Course guide
200900 - ML - Machine Learning

Last modified: 28/05/2026

Unit in charge: School of Mathematics and Statistics
Teaching unit: 748 - FIS - Department of Physics.
749 - MAT - Department of Mathematics.
715 - EIO - Department of Statistics and Operations Research.

Degree: MASTER'S DEGREE IN ADVANCED MATHEMATICS AND MATHEMATICAL ENGINEERING (Syllabus 2010). (Optional subject).
MASTER'S DEGREE IN STATISTICS AND OPERATIONS RESEARCH (Syllabus 2013). (Optional subject).

Academic year: 2026 **ECTS Credits:** 7.5 **Languages:** English

LECTURER

Coordinating lecturer: PEDRO FRANCISCO DELICADO USEROS

Others: Segon quadrimestre:
PEDRO FRANCISCO DELICADO USEROS - A
MANUEL MOLANO MAZÓN - A
ADRIÁN FERNANDO PONCE ÁLVAREZ - A
ADRIÁN FRANCISCO TAUSTE CAMPO - A

PRIOR SKILLS

The student should have knowledge of fundamental mathematical and statistical topics, such as linear algebra, calculus, probability distributions, optimization algorithms and basic multivariate statistical methods.

REQUIREMENTS

The student should have knowledge of basic machine learning concepts. These concepts can be acquired simultaneously, for example being enrolled in the "Statistical Learning" subject offered in the MESIO master.

DEGREE COMPETENCES TO WHICH THE SUBJECT CONTRIBUTES

Specific:
MAMME-CE2. MODELLING. Formulate, analyse and validate mathematical models of practical problems by using the appropriate mathematical tools.
MAMME-CE4. CRITICAL ASSESSMENT. Discuss the validity, scope and relevance of these solutions; present results and defend conclusions.

TEACHING METHODOLOGY

On-Site Learning: Learning will be organized into sessions that combine theory and practice. All sessions will be held in a standard classroom. Students should bring their own laptops to class. Lectures will consist of 75% expository instruction and 25% guided practical work. During the expository portion, theoretical aspects will be presented and discussed alongside practical examples using slides that will be provided to students in advance. The practical portion of the sessions will primarily use R and/or Python environments. Students should have intermediate knowledge of these environments and basic programming skills. Participation in class will be positively valued.

Off-site learning consists of studying and solving problems and real-world cases that students should submit throughout the course. Some of these exercises will require students to complete programming tasks in R or Python and prepare short reports using R Markdown (or a similar tool). In some cases, students will be required to give an oral presentation on their work in class.

LEARNING OBJECTIVES OF THE SUBJECT

Upon completion of the course, students should have acquired advanced competencies in general topics and techniques for unsupervised and supervised machine learning. In unsupervised learning, the course covers density estimation, clustering, and dimensionality reduction. Supervised learning topics include Bayesian methods, kernel-based learning, support vector machines, feedforward neural networks, recurrent neural networks, and reinforcement learning.

STUDY LOAD

Type	Hours	Percentage
Hours large group	60,0	32.00
Self study	127,5	68.00

Total learning time: 187.5 h

CONTENTS

Introduction to the course

Description:
Overview of the course objectives, structure, and main topics. Introduction to the importance of Machine Learning and its applications in various fields.

Full-or-part-time: 1h
Theory classes: 1h

Introduction to unsupervised learning

Description:
Introduction to key concepts in unsupervised learning. Focus on density estimation, clustering techniques, and dimensionality reduction.

Topics:
a. Introduction to the basic concepts of unsupervised learning.
b. Density estimation.
c. Clustering techniques.

Full-or-part-time: 8h
Theory classes: 8h

Dimensionality reduction

Description:
Advanced dimensionality reduction techniques for high-dimensional and infinite-dimensional datasets. Applications in visualization and exploratory analysis.

Topics:
a. Classical Dimensionality Reduction.
- Principal Component Analysis (PCA).
- Multidimensional Scaling (MDS).
b. Nonlinear Dimensionality Reduction.
- Principal curves.
- Local Multidimensional Scaling.
- ISOMAP.
- t-Stochastic Neighbor Embedding.

Full-or-part-time: 9h
Theory classes: 9h

Bayesian Methods

Description:
(Theory) Introduction to Bayesian thinking for machine learning. Concepts of Bayesian inference, including conjugate probability distributions, Bayesian linear regression, and Bayesian networks. Illustrative examples of probabilistic modeling and model comparison.
(Practice) Hands-on Bayesian inference using known software libraries for visualizing prior and posterior distributions and for Bayesian regression. Implementing the Bayesian regression solution and comparing it with the maximum likelihood estimator. Applications to real-world data.

Topics:
a. Bayesian Thinking and Inference.
b. Bayesian linear regression.
c. Bayesian Networks.

Full-or-part-time: 6h
Theory classes: 6h

Learning in Functional Spaces

Description:
(Theory) Exploration of reproducing kernel Hilbert spaces. The representer theorem and practical examples of kernel-based learning.
(Practice) Implementing Kernel Ridge Regression from scratch or using known libraries and exploring the effect of different kernels (linear, polynomial, Gaussian).

Topics:
a. Reproducing kernel Hilbert spaces (RKHS).
b. The representer theorem.
c. Kernel ridge regression as an example of kernel-based learning and comparison with Bayesian linear regression.

Full-or-part-time: 4h
Theory classes: 4h

Support Vector Machines

Description:
Exploration of support vector machines (SVM) as a flagship kernel method. Focus on classification, regression, and novelty detection.

Topics:
a. Introduction to SVMs: concept of separating hyperplanes and maximum margin.
b. Kernel trick: transforming data into higher-dimensional spaces.
c. Soft margin classification.

Full-or-part-time: 4h
Theory classes: 4h

Feedforward Neural Networks

Description:
Feedforward Neural Networks
Introduction to feedforward neural networks, including multi-layer Perceptrons and convolutional neural networks. Key aspects such as backpropagation, regularization, and examples are covered.

Topics:
a. Multi-layer Perceptron (MLP).
b. Backpropagation and regularization.
c. Convolutional Neural Networks (CNN).

Full-or-part-time: 6h
Theory classes: 6h

Recurrent Models

Description:
Recurrent Models
Exploration of sequential data modeling using recurrent neural networks (RNN), including advanced architectures such as LSTMs, GRUs, and transformers.

Topics:
a. Recurrent Neural Networks (RNN).
b. LSTM and GRU layers.
c. Interpretation of the networks.
d. Transformers as an alternative to RNN.

Full-or-part-time: 8h
Theory classes: 8h

Reinforcement Learning

Description:
Introduction to reinforcement learning (RL) and its foundational concepts, focusing on decision-making models and their mathematical underpinnings.

Topics:
a. Introduction to RL: basic principles and agent-environment interactions.
b. Dynamic Programming. Bellman equations: recursive relationships in value estimation.
c. Monte Carlo Methods.
d. Temporal-Difference Learning: Q-learning and SARSA algorithms.

Full-or-part-time: 6h
Theory classes: 6h

GRADING SYSTEM

The grading method will be based on two marks:

1) Participation in class: 15%.
2) Practical work completed throughout the course: 25%.
3) Final exam: 60%.

The practical work consists of studying and solving problems and real-world cases, which students should submit throughout the course. Some of these exercises will require students to complete programming tasks in R or Python and prepare short reports using R Markdown or a similar tool. In some cases, students will be required to give an oral presentation on their work in class. The practical work can be done in groups, and the format will be specified onsite. However, the exam must be completed individually.

EXAMINATION RULES.

The precise format for the exam will be specified with sufficient advance. It may include restrictions on the allowed knowledge sources, such as written notes, books, internet connection, etc.

BIBLIOGRAPHY

Basic:
- Bishop, Christopher M. Pattern recognition and machine learning. New York: Springer, cop. 2006. ISBN 9780387310732.
- Goodfellow, Ian; Bengio, Yoshua; Courville, Aaron. Deep learning [on line]. 2016 [Consultation: 19/06/2025]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=6287197. ISBN 9780262035613.
- Hastie, Trevor; Tibshirani, Robert; Friedman, Jerome. The Elements of statistical learning : data mining, inference, and prediction [on line]. 2nd ed. New York, NY: Springer New York, 2009 [Consultation: 19/06/2025]. Available on: https://link-springer-com.recursos.biblioteca.upc.edu/book/10.1007/978-0-387-84858-7. ISBN 9780387848587.
- Kung, S. Y. Kernel methods and machine learning [on line]. Cambridge University Press, 2014 [Consultation: 19/06/2025]. Available on: https://www-cambridge-org.recursos.biblioteca.upc.edu/core/books/kernel-methods-and-machine-learning/4B52092A98E1553A26EB5271D832D29E. ISBN 9781107024960.
- Murphy, Kevin P. Probabilistic machine learning: advanced topics [on line]. Cambridge: The MIT Press, 2023 [Consultation: 09/09/2025]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=30237492. ISBN 9780262048439.
- Schölkopf, Bernhard; Smola, Alexander J. Learning with Kernels : support vector machines, regularization, optimization, and beyond. Cambridge ; London: The MIT Press, 2002. ISBN 9780262194754.
- Sutton, Richard S.; Barto, Andrew G. Reinforcement learning : an introduction [on line]. Second edition. Cambridge, Massachusetts: The MIT Press, 2020 [Consultation: 19/06/2025]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=6260249. ISBN 9780262039246.
- Vapnik, Vladimir Naumovich. The Nature of statistical learning theory [on line]. 2nd ed. New York ; Barcelona [etc.]: Springer, 2000 [Consultation: 09/09/2025]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=3084784. ISBN 9780387987804.
- Aggarwal, Charu C.. Neural networks and deep learning : a textbook [on line]. Second edition. Springer, 2023 [Consultation: 01/06/2026]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=30620507. ISBN 9783031296413.
- Bosch Rué, Anna; Casas Roma, Jordi; Lozano Bagén, Toni. Deep Learning : principios y fundamentos [on line]. Universidad Oberta de Cataluña (UOC), 2019 [Consultation: 01/06/2026]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=7025971. ISBN 9788491806561.
- Boehmke, Bradley; Greenwell, Brandon. Hands-on machine learning with R [on line]. Chapman and Hall/CRC, 2019 [Consultation: 01/06/2026]. Available on: https://bradleyboehmke.github.io/HOML/. ISBN 9780367816377.

Complementary:
- Smola, Alexander J. Advances in large margin classifiers [on line]. Cambridge, Mass: MIT Press, 2000 [Consultation: 19/06/2025]. Available on: https://ebookcentral-proquest-com.recursos.biblioteca.upc.edu/lib/upcatalunya-ebooks/detail.action?pq-origsite=primo&docID=5966303. ISBN 9780262194488.