Technical University of Crete
School of Electronic and Computer Engineering

Graduate Course on Machine Learning

TUC ECE COMP 604 | **Course Syllabus** | Fall Semester 2015

Lectures:	Tuesday, 2pm–4pm, 141.A14-2 (Intelligent Systems Lab) Friday, 3pm–5pm, 141.A14-2 (Intelligent Systems Lab)
Instructor:	Michail G. Lagoudakis
Contact:	141.A22, 28210-37244, `lagoudakis at intelligence tuc gr`
Info:	www.intelligence.tuc.gr/~lagoudakis
Web Site:	courses.ece.tuc.gr
Textbook:	Christopher M. Bishop *Pattern Recognition and Machine Learning*, Springer, 2006 http://research.microsoft.com/~cmbishop/PRML

Description

Data production in the digital era is constantly growing and so does the need for automated data processing and analysis methods. Machine learning studies methods for automatically extracting useful knowledge, patterns, and structures from data. Nowadays, machine learning is a broad field covering a wide range of research topics (supervised learning, unsupervised learning, reinforcement learning, learning theory) with applications to several disciplines (computer science, engineering, telecommunications, bioinformatics, robotics). Modern machine learning methods signal a departure from symbolic representations and methods and focus on numeric representations and statistical methods. The emergence of Statistical Machine Learning came as a confluence of areas such as statistics, pattern recognition, and signal processing. This course aims to provide a concise introduction to modern machine learning by covering core methods and approaches without explicit focus on any specific application area. The idea is that students with a clear understanding of such methods can subsequently apply them to domains of their interest and benefit from their results.

Participation

The course is open to all graduate students with basic background in mathematics (multivariate calculus, probability, and linear algebra), algorithms (design and analysis), and programming (coding in C/C++, Matlab, or similar languages). Senior undergraduate students may be allowed to register and join the class only if there is space, provided that they have completed successfully the undergraduate core course on "Probability and Random Signals".

Topics

1. Probability (Random Variables, Expectations, Distributions, Densities)
2. Parameter Estimation (Maximum A Posteriori, Maximum Likelihood, Expectation Maximization, Bayesian Estimation, Discriminative Training)
3. Supervised Learning (Bayesian Models, Linear Models, Discriminant Analysis)
4. Unsupervised Learning (Principal/Independent Component Analysis)
5. Kernel Methods (Gaussian Processes, Support Vector Machines, Relevance Vector Machines)
6. Reinforcement Learning (Sequential Decision Making, Value Function and Policy Learning)

Grading

Class participation (10%), Semester Project (40%), Final Written Examination (50%)

Active class participation will be taken into consideration and a final written examination will ensure sufficient breadth of study. To encourage deeper individual study on at least one topic, each student will have to complete and present a semester project involving application of some method or algorithm covered in class to data drawn from a domain of interest related to their area of research.