

Introduction to Machine Learning

Andrea Passerini
passerini@disi.unitn.it

Machine Learning

What is Machine Learning?

A classical definition (T. 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 .

A working definition

Machine learning studies algorithms that improve their behaviour from data or interaction, rather than being explicitly programmed for every case.


Mitchell's classical formulation separates three ingredients:

experience data or interaction with the environment;

task what the system should do;

performance how success is measured.

Learning = improving performance from experience



A Learning Problem

Data

A training set contains examples

$$\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$$

or, in unsupervised learning, only inputs

$$\mathcal{D} = \{\mathbf{x}_i\}_{i=1}^n.$$

Model

A hypothesis or model

$$f_{\theta} : \mathcal{X} \rightarrow \mathcal{Y}$$

maps inputs to predictions, controlled by parameters θ .

Objective


Training chooses parameters that perform well on the training data:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(f_{\theta}(\mathbf{x}_i), y_i).$$

Goal

The goal is not to memorize the training set, but to perform well on future examples drawn from the same problem.


Machine learning is an iterative design process



Modern view

Machine learning is an iterative process: define the problem, collect data, choose a representation and model class, train, evaluate, and revise the design based on errors.

Generalization



Training versus test performance


Training error measures performance on examples used to fit the model.

Test error estimates performance on new, unseen examples.


Key question

Does the model learn a pattern that generalizes, or does it only memorize the training set?


Underfitting and Overfitting



Underfitting: the model is too simple to capture the signal.




Good fit: the model captures the main regularity.



Overfitting: the model also captures noise and accidents of the sample.


From Feature Engineering to Representation Learning

Classical pipeline



less manual
feature design

Representation learning



Classical versus deep learning

Classical systems often rely on human-designed features. Deep learning methods can learn internal representations directly from raw data.

Learning Settings

Supervised

(x_i, y_i)
learn $x \mapsto y$

classification
regression

Unsupervised

x_i only
discover structure

clustering
density, embeddings

Semi-supervised

few labels
many inputs

use unlabeled data
to improve learning

Reinforcement

states, actions
rewards

learn policies
by interaction

What changes across settings?

The available feedback changes: labels, no labels, few labels, or rewards from interaction. The learning algorithm must exploit the information that is available.

Setting

The learner observes input/output pairs

$$(\mathbf{x}_i, y_i) \in \mathcal{X} \times \mathcal{Y}$$

and learns a function $f : \mathcal{X} \rightarrow \mathcal{Y}$.

- The labels define the task: class, score, ranking, decision.
- A performance measure evaluates predictions on new examples.
- Examples: image classification, speech recognition, spam detection, medical prediction.

Unsupervised and Semi-supervised Learning

Unsupervised learning

Only inputs are observed:

$$\mathbf{x}_i \in \mathcal{X}.$$

The goal is to discover structure:
clusters, low-dimensional
representations, densities, anomalies.


Semi-supervised learning

- A small labelled set is combined with a larger unlabelled set.
- Unlabelled data can help:
 - learn better representations
 - enforce consistency between similar examples

Reinforcement Learning

Setting

An agent interacts with an environment through states, actions, and rewards.



- The objective is to learn a policy that maximizes long-term reward.
- The learner must balance exploration and exploitation.
- Examples: robotics, games, recommendation, sequential decisions.

Learning Tasks

Supervised tasks

Classification
discrete labels

Regression
real values

Ranking
ordered outputs

Clustering
groups

Dimensionality
reduction

Density and
novelty detection

Unsupervised tasks

Important distinction

A setting specifies what information is available; a task specifies what we want to predict, discover, or optimize.

- Reasoning in the presence of uncertainty.
- Evaluating how evidence about some variables changes beliefs about other variables.
- Estimating probabilities and relations between variables from observations.

Connection to machine learning

Many learning algorithms can be understood as estimating uncertain quantities from data and making decisions under uncertainty.

Choice of Learning Algorithms

Information available

- Full knowledge of data distributions:
 - Bayesian decision theory.
- Distributional form known, parameters unknown:
 - parameter estimation from training data.
- Training examples available, distributional form unknown:
 - discriminative and generative supervised methods.
- Labels unavailable:
 - unsupervised learning, representation learning, clustering, density estimation.

Machine learning in one slide

- Machine learning uses data or interaction to improve performance.
- The central goal is **generalization** to unseen examples.
- Overfitting is controlled by model choice, data, regularization, and evaluation.
- Modern ML increasingly learns useful **representations** instead of relying only on hand-crafted features.
- Different settings and tasks require different assumptions and algorithms.