Machine Learning			
Credits: 4 **Semester:** 3 **Compulsory:** No			
Format	Lectures: 30 h	Tutorials: 18 h	Private study: 102 h
Lecturers: S. Rovetta, A. Verri			

Objectives:

The goal of the class is to present Artificial Neural Networks and other well known Machine Learning techniques (e. g. Gaussian Processes, Bayesian Learning, hidden Markov models, etc.) as systems for solving supervised and unsupervised learning problems, with a specific emphasis on Robotics applications. Such learning systems can be applied to pattern recognition, function approximation, time-series prediction and clustering problems. Some mention will be made to the use of ANNs as static systems for information coding, and dynamical systems for optimization and identification.

Contents:

The course will cover the following topics.

- Classification and identification for contemporary versions of Rosenblatt's Perceptron, Multi-Layer Perceptrons, Support Vector Machines and other Kernel Methods, and multi-layer perceptrons.
- Approximation properties of neural networks for multilayer perceptrons and for radial basis function networks.
- Insights on Machine Learning and Statistical Learning Theory: in particular, approximation quality and generalization problems.
- Learning algorithms like Back-propagation and Sequential Minimal Optimization to solve unconstrained and constrained optimization problems. Practical learning examples are discussed applied to Robotics.
- Neural ARMA models are derived as a generalization of ARMA models, and their properties analysed.

Practical Work:

Exercises on the application of architectures and learning algorithms to Robotics domains.

Abilities:

After completing this course, the students will gain the following capabilities.

- Understand the commonly used Machine Learning (ML) architectures and learning algorithms.
- Distinguish classes of problems to which MLs offer solutions, which are superior to other methods.
- Evaluate which ML architectures and learning algorithms are better suited to be applied to a specific problem.
- Design a ML architectureable to solve a practical scientific or industrial problem.
- Asses the performance of the designed ML architecture.

Assessment:

Continuous assessment (30%); final exam (70%).

Recommended texts:

- S.Haykin. Neural Networks and Learning Machines (3rd Ed.), Prentice Hall, 2008.
- M.H. Hassoun. Fundamentals of Artificial Neural Networks. The MIT Press, 1995.
- C.Bishop. Pattern Recognition and Machine Learning. Springer-Verlag, 2007.

Further readings:

Reference material will be provided by the lecturer.