

CSC 581 -- Special Topics in AI – Spring 2009
An Introduction to Machine Learning with Support Vector Machines

Description:

Support vector machines (SVMs) belong to a new class of machine learning algorithms with their origins firmly rooted in statistical learning theory. Due to the strong theoretical foundation these algorithms possess desirable properties such as the ability to learn from very small sample sets and a firm estimation of the generalization capacity of the learned model. These properties make this new class of learning algorithms extremely attractive to the practitioner who is frequently faced with "not enough data" and needs to understand "how good a constructed model" actually is. The fact that SVMs have surpassed the performance of artificial neural networks in many areas such as text categorization, speech recognition and bioinformatics bears witness to the power of this new class of learning algorithms.

This course is an introduction to machine learning and SVMs. We begin by framing the notion of machine learning and then develop basic concepts such as hyperplanes, features spaces and kernels necessary for the construction of SVMs. Once the theoretical groundwork has been laid we look at practical examples where this class of algorithms can be applied. Here we use machine learning as a *knowledge discovery tool*. We will use the statistical computing environment R for our experiments.

The goals of this graduate level course are for you,

1. To have a basic understanding of machine learning and knowledge discovery.
2. To be familiar with the mathematical framework of describing data and constructing models.
3. To be able to apply SVM packages in R to real-world problems.

Required Text:

Knowledge Discovery with Support Vector Machines, Hamel, Wiley, to appear.

Optional Texts:

Kernel Methods in Computational Biology, B. Schoelkopf, K. Tsuda, and J-P Vert (editors), MIT Press, 2004.

Kernel Methods for Pattern Analysis Christianini and Shawe-Taylor, Cambridge, 2004.

An Introduction to Support Vector Machines, Christianini and Shawe-Taylor, Cambridge, 2000.

The Nature of Statistical Learning, 2nd Edition, Vapnik, Springer, 1999.

Introductory Statistics with R, P. Dalgaard, Springer, 2002.

Website: <http://homepage.cs.uri.edu/faculty/hamel/courses/2009/spring2009/csc581/>

Prerequisites:

There are no formal requirements except that the student should possess a certain amount of mathematical maturity, exposure to linear algebra, probability theory, and functional analysis would be helpful. We will introduce the mathematical notions as we need them.

Grading:

Homework 50%

Midterm	25%
Final	25%

Scheduling: TuTh 9:30 -10:45 am Tyler 108,

Instructor: Prof. Lutz Hamel, Tyler Hall Rm 251, email: hamel@cs.uri.edu

Policies:

- Check the website (often)! I will try to keep the website as up-to-date as possible.
- Class attendance, promptness, participation, and adequate preparation for each class are expected. If you are absent, it is your responsibility to find out what you missed (e.g. handouts, announcements, assignments, new material, etc.)
- **Late assignments:** Late assignments will **not** be accepted.
- Make-up quizzes and exams will **not** be given without a valid excuse, such as illness. If you are unable to attend a scheduled examination due to valid reasons, please inform myself, or the department office in Tyler Hall, prior to the exam time. Under such circumstances, you are not to discuss the exam with any other class member until after a make-up exam has been completed.
- All work is to be the result of your own individual efforts unless explicitly stated otherwise. Plagiarism, unauthorized cooperation or any form of cheating will be brought to the attention of the Dean for disciplinary action. See the appropriate sections (8.27) of the University Manual.
- Software piracy will be dealt with exactly like stealing of university or departmental property. Any abuse of computer or software equipment will be subject to disciplinary action.

Tentative Course Outline:

- What is Knowledge Discovery?
- Knowledge Discovery Environments
- Describing Data Mathematically
- Linear Decision Surfaces and Functions
- Perceptron Learning
- Maximum Margin Classifiers
- Support Vector Machines
- Implementation
- Evaluating What has been Learned
- Elements of Statistical Learning Theory
- Multi-Class Classification
- Regression with Support Vector Machines
- Novelty Detection