Syllabus

Fall 2023 CSI 873 / MATH 689

Computational Learning and Discovery

Schedule: T 7:20-10 pm, Exploratory Hall 4106, There is no class on October 10.
Instructor: Igor Griva, igriva@gmu.edu , (703) 993-4511.
Office hours: T 10 – 11 pm, Exploratory Hall, rm 4106.
Prerequisite: Permission of instructor. Students are expected to have familiarity with the basics of calculus, linear algebra, probability theory and statistics; understanding of basic programming principles and skills.
Text: Tom M. Mitchell, "Machine Learning," McGraw-Hill, 1997
Exams: There is one midterm exam: October 24 (points 0 - 100)
Final Exam: December 12 (points 0 - 100)
Final score: F = 0.3*(Midterm) + 0.4*(Homework / Projects) + 0.3*(Final Exam)

General description:
The course surveys algorithms of machine (computational) learning. The main goal of this class is to familiarize students with basic concepts and algorithms. Students who complete this course should be able to identify problems where computational learning algorithms can be useful and to apply these algorithms for finding the solution. We discuss the following topics: parametric/non-parametric learning, decision tree learning, neural networks, Bayesian learning, instance-based learning, Vapnik-Chernovenkis theory, support vector machines, and reinforcement learning. The class provides some necessary background introducing basic concepts from statistics, optimization, and information theory, relevant to computational learning.

Supplement recommended reading

Sergios Theodoridis, "Machine Learning: a Bayesian and Optimization Perspective", Academic Press, 2015.
Vladimir Vapnik, "The nature of statistical learning theory", Springer, 1999.
Trevor Hastie, Robert Tibshirani and Jerome Friedman, "The Elements of Statistical Learning: Data Mining, Inference, and Prediction," Second Edition, Springer Series in Statistics, 2009.
Foundations of Machine Learning, Second Edition, by Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalker, 2018 (click here to purchase a paper version with 30% off using the discount code MTSR20, or to rent a digital version for 4 or 12 months).

Topics week by week

Week 1. Survey of computational learning challenges.

Week 2. Concept learning. Version space and candidate elimination algorithm. Inductive bias.

Week 3. Data entropy. Information gain. Decision tree learning. Occam's razor.

Week 4. Artificial Neural Networks. Perceptron. Sigmoid units. Significance of the hidden nodes. Deep learning.

Week 5. Back propagation algorithm. Derivations of basic formulas. Strategies to amend overfitting. Recursive networks. Dynamically evolving networks.

Week 6. Hypothesis testing. Survey of basic material on probability related to computational learning. Confidence intervals. Training and true errors of computational learning algorithms.

Week 7. Bayesian learning. Bayes theorem. Maximum a posteriori hypothesis. Optimal Bayes classifier. Minimum description length principle.

Week 8. Naïve Bayes classifier. Conditional independence. Bayesian believe networks. Learning probabilities. Estimation – minimization algorithm.

Week 9. Computation learning theory. Probably approximately correct (PAC) setting. Vapnik-Chervonenkis (VC) dimension. Complexity bounds.

Week 10. Support Vector Machines. Theoretical justification. Practical considerations.

Week 11 Instance-based learning. Lazy algorithms. K-nearest neighbor (KNN) algorithm. Radial basis function networks.

Week 12. Evolutionary learning. Genetic algorithm. Evolutionary programming. Schemas.

Week 13. Brief survey of other topics of computational learning: unsupervised learning, reinforcement learning, analytical learning etc.

Week 14. Additional specific topics based on students' interest.