

CSE 403: Machine Learning

Chapter 01: Introduction to Machine Learning

Md Atikuzzaman

Lecturer

Department of Computer Science & Engineering
atik@cse.green.edu.bd



Outline

- 1 Introduction
- 2 What is Machine Learning?
- 3 Traditional Programming vs. Machine Learning
- 4 Approaches to Machine Learning
- 5 Supervised Learning
- 6 Unsupervised Learning
- 7 Reinforcement Learning
- 8 References

Our Everyday Life

1. Content Entertainment

- **Netflix / YouTube:**
"Because you watched X, you might like Y." (Recommendation Systems).
- **Spotify / Apple Music:**
Curated "Discover Weekly" playlists based on your listening history.
- **Social Media Feeds:**
Facebook/Instagram ranking posts to show you what matters most.

2. Utilities Productivity

- **Email (Gmail/Outlook):**
Automatically filtering "Spam" vs. "Primary" inbox.
- **Voice Assistants:**
Siri, Alexa, and Google Assistant understanding your speech (Speech Recognition).
- **Navigation (Google Maps):**
Predicting ETA and traffic jams based on historical data.

Machine Learning is not just a futuristic concept; it is the engine behind the apps we use every single day.

What is Machine Learning?

Traditional Definition (Arthur Samuel, 1959)

"Field of study that gives computers the ability to learn without being explicitly programmed."

Formal Definition (Tom Mitchell, 1997)


A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .

Core Idea:

Experience (Data) $\xrightarrow{\text{Generalization}}$ Prediction

Traditional Programming vs. Machine Learning


Traditional Programming



Traditional Paradigm:

- We manually write the code (rules).
- Data is processed to produce answers.

Machine Learning



ML Paradigm:

- The system "learns" from data and answers.
- The output is the **rules** (the model).

Example: Celsius to Fahrenheit (Traditional)

1. The Concept

The Logic is Known:

In traditional programming, we (the humans) understand the relationship beforehand.

We explicitly program the mathematical formula:

$$F = C \times 1.8 + 32$$

The rules are hard-coded by the developer.

2. The Implementation

We write the rule:

```
def to_fahrenheit(c):  
    # The rule is explicit:  
    return c * 1.8 + 32  
# Pass input, get answer  
print(to_fahrenheit(10))  
# Output: 50
```

We pass input data (10), and the function returns the calculated answer (50).

Example: Celsius to Fahrenheit (Machine Learning)

2. Machine Learning Approach

The Logic is Unknown (initially):

Instead of writing the formula, we provide the computer with raw examples (Data Pairs).

The Training Data:

Input (C)	Output (F)
-40	-40
0	32
10	50
100	212

The Process:

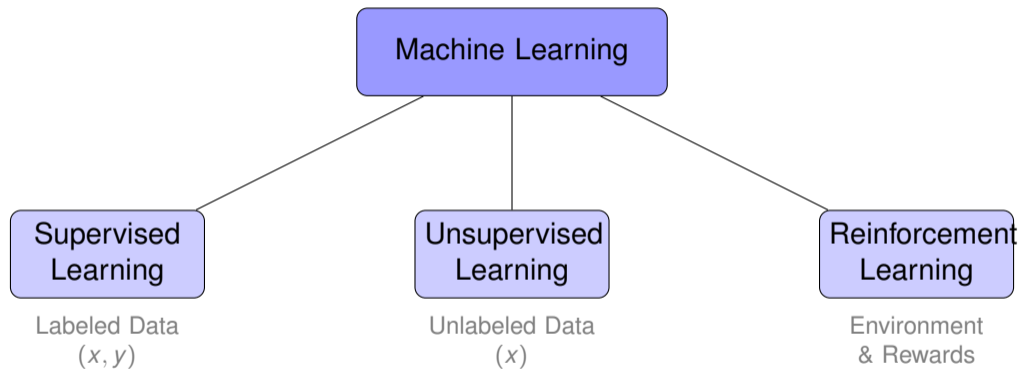
- 1 We assume a generic model structure:

$$\hat{y} = \mathbf{w} \cdot \mathbf{x} + \mathbf{b}$$

- 2 The algorithm iterates through the data to minimize error.
- 3 **Result:** It learns the parameters: $\mathbf{w} \approx 1.8$, $\mathbf{b} \approx 32$

Key Insight: In ML, the computer "figures out" the relationship (1.8 and 32) purely by analyzing the data, without being explicitly told the formula.

Three Approaches to Machine Learning



Supervised Machine Learning

Definition: Learning a mapping function f from input variables x to output variables y based on example input-output pairs.

- **Goal:** Approximate $y = f(x)$ well enough to predict y for new x .
- **Dataset:** $D = \{(x^{(i)}, y^{(i)})\}_{i=1}^m$

Regression

Predicting continuous values.

- *Ex:* House prices, Temperature.
- $y \in \mathbb{R}$

Classification

Predicting discrete labels.

- *Ex:* Spam detection, Diagnostics.
- $y \in \{0, 1, \dots, K\}$

Applications of Supervised Learning

1. Classification (Discrete Output)

Predicting a category or class label.

- **Image Recognition:** Face detection, identifying objects in photos (e.g., Google Photos).
- **Spam Filtering:** Classifying emails as "Spam" or "Not Spam" (e.g., Gmail).
- **Medical Diagnosis:** Predicting if a tumor is malignant or benign based on scans.
- **Fraud Detection:** Flagging credit card transactions as fraudulent or legitimate.

2. Regression (Continuous Output)

Predicting a numerical value.

- **Real Estate Prediction:** Estimating house prices based on square footage, location, and rooms.
- **Stock Market Forecasting:** Predicting future stock prices based on historical trends.
- **Weather Forecasting:** Predicting temperature, humidity, or rainfall amount.
- **Demand Forecasting:** Estimating product sales for retail inventory planning.

Supervised Learning is the most widely used form of ML in industry today.

Unsupervised Learning

Definition: Inferred function to describe hidden structure from "unlabeled" data.

- **Data:** $D = \{x^{(i)}\}_{i=1}^m$ (No labels y).
- **Goal:** Find patterns, groupings, or compressed representations.

Main Types:

- 1 **Clustering:** Grouping similar data points together (e.g., k-Means).

$$\min_{c, \mu} \sum_{i=1}^m \|x^{(i)} - \mu_{c^{(i)}}\|^2$$

- 2 **Dimensionality Reduction:** Reducing the number of random variables under consideration (e.g., PCA).
- 3 **Association:** Discovery of rules (e.g., Market Basket Analysis).

Applications of Unsupervised Learning

1. Clustering (Grouping)

Grouping similar data points together.

- **Customer Segmentation:** Grouping customers by purchasing behavior for targeted marketing (e.g., "Budget Shoppers" vs. "High Spenders").
- **Anomaly Detection:** identifying unusual patterns in network traffic (Cybersecurity) or manufacturing (Fault detection).
- **Social Network Analysis:** Detecting communities or circles of friends.

2. Association & Dim. Reduction

Finding rules and simplifying data.

- **Market Basket Analysis:** Discovering items frequently bought together (e.g., "People who buy Bread also buy Butter").
- **Recommendation Systems:** Suggesting movies/products based on user similarity (e.g., Netflix, Amazon).
- **Feature Compression:** Reducing file size or complexity without losing information.

Unsupervised Learning excels at discovering hidden structures in unlabeled data.

Reinforcement Learning (RL)

Definition

Learning **what to do how** to map situations to actions to maximize a numerical **reward signal**.

Key Components:

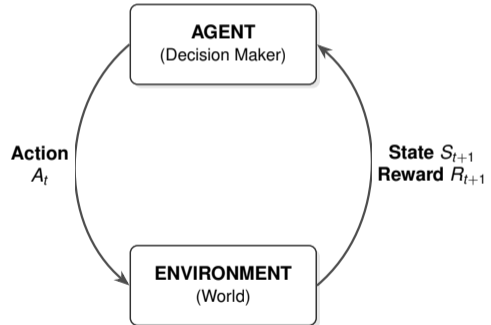
Agent: The learner & decision maker.

Environment: Everything the agent interacts with.

Action (A_t): What the agent chooses to do.

State (S_t): The current situation/context.

Reward (R_t): Immediate feedback score (Good/Bad).



Goal: Maximize the total cumulative reward over time.

Applications of Reinforcement Learning

1. Robotics & Control

Interacting with the physical world.

- **Autonomous Vehicles:** Learning to navigate traffic, merge lanes, and park without crashing.
- **Industrial Robotics:** Arms learning to pick and place objects of varying shapes and sizes.
- **Data Center Cooling:** Google's DeepMind reduced cooling costs by 40% via RL.

2. Games & Strategy

Beating humans in complex environments.

- **Game Playing (AlphaGo):** Mastering Go, Chess, and StarCraft II by playing millions of games against itself.
- **Recommendation Systems:** Optimizing long-term user engagement (e.g., YouTube next video suggestion).
- **Personalized Education:** Tutors that adapt lesson difficulty based on student performance.

RL learns complex behaviors through trial and error to maximize rewards.

References

- 1 C. M. Bishop, *Pattern Recognition and Machine Learning*, 1st Ed., Springer, 2006.
- 2 R. O. Duda, P. E. Hart, D. G. Stork, *Pattern Classification*, 2nd Ed., Wiley-Interscience, 2000.
- 3 I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, 1st Ed., The MIT Press, 2016.
- 4 A. Zhang, Z. C. Lipton, M. Li, A. J. Smola, *Dive into Deep Learning*, 1st Ed., Cambridge University Press, 2023.