to download the **ABHA App** and create your **ABHA**

Scan to see how **Tau Ji** created his **ABHA**

Ayushman Bharat Digital Mission

Call **14477** for more info (Toll-Free)

© NCERT not to be republished

Scan, download, and explore endless stories with Rashtriya e-Pustakalaya – your free digital library!

© NCERT not to be republished

UN514

विद्यया ऽमृतमश्नुते

एन सी ई आर टी
NCERT

राष्ट्रीय शैक्षिक अनुसंधान और प्रशिक्षण परिषद्
NATIONAL COUNCIL OF EDUCATIONAL RESEARCH AND TRAINING

ISBN 978-93-5729-286-3