

Introduction to Machine Learning

How learning algorithms turn data into predictions, discover structure, and make sequential decisions.

COURSE CONTENT

Objectives and optimization	Empirical risk, convexity, gradient methods, constrained optimization, and regularization.
Regression and classification	Linear and logistic regression, nearest neighbors, Bayes classifiers, LDA, and naive Bayes.
Kernels and Gaussian processes	Feature maps, kernel ridge regression, support vector machines, and Gaussian-process prediction.
Unsupervised learning	PCA and kernel PCA; k-means, spectral clustering, Gaussian mixtures, EM, and density estimation.
Model selection	Bias-variance trade-offs, validation, cross-validation, and choosing model complexity.
Trees and ensembles	Decision trees, bagging, random forests, AdaBoost, and gradient boosting.
Neural and reinforcement learning	Multilayer perceptrons; Markov decision processes, Bellman equations, Q-learning, policy gradients, and actor-critic methods.

LEARNING FOCUS

Connect each model's assumptions to its objective, optimization algorithm, and generalization behavior.

Mathematical derivations, worked examples, and algorithm pseudocode.