Machine Learning

Core models and principles for learning from data.

A rigorous first course in machine learning, moving from mathematical foundations and probabilistic modeling to supervised learning, kernels, unsupervised learning, neural networks, and reinforcement learning. The course emphasizes how objectives are derived, how algorithms are implemented, and how modeling assumptions shape empirical behavior.

TOPIC MAP

Mathematical Foundations

- Linear algebra, probability, and optimization.
- Loss functions, empirical risk, regularization, and validation.
- Maximum likelihood, Bayesian decision rules, and risk minimization.
- Model assumptions and their effect on prediction behavior.

Supervised Learning

- Bayes classifiers, LDA, Naive Bayes.
- Logistic regression and linear regression.
- Classification versus regression objectives.
- Empirical risk minimization and model selection.

Kernel and Nonparametric Methods

- Kernel methods and support vector machines.
- Margins, feature maps, and regularized optimization.
- Gaussian processes.
- Bias–variance trade-offs beyond linear models.

Unsupervised Learning

- PCA and dimensionality reduction.
- Clustering, density estimation, and mixture-model ideas.
- Spectral methods and model-selection criteria.
- Interpreting structure without labeled responses.

Later Topics

- Neural networks, activations, and representation learning.
- Reinforcement learning and sequential decisions.
- Decision trees, random forests, and ensemble methods.
- Learning theory and generalization as time permits.

Course Emphasis

- Derive the objective.
- Implement the algorithm.
- Compare baselines and tune hyperparameters.
- Evaluate the model carefully with held-out data.

Probability / Optimization / Linear Models / Kernels / Unsupervised Learning