Technical University of Crete
School of Electrical and Computer Engineering

Graduate Course on Machine Learning

MLDS/INF 905 **Course Syllabus** Spring Semester 2024

Lectures: Tuesday, 17:00–19:00, 145.P58 (Science Building, 2nd floor)
Thursday, 17:00–19:00, 145.P42 (Science Building, 2nd floor)

Instructor: Prof. Michail G. Lagoudakis
Contact: 145.A35, 28210-37244, `lagoudakis at tuc gr`
Info: www.intelligence.tuc.gr/~lagoudakis

Web Site: www.eclass.tuc.gr/courses/MLDS109/

Textbooks: Kevin Patrick Murphy
Probabilistic Machine Learning: An Introduction, MIT Press, 2022
https://probml.github.io/pml-book/book1.html
Probabilistic Machine Learning: Advanced Topics, MIT Press, 2023
https://probml.github.io/pml-book/book2.html

Christopher M. Bishop
Pattern Recognition and Machine Learning, Springer, 2006
https://www.microsoft.com/en-us/research/people/cmbishop/prml-book/

A. Lindholm, N. Wahlström, F. Lindsten, and T. B. Schön
Machine Learning - A First Course for Engineers and Scientists, CUP, 2022
http://smlbook.org/

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 overview of 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, python, 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 course on "*Statistical Modeling and Pattern Recognition*".

Grading

Semester Project (30%), Midterm Exam (30%), Final Exam (40%)

Active class participation will be taken into consideration and the two (midterm and final) written examinations 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.