Grade 8

FfathomKey

AI Series

The AI ADVENTURE

Future - Ready kids

Think
Smart
Build
Better

An Interactve Journey into AI and Algorithms

NEP 2020 • NCF-SE 2023 • CT & AI Competency Aligned

Young Problem Solver

I solve problems step by step.

Technology Explorer

I explore technology and use it wisely

Responsible Digital Citizen

I use technology safely and kindly

Authored by alumni of IIM Bangalore and University of Texas

FathomKey

The AI Adventure -Grade 8
First Edition: 2026

Published by: Fathomkey
Website: www.fathomkey.in
Email: info@fathomkey.in
Written by: Nadia Afreen
Editorial Support: Fathomkey Academic Team
Designed and Illustrated by: Nadeemur R. Choudhary
ISBN: 978-81-180798-5-9
Printed in India by:

Curriculum Alignment
This book is part of a grade-wise AI learning programme developed for Indian schools. Its content is informed by the principles of **NEP 2020, NCF** and relevant **CTAI** guidance. It supports the development of computational thinking, problem-solving, digital literacy, creativity and responsible use of Artificial Intelligence across school curricula.

LMS Support
This book is supported by the FATHOMKEY Learning Management System (LMS), which provides digital learning resources, interactive activities, assessments and additional support for students and teachers. LMS access is available through registered schools. For access and assistance, visit https://lms.fathomkey.in or contact info@fathomkey.in.

Disclaimer
This book is intended for educational purposes. AI technologies, digital tools, interfaces and features may change over time. Teachers and students should verify important AI-generated information using reliable sources and follow appropriate privacy and online-safety practices.
All examples, names, datasets and situations used in this book are fictional or included solely for learning. Product names, logos and trademarks mentioned in the book belong to their respective owners. Their inclusion does not imply endorsement, sponsorship or affiliation.

FATHOMKEY
AI for everyone.

The AI ADVENTURE

Grade 8

An Interactive Journey into AI and Algorithms

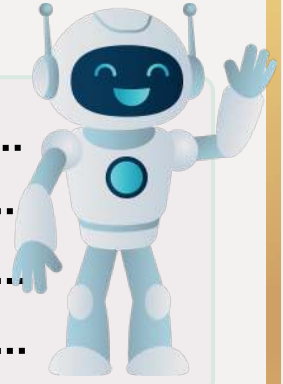

Name: .....

Class: ..... **Sec:** .....

Subject: ..... **Roll No:** .....

School: .....

CTAI STUDENT BOOK

FathomKey

AI book by IIM Bangalore & Austin University Alumni

PREFACE

Artificial Intelligence is transforming how people analyse information, recognise patterns, solve problems, create content and make decisions. To participate confidently in this changing world, students must understand how AI systems work, evaluate their outputs critically and use them responsibly.

AI for Young Minds – Grade 8 develops this understanding through a structured and age-appropriate exploration of the AI project lifecycle. Students learn to identify a problem, define project goals, prepare suitable data, select an AI approach, develop a solution and evaluate its performance.

The book strengthens learners' understanding of data, machine learning and intelligent systems. It introduces concepts such as features, labels, training and testing data, classification, prediction, accuracy, confidence and error analysis. Students also examine how data quality and human choices can influence AI results.

Learners explore data science, computer vision, natural language processing and generative AI. They discover how tools such as ChatGPT can support research and content creation, while platforms such as Lovable can help transform ideas into prototypes. Students are encouraged to verify information, protect sensitive data and retain control over final outputs.

Responsible AI is integrated throughout the book. Students investigate fairness, bias, privacy, security, transparency, accountability, copyright and the environmental impact of technology. They learn that AI may produce incomplete, inaccurate or unfair results and that human supervision remains essential.

Students also examine why deployed AI systems require monitoring, maintenance and improvement. They learn how changes in data, user behaviour or real-world conditions can affect performance and why a system may need to be updated, paused or retired.

Activities, case studies, exercises and AI Labs connect concepts with practical investigation. Learners think independently, use only fictional, anonymous or teacher-approved information, test solutions under different conditions and support conclusions with evidence.

The concluding project journey brings together problem-solving, data literacy, responsible design, prototyping, testing and reflection. It develops creativity, collaboration, communication and informed decision-making.

We hope this book inspires every learner to explore Artificial Intelligence with curiosity, question its outputs thoughtfully and create technology that is useful, inclusive and responsible.

— **Team Fathomkey**

Inside the Series

🎯 What You Will Learn
The section tells you what you will learn in the chapter and the skills you will develop.

🎨 Activity Corner
Enjoy hands-on activities, drawing, puzzles, and mini projects to reinforce your learning.

🤖 Byte's Smart Tip
Byte shares smart tips to help you to solve problems using Computational Thinking.

📖 Let's Begin with Introduction
Every chapter begins with an interesting introduction that connects the lesson to everyday life.

💎 Real-Life Connection
Discover how Computational Thinking helps us to solve problems in our daily lives.

📘 Definition
Learn important words and their meanings in simple and easy language.

💻 AI Connection
Explore how Artificial Intelligence and smart machines use instructions to complete tasks.

⭐ Did You Know?
Read amazing facts that make learning more interesting and fun.

🧠 Think About It
Answer thought-provoking questions that encourage observation, reasoning, and discussion.

🏆 Challenge Yourself
Apply your learning through HOTS, analytical thinking, and real-life problem-solving activities.

📝 Exercise Time
Practice what you have learned through fun and engaging questions.

📚 Let's Revise
Quickly revise the chapter using a summary, glossary, and a simple mind map.

CONTENTS

UNIT 1

Introduction to *Artificial Intelligence*

Chapter 1

Learning Outcome 🎯

✅ Define Artificial Intelligence in their own words and identify its common applications in everyday life.
✅ Distinguish AI systems from ordinary rule-based computer programs.
✅ Compare the strengths of human intelligence and Artificial Intelligence.
✅ Explain at an introductory level how an AI model learns from data.
✅ Describe the benefits and limitations of AI and suggest responsible ways of using AI systems.

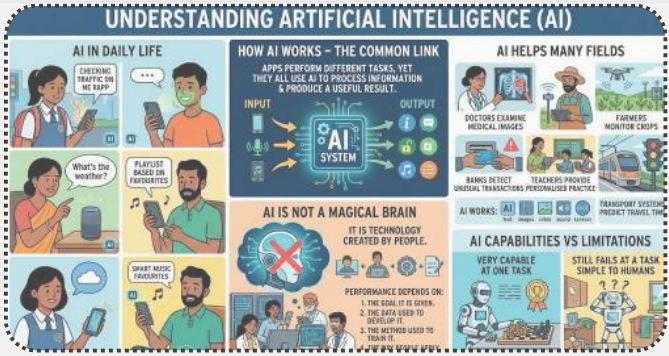

Chapter Introduction

Living in a World of AI

Before leaving for school, Meera checks a navigation app to see whether there is heavy traffic. Her brother unlocks a smartphone by looking at its screen.

Their mother asks a voice assistant about the weather, while a music application prepares a playlist based on songs the family has enjoyed before.

These applications appear to perform different tasks, yet they have something in common: each uses Artificial Intelligence, or AI, to process information and produce a useful result.

AI is no longer limited to research laboratories or science-fiction films. It helps doctors examine medical images, farmers monitor crops, banks detect unusual transactions, teachers provide personalised practice, and transport systems predict travel time. It can work with text, numbers, images, video, sound and information collected by sensors.

However, AI is not a **magical electronic brain**. It is technology created by people. Its performance depends on the **goal** it is given, the **data** used to develop it, the **method** used to **train** it and the way people apply its output.

An **AI system** can be very capable at one task and still fail at another task that seems simple to a human.

Keywords

Artificial Intelligence • Data • Algorithm • Machine Learning • Model • Training • Pattern • Prediction • Bias • Privacy

What Is Artificial Intelligence?

Artificial Intelligence is a field of computer science concerned with building systems that can perform tasks that usually require aspects of human intelligence. Such tasks may include recognising objects, understanding spoken or written language, finding patterns, solving problems, making predictions and recommending actions.
An ordinary computer program often follows rules written directly by a programmer. For example, a basic calculator follows exact mathematical instructions. When the same numbers and operation are entered, it produces the same result.

Three Important Definitions

Many modern AI systems work differently. Instead of receiving a separate instruction for every possible situation, they learn useful patterns from examples. A spam filter, for instance, can study large numbers of messages labelled “spam” or “not spam.” From these examples, it develops a model for estimating whether a new message is likely to be unwanted.
This does not mean that every AI system learns continuously or works without rules. AI is a broad field. Some systems use human-written rules, some learn from data, and many combine both approaches.

Artificial Intelligence (AI): The field of creating computer systems that can perform tasks associated with human intelligence, such as recognising patterns, understanding language or making predictions.

Algorithm: A clear sequence of instructions used to solve a problem or complete a task.

Machine Learning (ML): A method within AI in which a computer system learns patterns from data instead of being given every rule directly.

AI, Machine Learning and Generative AI

Artificial Intelligence is the broadest term. **Machine Learning** is one approach used to build AI systems. **Generative AI** is a type of AI that can create new content—such as text, images, audio, video or computer code—by using patterns learned from large collections of examples.
Generative AI can be helpful for brainstorming, explaining ideas and creating drafts. Its output can also contain errors or invented information. Therefore, important facts should always be checked using reliable sources.

KNOWLEDGE BOX **Does AI “think”?** 

Generative AI can be helpful for brainstorming, explaining ideas and creating drafts. Its output can also contain errors or invented information. Therefore, important facts should always be checked using reliable sources.

Ordinary Programs and AI Systems

The boundary between an ordinary program and an AI system is not always perfectly sharp. The comparison below shows a common difference.

Feature	Ordinary rule-based program	Data-driven AI system
How it is built	A programmer writes explicit rules	Developers select a learning method and train a model using data
Best suited for	Clear, fixed and predictable tasks	Tasks involving complex patterns or uncertainty
Example	Calculating the total of five numbers	Recognising an object in a photograph
Response to new situations	Follows the programmed rules	Uses learned patterns to estimate an answer
Possibility of error	May fail when a rule is missing or incorrect	May produce an incorrect prediction even when the software runs normally

Example: Sorting Fruit

Imagine a machine that must separate apples from oranges.
In a rule-based approach, a programmer might write: "If the fruit is orange in colour, place it in the orange basket." This rule could fail when an Apple is green or an apple is reddish-pink
In a machine-learning approach, the system studies many labelled images of apples and oranges. It may learn to use several features—including colour, shape, texture and size—to classify a new fruit.
Its answer is a prediction based on learned patterns, not a guarantee.

THINK & DISCUSS **Does AI "think"?**

A digital clock displays the time, while a face-unlock system identifies its owner. Which one is more likely to use AI? Explain what makes the tasks different.

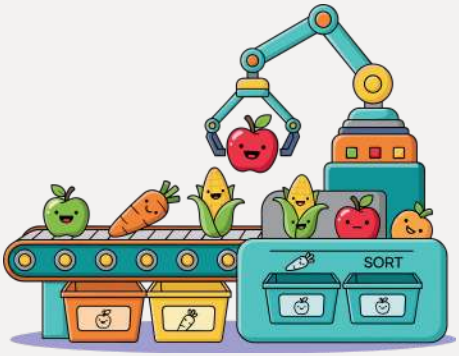

Human Intelligence and Artificial Intelligence

AI is inspired by certain abilities associated with human intelligence, but the two are not the same. Humans and machines have different strengths.

Human intelligence	Artificial intelligence
Learns from experience, instruction, observation and social interaction	Learns from data, programmed rules or a combination of both
Uses common sense and understands context across many situations	Usually performs best within the task and conditions for which it was designed
Experiences emotions and can show empathy	Does not experience emotions, even if it produces emotional language
Can consider values, relationships and moral responsibility	Optimises according to its design and requires human guidance for ethical decisions
Can transfer knowledge flexibly to unfamiliar situations	May struggle when a new situation differs greatly from its training data
Works more slowly when analysing enormous datasets	Can process large quantities of data quickly
Can become tired or distracted	Can repeat a defined task consistently without fatigue

AI may outperform people at a narrow task, such as finding a pattern in millions of records. A person, however, can combine common sense, empathy, life experience and moral judgement when making a decision.

This is why important areas—such as healthcare, education, recruitment and law—need meaningful **human oversight**.

Human–AI Collaboration

The most useful question is often not "**Will AI replace humans?**" but "**How can people and AI work together responsibly?**"

A doctor may use an AI tool to highlight a suspicious area in an X-ray, but the doctor examines the patient's complete condition and makes the final medical decision.

The **AI provides assistance**; the trained professional remains responsible.

Characteristics of AI Systems

AI systems do not have the same abilities. Depending on their purpose, they may show one or more of the following characteristics.

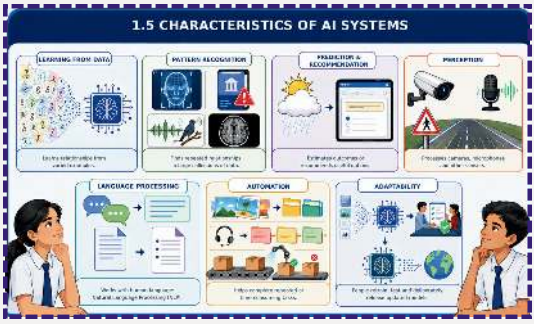

1. Learning from Data

Many AI systems **learn relationships from examples**. A handwriting-recognition model can be trained using thousands of images of handwritten letters and numbers. The quality, relevance and variety of those examples strongly affect its performance.

2. Pattern Recognition

AI can **identify repeated relationships** in large collections of **data**. It may recognise a face in an image, detect an unusual banking transaction, classify an animal sound or find a possible sign of disease in a medical scan.

3. Prediction and Recommendation

After learning from earlier examples, an AI model can estimate what may happen next or which option may be useful. A **weather model** can support rainfall prediction, while a learning application can recommend the next practice question.

4. Perception

Some AI systems process information from cameras, microphones or other sensors. They may identify objects, transcribe speech, detect movement or interpret road signs. This field includes **computer vision** for images and **speech recognition** for spoken language.

5. Language Processing

AI can work with human language. It can translate text, answer questions, summarise a passage or identify the topic of a message. Working with language is called **Natural Language Processing (NLP).**

6. Automation

AI can help complete repeated or time-consuming tasks, such as organising photographs, sorting support requests or checking products for visible defects.
Automation is most valuable when the task and responsibility are clearly defined.

7. Adaptability

Some AI models can be retrained or updated when new data becomes available. This is different from saying that every AI system improves automatically.
In many cases, people must collect new data, test the model and deliberately release an updated version.

Chapter Check-Up

1. In one sentence, explain why AI is not magic.
2. What is the relationship between AI and Machine Learning?
3. Give one task suited to a rule-based program and one suited to a data-driven AI system.
4. Why does an important AI-supported decision still require human oversight?
5. Name any three characteristics of an AI system.

How Does AI Learn?

Humans learn through instruction, observation, practice and experience.
A machine-learning model learns differently. It uses a mathematical process to detect patterns in data.

Data, Model, Training and Prediction

Data:
Facts or observations collected in forms such as numbers, words, images, sounds or sensor readings.

Model:
A mathematical representation of patterns learned from data. It is used to produce a prediction or another output.

Training:
The process of using data to adjust a model so that it performs a particular task.

Prediction (or inference):
The output produced when a trained model receives new input.

A Simple AI Learning Process

1. Define the problem: Decide exactly what the system should do.
2. Collect data: Gather relevant examples in a lawful and responsible way.
3. Prepare the data: Remove errors, organise the examples and add labels when required.
4. Train the model: Allow an algorithm to identify patterns in the training data.
5. Test the model: Check its performance using examples it did not see during training.
6. Use the model: Provide new input and receive a prediction or recommendation.
7. Monitor and improve: Observe errors, check fairness and retrain the model when necessary.

Flow for the designer:

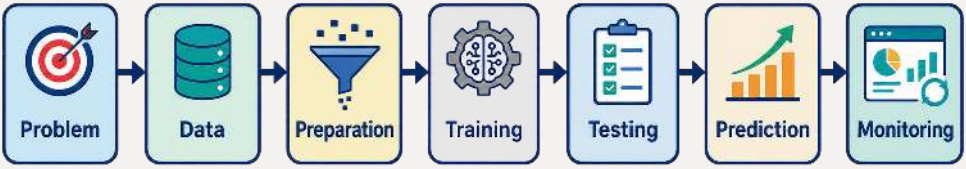

Suppose a school wants a smart bin to identify whether an item is paper, plastic or metal. First, the team defines the task: classify one visible item into one of three categories. It collects photographs of different paper, plastic and metal objects under varied lighting and from different angles. Each image is given the correct label.

During training, the model searches for useful visual patterns such as shape, surface and colour.

The team then tests it using new photographs that were not used during training.

If the model performs poorly on crushed cans or transparent plastic, the team may need more varied examples and another round of training.

When the trained model receives a new photograph, it might produce:

Paper: 8%

Plastic: 87%

Metal: 5%

The highest score suggests "**plastic**," but the score does not prove that the answer is correct.

AI predictions involve uncertainty.

Important:
Testing a model on its training examples is like giving a student the same questions used during practice. To measure learning fairly, the test should include new examples.

Data Quality Matters

A model learns from the examples it receives. If its data is incorrect, too small, one-sided or unrelated to the real task, its output may be unreliable.

This idea is often described as "**Garbage in**, **Garbage out**." Better quantity alone is not enough,data must also be accurate, relevant, representative and collected responsibly.

AI Around Us

AI is used in many places, although not every digital device or automated machine is an AI system.

Area	Example use of AI	What the AI does
Home	Smart speaker	Recognises speech and responds to a request
Education	Adaptive learning platform	Recommends practice based on a learner's performance
Healthcare	Medical-image support tool	Highlights patterns for a trained professional to examine
Transport	Navigation application	Estimates travel time and suggests routes
Banking	Fraud-detection system	Flags transactions that differ from expected patterns
Agriculture	Crop-monitoring system	Examines images or sensor data for possible crop stress
Entertainment	Video or music platform	Recommends content based on user behaviour and other signals
Manufacturing	Visual quality inspection	Detects possible defects in products
Environment	Wildlife camera system	Identifies animals in large collections of images

Is Every Smart Device AI?

No. A motion-sensor light may switch on whenever movement is detected.

If it simply follows the rule "**movement detected → light on**," it is automated but not necessarily intelligent. A more advanced system might learn when rooms are normally occupied and adjust lighting to save energy. The presence of a sensor or automation alone does not prove that a system uses AI.

Try It Yourself: AI or Ordinary Automation?

Classify each example and explain your reasoning.

1. A washing machine runs for exactly 30 minutes after a button is pressed.
2. An email service estimates whether a new message is spam.
3. A school bell rings at fixed times stored in a timetable.
4. A phone groups photographs by the people appearing in them.
5. A calculator multiplies two numbers using fixed instructions

Four Worked Examples from Everyday Life

Worked Example 1: Navigation and Route Suggestions

When a user enters a destination, a navigation application compares possible routes. It may use map data, current traffic signals, reported road conditions and patterns from earlier journeys to estimate travel time. If conditions change, the application can calculate another route.

Input: Starting point, destination, map and traffic-related data
AI task: Travel-time prediction and route recommendation
Output: A suggested route and estimated arrival time
Human role: The traveller decides whether the route is safe and suitable

A video platform may use viewing history, searches, likes, subscriptions and other signals to estimate which videos a person may want to watch. It compares patterns across content and viewing behaviour, then ranks possible recommendations.

Worked Example 2: Video Recommendations

Input: Viewing and content-related data
AI task: Pattern matching and ranking
Output: A personalised list of videos
Human role: The user chooses what to watch and manages privacy and recommendation settings

Recommendations can be convenient, but they may also repeatedly show similar content. Users should actively explore different reliable sources instead of allowing an algorithm to make every choice.

Worked Example 3: Face Unlock

During enrolment, a face-unlock system captures features of the authorised user's face and creates a numerical representation.
When someone tries to unlock the device, the system compares newly captured facial information with the stored representation and estimates whether they match.

Input: Facial information captured by the device
AI task: Pattern comparison
Output: Match or no-match decision
Human role: The owner protects the device and uses a secure backup method

Face recognition is not perfect. Lighting, camera angle, appearance changes and the design of the system can affect performance. Biometric information is sensitive personal data and must be protected.

Worked Example 4: Voice Assistants

When a person says, "**Set an alarm for 6:30 a.m.**," a voice assistant completes several steps. It captures the sound, converts speech into words, identifies the requested action and extracts the time. The device then uses ordinary software to create the alarm.

Input: Spoken command
AI tasks: Speech recognition and language processing
Output: Interpreted instruction
Human role: The user checks that the assistant understood the command correctly

This example shows that one application can combine AI components with ordinary programmed functions.

Chapter Check-Up 2
Answer the questions:

1. Is the wrong label necessarily a software crash? Explain.
2. Which learned patterns might have confused the model?
3. What kind of new training examples could help improve it?
4. Why should users check an AI-generated description before relying on it?

Read the situation below.

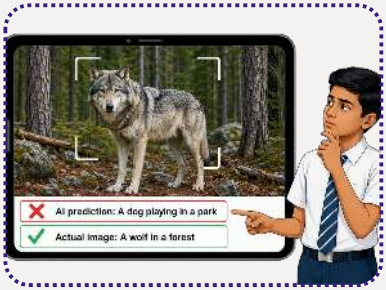

Benefits of Artificial Intelligence

Speed and Scale

AI can process far more records or images than a person could examine in the same amount of time.
This is useful when people need help finding important patterns in large datasets.

Consistency

A well-designed system can apply the same method repeatedly without becoming tired. Consistency does not guarantee correctness, but it can reduce variation in repetitive tasks.

Automation of Repetitive Work

AI can help sort documents, organise information and perform initial checks. This allows people to spend more time on creative work, communication and complex decisions.

Personalisation

Learning platforms can adjust practice questions, accessibility tools can support different user needs, and entertainment platforms can recommend content based on preferences.

Accessibility

Speech-to-text, text-to-speech, image descriptions and real-time captions can make digital information easier to use for people with different abilities.

Support in Risky Environments

AI-enabled robots or inspection systems can support work in mines, disaster areas, deep oceans or industrial environments where conditions may be dangerous for humans.
The value of AI depends on how it is designed and used. A fast system is not useful if it produces unfair, unsafe or irrelevant results.

Limitations and Responsible Use of AI

AI Can Be Wrong

An AI model predicts an answer from patterns; it does not guarantee the truth.
A language model may produce a confident but incorrect statement, while an image model may misidentify an object.

Dependence on Data

Incomplete, incorrect or unrepresentative data can reduce accuracy.
A model trained mainly on bright daytime photographs may perform poorly at night.

Bias and Fairness

Bias is a systematic tendency that can lead to unfair results. Bias may enter through unbalanced data, historical inequalities, unsuitable labels or design decisions.
Developers should test performance across relevant groups and situations.

Lack of Human Understanding

AI does not possess human experience, empathy or moral responsibility. It can generate a polite response without feeling concern and can imitate reasoning without understanding a situation as a person does.

Privacy and Security

Some AI applications use personal data such as location, voice, photographs or browsing behaviour.
People should know what information is collected, why it is needed, how it is protected and whether they have a meaningful choice.

Cost and Resources

Developing and operating advanced AI can require skilled people, computing equipment, electricity and careful maintenance. Schools and organisations should consider whether the benefit justifies these resources.

Need for Human Oversight

High-impact decisions should not be left blindly to an automated system. Humans must review important outputs, allow mistakes to be challenged and remain accountable for final decisions.

A Responsible AI Checklist for Students

Before using AI-generated information, ask:

- **Purpose**: Am I using the tool for an appropriate task?
- **Privacy**: Am I sharing personal, confidential or sensitive information?
- **Accuracy**: Have I checked important claims using reliable sources?
- **Fairness**: Could this output be one-sided or unfair?
- **Ownership**: Am I respecting copyright and acknowledging help received?
- **Responsibility**: Am I making the final decision instead of blindly following the tool?

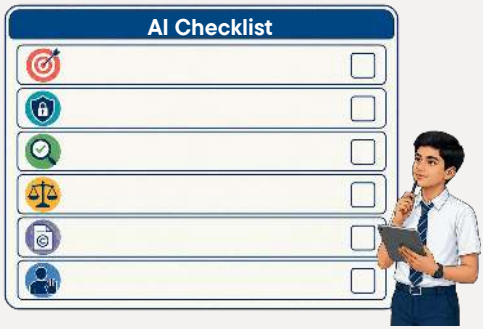

Digital Citizenship Note

Never upload another person's photograph, voice, school record or private message to an AI service without appropriate permission. Follow your school's rules and the service's age requirements.

Real-Life Case Study: How a Recommendation System Learns Your Interests

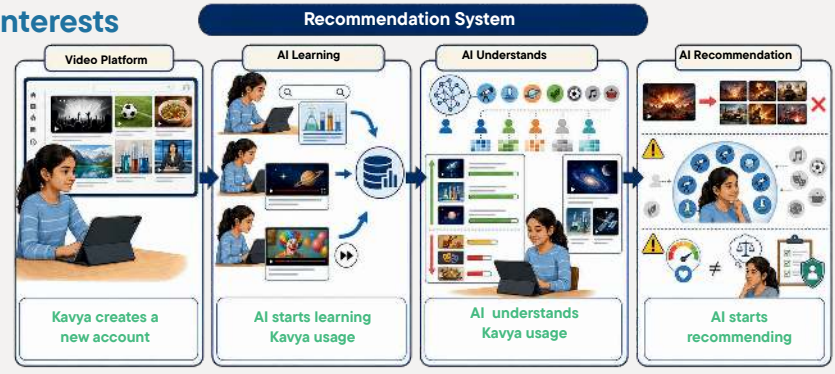

Imagine that Kavya creates a new account on a video platform. At first, the platform knows very little about her interests, so its recommendations may be broad or based on generally popular content. Kavya searches for science experiments, watches several astronomy videos to the end and skips some unrelated videos quickly.

The recommendation system receives signals from these actions. It may compare the topics, titles and viewing patterns of many videos and users. It then ranks content that appears more likely to interest Kavya. As she continues using the platform, the recommendations may become more personalised.

This convenience has limitations. Watching one dramatic video does not necessarily mean that Kavya wants many similar videos. Repeated recommendations can also create a narrow **filter bubble**, in which a user sees fewer viewpoints or subjects. Recommendation systems are designed to predict engagement, but an engaging item is not automatically accurate, healthy or educational.

Analyse the Case

1.What types of data or signals might the recommendation system use?
2.Why are recommendations less personalised for a new user?
3.What pattern might the AI infer from Kavya's viewing behaviour?
4.Why is a recommendation a prediction rather than a fact?
5.What is a filter bubble?
6.Suggest two actions Kavya can take to remain in control of what she watches.
7.Which information should the platform explain to users to support transparency?

Activities

Activity 1: Train a Human Classifier

Collect 20 simple pictures of leaves. Choose two categories, such as "smooth edge" and "toothed edge."

1. Use 14 pictures as training examples.
2. Ask one student—the "model"—to study the labelled examples and develop a classification rule.
3. Test the student using the remaining six unseen pictures.
4. Record correct and incorrect classifications.
5. Discuss which examples were difficult and how more varied data might help.

Python Code for the activity.

1.Project structure

```
.eaf-classifier/
│
├── leaf_classifier.py
│
└── leaf_data/
    ├── train/
    │   ├── smooth_edge/
    │   │   ├── leaf01.jpg
    │   │   └── ...
    │   └── toothed_edge/
    │       ├── leaf08.jpg
    │       └── ...
    │
    └── test/
        ├── smooth_edge/
        │   ├── leaf15.jpg
        │   └── ...
        └── toothed_edge/
            ├── leaf18.jpg
            └── ...
```

2.Prerequisites

1. Install Python 3.10 or later.
2. Open Terminal or Command Prompt in the project folder and run:
   a. For some Windows installations, use:
   b. py -m pip install numpy opencv-python scikit-learn

3. Preparing the pictures

For best results:

- Photograph one leaf at a time.
- Use a plain white or contrasting background.
- Keep the complete leaf visible.
- Avoid shadows and other objects.
- Include leaves of different shapes and sizes.
- Use .jpg, .jpeg, .png, or .bmp files.

4. Run the code

From the project folder, run:

python leaf_classifier.py

```
"""Activity 2: Train a Leaf Classifier.

Place 14 labelled training images and 6 labelled test images in the folders
described in the accompanying instructions, then run this file.
"""
from _future_ import annotations

import argparse
import csv
from pathlib import Path

import cv2
import numpy as np
from sklearn.metrics import accuracy_score, confusion_matrix
from sklearn.pipeline import make_pipeline
from sklearn.preprocessing import StandardScaler
from sklearn.tree import DecisionTreeClassifier, export_text

CLASSES = ("smooth_edge", "toothed_edge")
VALID_EXTENSIONS = {".jpg", ".jpeg", ".png", ".bmp"}

def image_files(folder: Path) -> list[Path]:
    """Return supported image files in predictable order."""
    if not folder.exists():
        return []
    return sorted(
        path for path in folder.iterdir()
        if path.is_file() and path.suffix.lower() in VALID_EXTENSIONS
    )

def largest_leaf_contour(image_path: Path) -> np.ndarray:
    """Separate a leaf from a plain background and return its largest contour."""
    image = cv2.imread(str(image_path))
    if image is None:
        raise ValueError("The image could not be opened.")

    gray = cv2.cvtColor(image, cv2.COLOR_BGR2GRAY)
    gray = cv2.GaussianBlur(gray, (5, 5), 0)
    # Otsu selects a threshold automatically. Try both polarities because a
    # leaf may be darker or lighter than its background.
    _, normal = cv2.threshold(gray, 0, 255, cv2.THRESH_BINARY + cv2.THRESH_OTSU)
    inverted = cv2.bitwise_not(normal)
 best = None
    image_area = gray.shape[0] * gray.shape[1]
    for mask in (normal, inverted):
        contours, _ = cv2.findContours(
            mask, cv2.RETR_EXTERNAL, cv2.CHAIN_APPROX_NONE
        )
```

```
for contour in contours:
area = cv2.contourArea(contour)
# Ignore tiny marks and a contour that is probably the whole frame.
if 0.02 * image_area < area < 0.95 * image_area:
if best is None or area > cv2.contourArea(best):
best = contour
if best is None:
raise ValueError(
"No clear leaf shape was found. Use one leaf on a plain, contrasting background."
)
return best


def extract_features(image_path: Path) -> list[float]:
"""Measure properties that can help distinguish smooth and toothed edges."""
contour = largest_leaf_contour(image_path)
area = max(cv2.contourArea(contour), 1.0)
perimeter = max(cv2.arcLength(contour, True), 1.0)
hull = cv2.convexHull(contour)
hull_area = max(cv2.contourArea(hull), 1.0)
hull_perimeter = max(cv2.arcLength(hull, True), 1.0)
x, y, width, height = cv2.boundingRect(contour)

# Toothed leaves often have a longer, less convex and less smooth outline.
circularity = 4.0 * np.pi * area / (perimeter * perimeter)
solidity = area / hull_area
edge_roughness = perimeter / hull_perimeter
extent = area / max(float(width * height), 1.0)
aspect_ratio = max(width, height) / max(float(min(width, height)), 1.0)

hu = cv2.HuMoments(cv2.moments(contour)).flatten()
hu = -np.sign(hu) * np.log10(np.abs(hu) + 1e-12)

return [
circularity,
solidity,
edge_roughness,
extent,
aspect_ratio,
*hu.tolist(),
]
```