

**Fall 2007 CSI 709 / MATH 689**

## **Elements of Computational Learning**

**Schedule:** T 7:20-10:00 pm, Research I, Rm 202

**Instructor:** Igor Griva, igriva@gmu.edu , (703) 993-4511

**Office hours:** T 5:00 - 6:00 pm, Research I, Rm 344

**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 23 (points 0 - 100)

**Final Exam:** December 11 (points 0 - 100)

**Final score:**  $F = 0.3*(\text{Midterm}) + 0.4*(\text{Homework / Projects}) + 0.3*(\text{Final Exam})$

### **General description:**

The course focuses on elements of computational learning theory, which studies algorithms that can automatically improve their performance of some task with experience. The main goal of this class is to familiarize students with basic concepts and algorithms of computational learning. 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, bias/variance tradeoffs, 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. Some popular real world applications of computational learning algorithms are also discussed.

### **Supplement recommended reading**

Vladimir Vapnik, "The nature of statistical learning theory", Springer, 1999

Richard Duda, Peter Hart and David Stork, Pattern Classification, 2nd ed. John Wiley & Sons, 2001.

Christopher Bishop, Neural Networks for Pattern Recognition. Oxford University Press, 1996.

Brian Ripley, Pattern Recognition for Neural Networks. Cambridge University Press, 1996.

Richard Sutton and Andrew Barto, Reinforcement Learning: An introduction. MIT Press, 1998