



UNIT 6

Introduction To Artificial Intelligence (AI), and Machine Learning (ML)

Student Learning Outcomes

By the end of this chapter, you will be able to:

- Define artificial intelligence (AI), and machine learning (ML).
- Differentiate between AI, and ML, understanding their unique characteristics and applications.
- Understand the concept of supervised learning and its application to teach systems to make predictions.
- Comprehend unsupervised learning, used for clustering and pattern recognition.
- Understand different algorithms and choose relevant algorithms based on specific business problems. This includes simple linear regression.
- Familiarity with key model performance metrics accuracy.
- Understand where AI and ML are applied in daily life, business, and education.
- Discuss bias, privacy, and responsibility when using AI.

Introduction

In the present digital era, technology is not just about machines, it is about smart systems that can think, learn, and help us in making better decisions. This chapter introduces the exciting fields of Artificial Intelligence (AI), and Machine Learning (ML). These are powerful technologies behind many commonly used things such as voice assistants like google assistant, ChatGPT, video recommendations on YouTube, and even predicting the weather. Understanding these fields helps you see how computers are becoming more intelligent and helpful in solving real-life problems.

In this chapter, Students will explore how data is used to train machines to recognize patterns, make predictions, and group information just like a human brain, but more quickly and accurately.



6.1 Overview of AI and ML

In the modern world, huge amounts of data are being collected every second (more than humans can handle alone). This is where AI and ML come in. These technologies learn from data and help machines think, decide, and improve (just like humans) but faster and more consistently.

What is AI?

AI means teaching machines to do tasks that usually require human intelligence such as understanding voice, recognizing images, or making decisions.

Definition: AI is the ability of a machine or computer program to perform tasks that normally require human thinking and ability.

Real Life Examples:

- Urdu speech-to-text apps used in educational institutes including schools
- Chatbots used by banks for customer support
- Traffic control systems in big cities such as one introduced by Safe City Authority in Lahore.

What is ML?

Machine Learning is a part of AI. It means training a machine to learn from data and make decisions on its own without being programmed every time.

Definition: ML is a method where computers learn from past data to predict or decide something.

Examples:

- Netflix or YouTube recommend videos based on watching history of users.
- A weather app uses past temperatures to predict weather for next day.

Differences between AI and ML

The Table 6.1 shows the differences between AI and ML.

Table 6.1: Differences between AI and ML

Feature	AI	ML
Definition	Making machines smart enough to mimic human behavior	Teaching machines to learn from data and improve



Main Goal	To perform smart tasks like humans	To make predictions or decisions from data
Needs Data?	Sometimes (uses data to act smartly)	Always (needs data to learn)
Human Like Behavior?	Yes (like speaking, seeing, or understanding)	No (just learns patterns)
Real Life Example	Chatbot on JazzCash that replies to customers	App that predicts electricity usage next week



TIDBITS

- Think of AI as the brain, and ML as the experience that teaches brain how to improve.
- AI is the broader concept of machines acting smart, while ML is when machines learn from data to make predictions or decisions.

6.2 Machine Learning Types

In the previous section, you learned that ML helps machines learn from data to make predictions or decisions. Just like students can learn with the help of a teacher or by exploring on their own, machines can also learn in two main ways, i.e., Supervised Learning and Unsupervised Learning.

Supervised Learning

In supervised learning, the machine is given data with correct answers (**labels**). It learns by comparing its own guesses with the correct answers and its learning improves over the time.

Definition: Supervised learning is when the machine is trained using labeled data, where both instance and correct output are provided.

Example 1: Suppose you are training a machine to predict whether a fruit is an apple or a mango. You show the machine many Apples and Mangoes and label it as that

- "This is an Apple"
- "This is a Mango"



The computer learns from these examples. Later, it can predict with high accuracy when shown a picture of apple and mangoes such pictures.

Example 2: A system can be trained based on marks and attendance of students to predict who might fail.

Unsupervised Learning

In unsupervised learning, the machine is given only data without any labels. It has to find patterns or group similar things (cluster) by itself.

Definition: Unsupervised learning is when the machine finds patterns or group or similar things (cluster) in unlabeled data without being told, what the correct answer is?

This helps the store in marketing products in a better way, even without knowing the exact customer types.

Example: A school wants to divide students into study groups. The system checks their interests and subjects and forms groups just based on similarities and without knowing which group is "best".

Supervised vs Unsupervised Learning

The comparison between both types of learning is provided in Table 6.3.

Table 6.3: Supervised vs Unsupervised Learning

Feature	Supervised Learning	Unsupervised Learning
Labeled Data?	Yes (machine is told the correct answers)	No (machine finds patterns on its own)
Main Task	Predict or classify	Group or cluster similar data
Example	Predict student results using marks and study hours	Group customers based on shopping behavior
Uses in School	Identify students who need help	Make automatic student groups for projects



ACTIVITY

Discuss and Match (Supervised or Unsupervised Learning)

Objective: Understand the difference between **Supervised** and **Unsupervised** Machine Learning by identifying the correct type for real-life situations.



Instructions:

- Read each situation below carefully (refer to Table 6.4)

Table 6.4: Match as Supervised (S) or Unsupervised (U) Learning

Sr. No	Situation Statement	S or U
1	Predicting if a student will pass or fail	
2	Grouping students based on subject preferences	
3	Identifying spam emails from inbox	
4	Dividing products into types based on features	

- In your notebook, write the correct ML type for each: whether each situation is an application of **supervised or unsupervised learning**.
- Write your answers in your notebook with a brief reason for each.
- Discuss your answers with a classmate and groups
- Explain your reasoning during class discussion.



TIDBIT

Think of **Supervised Learning** as learning with a teacher, and **Unsupervised Learning** as exploring on your own.



REMEMBER

1. **Supervised Learning** uses labeled data to **predict or classify**.
2. **Unsupervised Learning** uses unlabeled data to **find hidden patterns**.
3. Both are used in schools, businesses, apps, and government systems in Pakistan.
4. If the system is learning from examples with correct answers, it is supervised. If it is grouping without given answers, it is unsupervised.

6.3 Practical Applications of AI/ML

Now that you understand what AI and ML are, and how machines learn, now it is time to explore where these technologies are being used around us. We use these technologies everyday similarly.

Smart Assistants and Recommendations

AI powers many smart tools and apps that help make our lives easier.

Smart Assistants: Smart assistants like Google Assistant, Siri, use AI to understand voice commands, translate languages, set reminders, and answer questions.



Example: Say "Kal 8 bajay mujhe jagana" to Google Assistant in Urdu, it will set an alarm using speech recognition and natural language processing techniques, both powered by AI.

Recommendation systems: These systems suggest what you might like, based on your past activities.

Examples:

- YouTube recommends videos after you watch a few clips.
- Shopping apps suggests clothes based on your browsing history.
- Instagram shows videos similar to what you have watched before.

These apps use **ML** to analyze what you like and improve their suggestions over time.



TIDBIT

If the app makes suggestions, understands your voice, answers you, or learns your habits.



ACTIVITY

AI in Our Daily Life

Objective: Recognize how AI and ML are used in everyday tools and applications.

Instructions:

- List **3 Apps** you have used during previous week that may use AI. For example, YouTube, Google Maps, ChatGPT, Snapchat.
- With a class partner, discuss **what the AI does** in each app. Does it suggest videos? Does it respond to voice commands? Does it help you search faster?
- Write your examples and the AI Role in your notebook.

Business Use Cases Using Customer Analysis and Chatbots

AI help businesses understand customers, answer questions, and improve their services.

Customer Analysis: Companies collect data about customers to find:

- Which product is best to sell.

Example: A grocery delivery app uses ML to suggest items you often buy, and predicts which items are needed in which city.

Chatbots and Automated Help: Many Pakistani banks, telecoms, and e-commerce sites now use chatbots which are smart systems that can answer



questions without needing a human agent.

These chatbots use AI to **understand typed questions** and respond automatically.



AI helps businesses **work faster**, **save money**, and **keep customers happy** using the power of data.

Bias, Privacy, and Responsibility

While AI is powerful, it must be used **responsibly**.

Bias in AI: Sometimes, machines **learn unfair patterns** from the data they are trained on. This is called **bias**. **For example**, if an AI system is trained only on English or male voices, it may not work well with Urdu speakers or female voices.



REMEMBER

Always use diverse and balanced data to reduce bias.

Privacy and Data Use: AI systems often use **personal data** (like name, location, or choices) which raises privacy concerns. For example, if an app tracks your movements or listens to your microphone without permission, it can be misused.

Best Practice:

- Ask for user permission
- Protect data using encryption
- Do not collect more than needed

Responsibility of Humans: Humans must use AI wisely as it is just a tool. If an AI powered system makes a wrong decision, the responsibility lies with the humans who created or used it. For example, if a facial recognition system wrongly accuses system should be fixed and detect.



REMEMBER

- AI is used in smart assistants, shopping apps, and chatbots.
- Businesses use them for customer service and sales analysis.
- We must always consider ethics: fairness, privacy, and responsibility.



6.4 Supervised and Unsupervised Algorithms

As stated in the previous Chapter, Data science uses various algorithms to find patterns, make predictions, or analyze large data. Choosing the right algorithm depends on the type of data and the kind of problem you are trying to solve. Some common data science algorithms are:

- **Linear Regression:** Predicts future values based on a straight-line relationship (e.g., predicting student scores from study hours).
- **Decision Trees:** Uses a tree-like model to make decisions based on yes/no questions (e.g., predicting if a customer will buy or not).
- **K-Means (Clustering):** Groups similar items together (e.g., grouping students with similar interests).

These algorithms help organizations solve real-world problems such as predicting exam results, grouping customers for better marketing and identifying patterns in fraud or health data.



TIDBIT

Think of an algorithm like a recipe in the kitchen. It tells the computer what ingredients (data) to use and what steps to follow to get the final dish (prediction or group).

Linear Regression

Simple Linear Regression algorithm helps us predict a value based on the relationship between two things. For example, a teacher wants to predict how well students will perform in a test based on the number of hours they studied. Here, study hours is taken as Input (X) while test score is used as Output (Y). Linear regression finds a line that best fits the data and helps predict marks if students studied for, say, 6 hours. Predicting sales, temperature, or student performance based on past data are among the few use cases of linear regression.

Scenario: A teacher wants to predict student test scores based on their study hours using data provided in Table 6.5. Linear Regression draws a best-fit straight line to predict future values as shown in Figure 6.1. In this graph, blue dots are actual student scores while blue line shows the predicted trend. If a student studies for 6.5 hours, the model can predict their score using the line.



Table 6.5: Student Data

Study Hours	Test Scores
01	35
02	40
03	45
04	50
05	55
06	60
07	65
08	70
09	75
10	80




Figure 6.1: Linear Regression: Study Hours vs Test Score



- Simple linear regression uses just one independent variable (input) to predict one outcome. However, in real life, outcomes usually depend on many things.
- Multiple linear regression helps the computer understand more complex patterns and improve prediction accuracy.



6.5 Evaluating Model Performance

Once a ML model is trained, (like predicting who will buy a product or whether a message is spam), then we are required to measure the performance of the model. This helps in deciding if the model is accurate and reliable. In this section, we will explore four important performance metrics with simple definitions and examples.

Confusion Matrix

A **confusion matrix** is a simple table that shows how many predictions/classifications were correct or incorrect.

Table 6.6 shows a simple confusion matrix considering whether a student will pass or fail (classification problem).

Table 6.6: Confusion Matrix

	Predicted: pass	Predicted: fail
Actual: pass	True Positive (TP)	False Negative (FN)
Actual: fail	False Positive (FP)	Ture Negative (TN)

- **TP** = Model predicted pass, and the student was pass in actual as well.
- **TN** = Model predicted fail and the student failed in actual as well.
- **FP** = Model predicted pass, but the student failed in actual.
- **FN** = Model predicted fail and the student was pass in actual.

The confusion matrix helps in understand whether a model is making classification or prediction mistakes and helps in calculating the accuracy (an important measure) of the model.

Accuracy

Accuracy (A) tells us how many predictions were correct out of all predictions made.

It is calculated as follows:
$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$$

Let us assume that we have data for 15 students. In actual, 8 students were pass while 7 failed. The model produced the following results:

- TP = Correctly predicted pass = 7
- TN = Correctly predicted fail = 5



- FP = Wrongly predicted pass = 2
- FN = Wrongly predicted fail = 1

Using the above formulas and model results, we calculate accuracy as follows:

$$A = \frac{7 + 5}{7 + 5 + 2 + 1} = \frac{12}{15} = 0.80 \text{ or } 80\%$$

Summary

- AI enables machines to perform smart tasks. ML, a part of AI, allows machines to learn from data to make predictions or identify patterns without being explicitly programmed every time
- Supervised Learning: The machine is given data with correct answers or labels (e.g., predicting student grades or customer purchases)
- Unsupervised Learning: The machine is given data without correct answers or labels (e.g., grouping people by preferences or behavior)
- Linear Regression for predicting values based on relationship between two values
- Decision Trees for making decisions based on yes / no questions
- K-Means Clustering for finding patterns in data (overview)

EXERCISE

Multiple Choice Questions:

1. What is Artificial Intelligence (AI)?

- (a) Solving math problems by hand
- (b) Machines performing tasks like humans
- (c) Sending emails using a computer
- (d) Watching videos online

2. Which of the following is an example of AI?

- (a) A chatbot replying to your questions
- (b) Using a whiteboard
- (c) A TV remote
- (d) Typing on a keyboard



3. What is Machine Learning (ML)?

- (a) Teaching humans to use machines
- (b) Teaching machines to learn from data like human
- (c) Learning computer hardware basics
- (d) A machine without internet

4. Which app recommends videos using ML?

- (a) Google Maps
- (b) Microsoft Paint
- (c) YouTube
- (d) Windows Media Player

5. In supervised learning, the data is:

- (a) Unorganized
- (b) Without any answers
- (c) Labeled with correct answers
- (d) Only text-based

6. Which of the following is an example of unsupervised learning?

- (a) Predicting student marks
- (b) Identifying spam emails
- (c) Marking quizzes
- (d) Grouping students by interest without labels

7. Google Assistant is an example of:

- (a) A game app
- (b) A word processor
- (c) A smart assistant using AI
- (d) A calculator

8. AI is mostly used for:

- (a) Handwriting notes
- (b) Making phone calls
- (c) Charging phones
- (d) Smart decision making and automation

9. A chatbot on JazzCash is an example of:

- (a) Cloud computing
- (b) Data visualization
- (c) AI in business
- (d) Unsupervised learning

Short Answer Questions

1. What is the difference between AI and ML?
2. Give one real-life example of AI used in Pakistan.



3. How does supervised learning work?
4. Define confusion matrix.
5. How can businesses use AI?
6. Give one example each of a smart assistant and a recommendation system.
7. What kind of data is used in unsupervised learning?
8. Define linear regression.

Long Answer Questions

1. Explain supervised and unsupervised learning with suitable examples.
2. Explain how model performance is evaluated in machine learning using the following metrics:
 - a) Confusion Matrix
 - b) Accuracy
3. A model predicts if students will pass (1) or fail (0). Out of 20 students:
True Positives (TP) = 8
True Negatives (TN) = 6
False Positives (FP) = 4
False Negatives (FN) = 2

You are required to (a) fill the confusion matrix, and (b) calculate accuracy.

Answer Key for Multiple Choice Questions

1. b) - Machines performing tasks like humans
2. a) - A chatbot replying to your questions
3. b) - Teaching machines to learn from data
4. c) - YouTube
5. c) - Labeled with correct answers
6. d) - Grouping students by interest without labels
7. c) - A smart assistant using AI
8. c) - Interpreting and visualizing data
9. d) - Smart decision making and automation
10. c) - AI in business