

Image Processing and Computer Vision

Introduction to Machine Learning for Computer Vision

Sanath Thilakarathna

BSc Eng (Hons), MSc (Reading)

CINEC Campus

March 5, 2026



What is Machine Learning?

Definition

Machine Learning is a field of artificial intelligence that enables computers to learn patterns from data and make predictions or decisions without being explicitly programmed.

- Traditional programs follow explicitly written rules.
- Machine learning systems learn patterns from examples.
- The system improves its performance as more data becomes available.

Example

Instead of writing rules to detect a cat in an image, a machine learning model can learn the visual patterns of cats from thousands of labeled images.

Traditional Programming vs Machine Learning

Traditional Programming	Machine Learning
Rules + Data $\rightarrow$ Output Explicit programming Human designed logic	Data + Output $\rightarrow$ Rules (Model) Learning from examples Automatically learned patterns

Important

In machine learning, the system automatically learns the decision rules from training data.

Why Machine Learning is Needed

- Many real-world problems are too complex to solve using fixed rules.
- Images, speech, and natural language contain complex patterns.
- Machine learning allows computers to recognize patterns and generalize from data.

Example

Detecting handwritten digits using fixed rules is extremely difficult, but a machine learning model can learn digit patterns from thousands of examples.

Types of Machine Learning

Machine learning algorithms are commonly categorized into three main types.

- **Supervised Learning**
- **Unsupervised Learning**
- **Reinforcement Learning**

Important

The classification is based on the type of data and feedback provided to the learning algorithm.

Supervised Learning

Definition

Supervised learning uses labeled training data where the correct output is known.

- The model learns a mapping between inputs and outputs.
- Training data contains both input features and corresponding labels.

Example

In an image classification task, the training dataset may contain images labeled as:

- Cat
- Dog
- Car

The model learns to classify new images based on these examples.

Pros and Cons of Supervised Learning

Advantages

- High accuracy when labeled data is available
- Clear performance metrics
- Well-established algorithms
- Suitable for classification and regression tasks

Disadvantages

- Requires large amounts of labeled data
- Labeling is time-consuming and expensive
- May not generalize well to new domains
- Sensitive to label quality

Supervised Learning Models: Statistical and Mathematical Approaches

Definition

Statistical and mathematical models learn decision boundaries using mathematical functions and probability theory.

- **Linear Regression:** Predicts continuous values using linear relationships
- **Logistic Regression:** Binary classification using sigmoid function
- **Support Vector Machines (SVM):** Finds optimal decision boundaries
- **Naive Bayes:** Probabilistic classifier based on Bayes' theorem

Example

Logistic regression can classify images as cat or dog by learning a mathematical boundary that separates the two classes.

Supervised Learning Models: Neural Networks

Definition

Neural networks are models inspired by biological neurons that learn hierarchical patterns through multiple layers of interconnected nodes.

- Composed of layers: input, hidden, and output layers
- Learn complex non-linear relationships in data
- Trained using backpropagation algorithm
- Particularly effective for image and pattern recognition

Example

Convolutional Neural Networks (CNNs) are specialized neural networks that automatically extract visual features from images for classification tasks.

Unsupervised Learning

Definition

Unsupervised learning uses data without labeled outputs and attempts to discover hidden patterns or structures in the data.

- The system identifies patterns automatically.
- Often used for clustering or dimensionality reduction.

Example

A clustering algorithm may group images into categories based on similarity, even when no labels are provided.

Pros and Cons of Unsupervised Learning

Advantages

- No need for labeled data
- Reduces manual labeling costs
- Discovers hidden patterns automatically
- Useful for exploratory data analysis

Disadvantages

- Results are harder to evaluate
- No clear performance metrics
- May produce unintuitive clusters
- Requires domain expertise to interpret

Unsupervised Learning Models: Statistical and Mathematical Approaches

Definition

Statistical models discover data structure using mathematical techniques like distance metrics and probability distributions.

- **K-Means Clustering:** Partitions data into K clusters based on similarity
- **Hierarchical Clustering:** Creates tree-like cluster structures
- **Gaussian Mixture Models:** Probabilistic clustering with soft assignments
- **Principal Component Analysis (PCA):** Reduces dimensionality while preserving variance

Example

K-Means can cluster images based on visual similarity without any labels.

Unsupervised Learning Models: Neural Networks

Definition

Unsupervised neural networks learn representations and patterns without labeled data.

- **Autoencoders:** Learn compressed data representations
- **Self-Organizing Maps (SOM):** Organize data into low-dimensional grids
- **Deep Belief Networks:** Learn hierarchical features layer by layer
- Effective for feature learning and anomaly detection

Example

Autoencoders can compress images into lower-dimensional representations and reconstruct them, discovering important visual features automatically.

Reinforcement Learning

Definition

Reinforcement learning is a learning paradigm where an agent learns by interacting with an environment and receiving rewards or penalties.

- The system learns through trial and error.
- Actions that lead to higher rewards are reinforced.

Example

A robot may learn how to navigate a room by trying different movements and receiving rewards when it reaches the target location.

Pros and Cons of Reinforcement Learning

Advantages

- Learns through interaction with environment
- No labeled data required
- Can discover novel strategies
- Suitable for sequential decision-making

Disadvantages

- Requires many interactions to learn
- Computationally expensive
- Reward design is challenging
- Can be unstable during training

Reinforcement Learning Models: Statistical and Mathematical Approaches

Definition

Statistical and mathematical models in reinforcement learning use probability and value functions to learn optimal decision policies.

- **Q-Learning**: Learns optimal action values through experience
- **SARSA**: On-policy temporal difference learning
- **Markov Decision Processes (MDP)**: Mathematical framework for sequential decisions
- **Dynamic Programming**: Solves optimization problems iteratively

Example

Q-Learning allows an agent to learn the best moves in a game by updating value estimates based on rewards received.

Reinforcement Learning Models: Neural Networks

Definition


Neural network approaches combine deep learning with reinforcement learning to handle complex environments and high-dimensional state spaces.

- **Deep Q-Networks (DQN):** Uses neural networks to approximate Q-values
- **Policy Gradient Methods:** Directly learn optimal policies using neural networks
- **Actor-Critic Methods:** Combine value and policy learning
- Effective for complex tasks like game playing and robotics

Example

Deep Q-Networks can learn to play complex games by using convolutional neural networks to process visual input.


Machine Learning Workflow



Important

Machine learning systems learn patterns during training through iterative optimization.

Model Deployment and Prediction



Important

After training and evaluation, the trained model is deployed to make predictions on new, unseen data.

Machine Learning in Computer Vision


Machine learning is widely used in image processing and computer vision to analyze and interpret visual data.

Common applications include:

- Image classification
- Object detection
- Object tracking
- Image segmentation
- Image generation and enhancement

Different machine learning models are designed to solve different computer vision tasks.

Machine Learning Pipeline for Computer Vision



Important

In computer vision systems, machine learning models learn visual patterns from large collections of images.

Major Computer Vision Tasks

Task	Purpose
Image Classification	Identify the main object or category in an image
Object Detection	Locate objects in an image and determine their class
Object Tracking	Follow objects across multiple frames in a video
Image Segmentation	Assign a label to every pixel in the image
Image Generation / Enhancement	Generate new images or improve image quality

Image Classification Models

Goal: Identify what object is present in an image.

Model	Type
K-Nearest Neighbors (KNN)	Classical Machine Learning
Support Vector Machine (SVM)	Classical Machine Learning
Logistic Regression	Statistical Model
LeNet	Convolutional Neural Network
AlexNet	Deep Learning Model
VGGNet	Deep Learning Model
ResNet	Deep Learning Model
Vision Transformer (ViT)	Transformer-based Model

Object Detection Models

Goal: Detect objects in an image and determine their locations.

Model	Type
Viola–Jones (Haar Cascade)	Classical Computer Vision
HOG + SVM	Classical Machine Learning
R-CNN	Deep Learning
Fast R-CNN	Deep Learning
Faster R-CNN	Deep Learning
YOLO (You Only Look Once)	Real-time Deep Learning Detector
SSD (Single Shot Detector)	Deep Learning Detector
RetinaNet	Deep Learning Detector

Object Tracking Methods

Goal: Follow an object as it moves across video frames.

Method	Type
Optical Flow (Lucas–Kanade)	Classical Computer Vision
MeanShift / CamShift	Classical Tracking Algorithm
Kalman Filter	Statistical Model
SORT (Simple Online Tracking)	Deep Learning Based Tracking
Deep SORT	Deep Learning Tracking System

Image Segmentation Models

Goal: Assign a label to every pixel in the image.

Model	Type
Watershed Algorithm	Classical Image Processing
U-Net	Deep Learning Segmentation Network
Mask R-CNN	Instance Segmentation Model
DeepLab	Semantic Segmentation Network
SegNet	Deep Learning Segmentation Model

Image Generation and Enhancement Models

Goal: Generate new images or improve image quality.

Model	Application
Autoencoders	Image reconstruction and denoising
Generative Adversarial Networks (GANs)	Image generation
Diffusion Models	High-quality image generation
Super-Resolution Networks	Image enhancement

- Machine learning enables computers to learn patterns from data.
- Learning algorithms improve performance as more data becomes available.
- Machine learning methods are classified into supervised, unsupervised, and reinforcement learning.
- Machine learning is widely used in image processing and computer vision applications.