

Introduction to Machine Learning

Introduction to Machine Learning

Key concepts, terminology, and foundational ideas

1. What Is Machine Learning?

Definition	Core Idea
Machine Learning (ML) is a sub-field of artificial intelligence that enables computers to learn patterns from data and make decisions or predictions without being explicitly programmed for each task.	Instead of hard-coding rules, an algorithm infers a function that maps inputs X to outputs Y from examples.

2. Types of Learning Paradigms

Paradigm	Goal	Typical Tasks	Example Algorithms
Supervised Learning	Learn a mapping from inputs to known labels.	Classification, regression	Linear regression, Decision trees, SVM, Neural networks
Unsupervised Learning	Discover hidden structure in unlabeled data.	Clustering, dimensionality reduction, density estimation	K-means, Hierarchical clustering, PCA, Autoencoders
Semi-Supervised Learning	Combine a small labeled set with a large unlabeled set.	Improved classification when labeling is expensive	Self-training, Graph-based methods
Self-Supervised Learning	Create surrogate labels from the data itself.	Representation learning for vision, NLP	Contrastive learning, Masked language modeling
Reinforcement Learning (RL)	Learn a policy that maximizes cumulative reward through interaction.	Game playing, robotics, recommendation	Q-learning, Policy gradients, Deep Q-Network (DQN)

Introduction to Machine Learning

Paradigm	Goal	Typical Tasks	Example Algorithms
Transfer / Multi-Task Learning	Leverage knowledge from related tasks to improve performance.	Domain adaptation, few-shot learning	Fine-tuning pretrained models, Multi-head networks

3. Core Components of an ML System

1. Data

- Raw observations (images, text, sensor readings, etc.)
- **Features** (raw or engineered) that capture informative aspects of the data.

2. Model

- A parameterized function f (e.g., linear model, neural network).
- Parameters θ are learned from data.

3. Loss / Objective Function

- Quantifies the error between predictions and true targets.
- Common choices: Mean Squared Error (MSE), Cross-Entropy, Hinge loss.

4. Optimization Algorithm

- Procedure to adjust θ to minimize the loss.
- Gradient-based methods (SGD, Adam) dominate deep learning; others include L-BFGS, coordinate descent.

5. Evaluation Metrics

- **Classification:** Accuracy, Precision, Recall, F1-score, ROC-AUC.
- **Regression:** RMSE, MAE, R^2 .
- **Clustering:** Silhouette score, Adjusted Rand Index.

6. Deployment & Monitoring

- Model serving (REST API, edge device).
- Continuous monitoring for drift, latency, and fairness.

Introduction to Machine Learning

4. The Machine Learning Workflow

1?? Define the problem ? 2?? Collect & explore data ? 3??

Preprocess & engineer features

?

?

4?? Split data (train/val/test)

5?? Choose model & loss

?

?

6?? Train (optimize) ? 7?? Validate & tune hyperparameters

?

8?? Test on heldout data ? 9?? Deploy

?

? Monitor ? Update (retrain) as needed

Key Practices

- **Data hygiene