



Core Courses Syllabi

ML701 - Machine Learning

Title	Machine Learning
Code	ML701
Loading	4 Credit-hours
Prerequisites	Basic Concepts in Calculus, Linear Algebra and Programming
Catalog Description	This course provides a comprehensive introduction to Machine Learning. It builds upon fundamental concepts in Mathematics, specifically probability and statistics, linear algebra, and calculus. Students will learn about supervised and unsupervised learning, various learning algorithms, and basics of learning theory, graphical models, and reinforcement learning.
Goal	This graduate level course aims to familiarize students with foundations of core machine learning algorithms. This course aims to instill in students a strong grasp of supervised and unsupervised as well as the variants of learning algorithms. In addition, this course aims to expose students to the basics of learning theory, graphical models, and reinforcement learning.
Content	This course covers the following major module: (I) supervised learning, (II) unsupervised learning, (III) learning theory (IV) graphical models and reinforcement learning
Recommended Textbooks	Tom M. Mitchell, <i>Machine Learning</i> , McGraw-Hill Science/Engineering/Math publishing, ISBN: 0070428077
Recommended References & Supplemental Material	The following textbooks may be useful: <ol style="list-style-type: none">1. C. Bishop, <i>Pattern Recognition and Machine Learning</i>, Berlin: Springer-Verlag, 2006, ISBN: 03873107382. K. Murphy, <i>Machine Learning: A Probabilistic Perspective</i>, MIT Press, 2012, ISBN-10: 0262018020



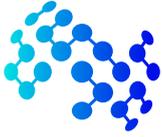
Teaching Week	Topics
1	Machine Learning Basics Lecture <ul style="list-style-type: none">Recap on probability theory and linear algebraIntroduction to optimizationOverview of libraries and programming recap Lab <ul style="list-style-type: none">Demonstration of programming environment and librariesPractice problems to get started with Python, Matlab, Jupyter Notebook, Numpy arrays etc.
2	Introduction to Supervised Learning Lecture <ul style="list-style-type: none">Regression vs classifier,Discriminative vs generative modelGeneral concepts: hypothesis, loss function (types of loss functions), cost functionGradient descent: stochastic vs batchLikelihoodNewton's algorithm Lab <ul style="list-style-type: none">Problem solving related to regression and fundamental optimization methodsImplementation of regressionData preprocessing lab – handling missing data, noisy labels etc.
3	Linear Models Lecture <ul style="list-style-type: none">Linear regression: LMS algorithm, locally weighted regressionClassification and logistic regression: sigmoid function, logistic regression, softmax regression,Classification and regression metricsGeneralized linear models: exponential family, assumption of GLMs Lab <ul style="list-style-type: none">Instructor-led walk-through of examples to demonstrate conceptsData visualization lab



Teaching Week	Topics
4	Support Vector Machines Lecture <ul style="list-style-type: none">• Optimal margin classifier• Lagrange duality• Support vectors, max-margin learning• Hinge loss, kernels, Lagrangian Lab <ul style="list-style-type: none">• Implement basic variant of SVM classifier from scratch without using any publicly available library
5	Generative Learning Lecture <ul style="list-style-type: none">• Gaussian discriminant analysis• Mixture models• Naïve Bayes model Lab <ul style="list-style-type: none">• Implementation of Naïve Bayes algorithm
6	Tree-based and Ensemble Models: Lecture <ul style="list-style-type: none">• Decision tree• Random forest• Boosting (adaptive boosting and gradient boosting) Lab <ul style="list-style-type: none">• Instructor-led demonstrations of random forest for classification
7	Learning Theory Lecture <ul style="list-style-type: none">• Union bound• Hoeffding inequality• Empirical risk• Probably approximately correct framework• Shattering• Upper bound theorem• Vapnik-Chervonenkis (VC) dimension and theorem Lab <ul style="list-style-type: none">• Mathematical Proofs and problem solving related to learning theory



Teaching Week	Topics
8	Introduction to Unsupervised Learning Lecture <ul style="list-style-type: none">• Introduction to unsupervised learning• Gaussian Mixture Model• Expectation-Maximization• k-means clustering Lab <ul style="list-style-type: none">• Instructor-led walkthrough of concepts taught in class
9	Clustering Lecture <ul style="list-style-type: none">• Hierarchical clustering• DB scan algorithm• Clustering assessment metrics (Silhouette coefficient, Calinski-Harabasz index)• Density estimation Lab <ul style="list-style-type: none">• Implement basic version of K-means clustering algorithm without help from any public library
10	Dimensionality Reduction Lecture <ul style="list-style-type: none">• Principal component analysis (Eigenvalues/vectors, Spectral theorem, PCA algorithm)• Independent component analysis (Bell and Sejnowski ICA algorithm)• Linear discriminant analysis• Factor analysis Lab <ul style="list-style-type: none">• Implement PCA and ICA from scratch, without using any library
11	Neural Networks: Lecture <ul style="list-style-type: none">• Basic architecture: multi-layer perceptron• Activation functions (Sigmoid, tanh, relu, lrelu etc.)• Losses (cross-entropy, l1/l2 loss, binary cross entropy, focal loss, hinge-loss etc.)• Back propagation, learning rate• Regularization (dropout, batch normalization) Lab <ul style="list-style-type: none">• Demonstration and familiarization with deep learning libraries



Teaching Week	Topics
12	Convolutional Neural Networks: Lecture <ul style="list-style-type: none">• Convolution operation• Normalization operations, novel layers• Training and inference Lab <ul style="list-style-type: none">• Train Neural Networks using standard libraries for basic problems
13	Recurrent Neural Networks: Lecture <ul style="list-style-type: none">• Recurrent networks• Backpropagation through time• Long short term memory• Gated recurrent units Lab <ul style="list-style-type: none">• Instructor-led walkthrough of concepts taught in class
14	Reinforcement Learning: Lecture <ul style="list-style-type: none">• Markov decision processes• Policy function• Value function• Bellman equation• Value iteration algorithm• Q-learning Lab <ul style="list-style-type: none">• Instructor-led demonstrations related to reinforcement learning
15	Graphical Models: Lecture <ul style="list-style-type: none">• Bayesian networks• Markov models Lab <ul style="list-style-type: none">• Instructor-led demonstrations related to graphical models