

Course guide

295921 - FAA - Fundamentals of Machine Learning

Last modified: 20/01/2026

Unit in charge: Barcelona East School of Engineering
Teaching unit: 749 - MAT - Department of Mathematics.

Degree: BACHELOR'S DEGREE IN BIOMEDICAL ENGINEERING (Syllabus 2009). (Optional subject).
BACHELOR'S DEGREE IN CHEMICAL ENGINEERING (Syllabus 2009). (Optional subject).
BACHELOR'S DEGREE IN ELECTRICAL ENGINEERING (Syllabus 2009). (Optional subject).
BACHELOR'S DEGREE IN ENERGY ENGINEERING (Syllabus 2009). (Optional subject).
BACHELOR'S DEGREE IN INDUSTRIAL ELECTRONICS AND AUTOMATIC CONTROL ENGINEERING (Syllabus 2009). (Optional subject).
BACHELOR'S DEGREE IN MECHANICAL ENGINEERING (Syllabus 2009). (Optional subject).
BACHELOR'S DEGREE IN MATERIALS ENGINEERING (Syllabus 2010). (Optional subject).

Academic year: 2025 **ECTS Credits:** 6.0 **Languages:** English

LECTURER

Coordinating lecturer: FRANCESC POZO MONTERO

Others: Pozo Montero, Francesc

PRIOR SKILLS

Students are expected to have a solid undergraduate-level background in linear algebra (vectors and matrices, matrix operations, norms, eigenvalues, and eigenvectors), calculus (functions of one and several variables, partial derivatives, gradients, and basic optimization concepts), and probability and statistics (random variables, common probability distributions, expectation, variance, and introductory estimation concepts). Students should also be able to program in at least one high-level language, such as Python, and be familiar with basic programming constructs, including variables, control flow, functions, and simple data structures. In addition, students are expected to be comfortable reading mathematical notation and to have basic algorithmic and computational thinking skills; no prior knowledge of machine learning is required, but students lacking part of this background are encouraged to review the relevant material during the first weeks of the course.

REQUIREMENTS

Students are strongly encouraged to bring a personal laptop to attend lectures and laboratory sessions, as no computers will be provided in the classroom; electrical power will be available for charging devices. In addition, students are highly recommended to register for the MIT Open Learning Library course MITx 6.036 Introduction to Machine Learning (<https://openlearninglibrary.mit.edu/courses/course-v1:MITx+6.036+1T2019/course/>), which provides exercises, homework assignments, and laboratory materials that can be used for self-assessment and practice. Registration and use of this online resource are recommended but not compulsory for successful completion of the course.

TEACHING METHODOLOGY

The course combines theoretical exposition with practical application to provide students with a solid understanding of the fundamental principles of machine learning. Core concepts will be introduced through lectures supported by visual materials and mathematical formulations, followed by illustrative examples and guided discussions. Practical sessions and assignments will emphasize hands-on experience with data, models, and algorithms, allowing students to apply theoretical ideas in realistic settings. Active student participation is encouraged through problem-solving activities and analysis of results, fostering critical thinking and a deeper understanding of model assumptions, limitations, and performance.

LEARNING OBJECTIVES OF THE SUBJECT

Upon successful completion of this course, students will be able to understand and explain the fundamental concepts and assumptions underlying machine learning methods; formulate supervised and unsupervised learning problems in mathematical terms; select appropriate models, loss functions, and evaluation criteria for a given task; apply core machine learning algorithms for classification, regression, clustering, and basic sequence and reinforcement learning problems; analyze model performance and generalization behavior; implement and evaluate machine learning models using appropriate computational tools; and critically interpret results, limitations, and potential sources of error in data-driven models.

STUDY LOAD

Type	Hours	Percentage
Hours large group	60,0	40.00
Self study	90,0	60.00

Total learning time: 150 h

CONTENTS

Introduction

Description:

This chapter introduces machine learning as the task of making predictions or decisions from data, with emphasis on generalization beyond the observed examples. It presents the motivation for machine learning, the problem of induction, and a unified framework to describe learning problems and solutions. The chapter outlines key concepts such as problem classes, assumptions, evaluation criteria, model types, model classes, and learning algorithms, establishing the conceptual foundation for the rest of the course.

Specific objectives:

By the end of this chapter, students will be able to:

- Explain the goals and scope of machine learning and its relationship with related fields.
- Identify the main classes of machine learning problems and their defining characteristics.
- Understand the role of assumptions and evaluation criteria in learning from data.
- Describe the fundamental components of a machine learning system and how they interact.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h



Linear classifiers

Description:

This chapter introduces linear classifiers as a fundamental class of models for supervised classification. It presents linear decision functions, geometric interpretations of classification in feature space, and the concept of separating data using linear boundaries. The chapter establishes the basis for understanding linear classification methods and their limitations, preparing the ground for learning algorithms introduced in subsequent chapters.

Specific objectives:

By the end of this chapter, students will be able to:

- Formulate linear classification problems using mathematical notation.
- Interpret linear classifiers geometrically in feature space.
- Understand the role of decision boundaries in classification tasks.
- Identify the strengths and limitations of linear classifiers.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h

The Perceptron

Description:

This chapter introduces the perceptron as a fundamental learning algorithm for linear classification. It presents the perceptron model, its update rule, and its interpretation as an iterative procedure for finding a separating hyperplane. The chapter discusses conditions for convergence, the relationship between the perceptron and linear classifiers, and highlights the limitations of the algorithm when data are not linearly separable.

Specific objectives:

By the end of this chapter, students will be able to:

- Describe the perceptron as a linear classification model.
- Understand and apply the perceptron learning algorithm.
- Interpret the perceptron update rule geometrically and algebraically.
- Analyze the convergence properties and limitations of the perceptron.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h



Feature representation

Description:

This chapter focuses on the role of feature representation in machine learning and how the choice of representation affects learning performance. It introduces feature transformations, basis expansions, and normalization techniques, emphasizing how raw data can be mapped into representations that make learning tasks easier or more effective. The chapter highlights the interaction between feature representation and model complexity.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the importance of feature representation in machine learning.
- Apply basic feature transformations and basis expansions.
- Analyze how representation choices affect model performance and generalization.
- Recognize the trade-offs between feature complexity and learning capability.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h

Margin Maximization

Description:

This chapter introduces margin maximization as a central principle in supervised classification. It formulates learning as an optimization problem in which the goal is to find a decision boundary that not only separates the data but does so with the largest possible margin. The chapter discusses the relationship between margins, robustness, and generalization, and introduces regularization as a mechanism to control model complexity.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the concept of margin in classification problems.
- Formulate margin-based learning objectives mathematically.
- Explain the role of regularization in controlling model complexity.
- Analyze how margin maximization influences generalization performance.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h



Gradient Descent

Description:

This chapter introduces gradient descent as a fundamental optimization method for training machine learning models. It presents the idea of minimizing an objective function through iterative parameter updates based on gradient information, and discusses practical aspects such as step size selection, convergence behavior, and computational efficiency. The chapter provides the optimization foundation used by many learning algorithms throughout the course.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand gradient descent as an optimization technique for learning models.
- Apply gradient-based methods to minimize learning objectives.
- Analyze the effects of step size and convergence properties.
- Recognize the role of optimization in training machine learning models.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h

Regression

Description:

This chapter addresses regression problems, where the objective is to predict continuous-valued outputs from input data. It presents regression as a supervised learning task, introduces common regression models and loss functions, and discusses how regression problems are formulated and evaluated. The chapter emphasizes interpretation of predictions and the relationship between model choice, loss functions, and generalization.

Specific objectives:

By the end of this chapter, students will be able to:

- Formulate regression problems within the supervised learning framework.
- Understand common regression loss functions and their implications.
- Apply regression models to continuous prediction tasks.
- Analyze regression model performance and generalization behavior.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h



Neural Networks

Description:

This chapter introduces neural networks as flexible parametric models for supervised learning. It presents neural networks as compositions of linear transformations and nonlinear activation functions, enabling the approximation of complex functions. The chapter discusses basic network architectures, the role of hidden layers, and training neural networks using gradient-based optimization methods.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand neural networks as function approximators for classification and regression.
- Describe the structure of feedforward neural network architectures.
- Explain the role of activation functions and hidden layers.
- Apply gradient-based methods to train neural network models.

Full-or-part-time: 15h

Theory classes: 3h

Laboratory classes: 3h

Self study : 9h

Convolutional Neural Networks

Description:

This chapter introduces convolutional neural networks (CNNs) as specialized neural network architectures designed to process structured data with spatial organization, such as images. It presents the principles of convolution, weight sharing, and pooling, and explains how these mechanisms exploit local structure and invariances in the data. The chapter highlights why CNNs are effective for high-dimensional inputs with spatial correlations.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the motivation behind convolutional neural networks.
- Describe convolution and pooling operations and their roles in CNN architectures.
- Explain how weight sharing and locality reduce model complexity.
- Analyze the advantages of CNNs for spatially structured data.

Full-or-part-time: 15h

Theory classes: 3h

Laboratory classes: 3h

Self study : 9h



Sequential models

Description:

This chapter introduces models for sequential data, where observations and predictions are ordered in time and may exhibit temporal dependencies. It presents state-based models as a framework for representing sequences, emphasizing how internal states summarize past information to influence future predictions. The chapter provides the conceptual basis for modeling time-dependent data in machine learning.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the challenges posed by sequential and time-dependent data.
- Describe state-based representations for sequence modeling.
- Formulate learning problems involving input and output sequences.
- Analyze how temporal dependencies affect prediction and learning.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h

Reinforcement learning

Description:

This chapter introduces reinforcement learning as a learning paradigm in which an agent learns through interaction with an environment by receiving reward signals. It presents the reinforcement learning framework, including states, actions, rewards, and policies, and explains how learning differs from supervised settings due to delayed feedback and sequential decision-making. The chapter establishes the foundations for understanding how agents learn to act optimally over time.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the reinforcement learning framework and its components.
- Formulate decision-making problems as reinforcement learning tasks.
- Distinguish reinforcement learning from supervised and unsupervised learning.
- Analyze the role of rewards and policies in sequential decision-making.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h



Recurrent Neural Networks

Description:

This chapter introduces recurrent neural networks (RNNs) as neural architectures designed to model sequential and temporal data. It explains how recurrence allows information from previous time steps to influence current predictions through hidden states. The chapter discusses the basic structure of RNNs, their training using gradient-based methods, and the challenges associated with learning long-term dependencies.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand recurrent neural networks as models for sequential data.
- Describe the role of hidden states and recurrence in RNNs.
- Explain how RNNs are trained using gradient-based optimization.
- Recognize the main challenges of training recurrent models, such as vanishing or exploding gradients.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h

Recommender systems

Description:

This chapter introduces recommender systems as machine learning methods designed to predict user preferences and suggest relevant items. It presents the basic principles behind recommendation tasks, including preference prediction and ranking, and discusses common approaches based on user-item interactions. The chapter highlights the challenges of sparsity, scalability, and evaluation in recommender systems.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the goals and structure of recommender systems.
- Formulate recommendation problems using user-item data.
- Identify common challenges such as data sparsity and scalability.
- Analyze the evaluation of recommendation performance.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h



Non-parametric methods

Description:

This chapter introduces non-parametric methods in machine learning, which make minimal assumptions about the functional form of the underlying model. It presents learning approaches where model complexity can grow with the amount of data, emphasizing flexibility and data-driven behavior. The chapter discusses the trade-offs between bias, variance, interpretability, and computational cost in non-parametric learning.

Specific objectives:

By the end of this chapter, students will be able to:

- Understand the defining characteristics of non-parametric learning methods.
- Distinguish non-parametric approaches from parametric models.
- Analyze the advantages and limitations of non-parametric methods.
- Evaluate the impact of data size and model flexibility on performance.

Full-or-part-time: 10h

Theory classes: 2h

Laboratory classes: 2h

Self study : 6h

GRADING SYSTEM

Student performance will be assessed through two partial exams and a final project. Each partial exam will account for 40% of the final grade and will evaluate the student's understanding of the theoretical foundations and problem-solving techniques covered in the corresponding part of the course. The remaining 20% of the final grade will correspond to a final project, which will involve the application of machine learning concepts to a practical problem and will include a written component and an oral presentation. The final grade will be computed as the weighted average of these assessment components, and detailed evaluation criteria will be provided in advance.

EXAMINATION RULES.

Partial exams and any other assessed tests will be conducted in accordance with the academic regulations of the institution. Students must attend examinations at the scheduled date and time and present valid identification if required. The use of unauthorized materials, devices, or any form of academic dishonesty will not be permitted and will be handled in accordance with the applicable disciplinary procedures. Students are responsible for following all instructions provided during the examination, and failure to comply may result in penalties or invalidation of the assessment. Any justified absence or special accommodation request must be communicated in advance and supported by appropriate documentation, in accordance with institutional guidelines.

BIBLIOGRAPHY

Basic:

- Mitchell, Tom M. Machine learning. New York [etc.]: The McGraw-Hill Companies, cop. 1997. ISBN 978-0070428072.
- Géron, Aurélien. Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow : concepts, tools, and techniques to build intelligent systems. 2nd edition. Sebastopol, CA: O'Reilly Media, Inc, 2019. ISBN 9781492032649.
- Bishop, Christopher M. Pattern recognition and machine learning. New York: Springer, cop. 2006. ISBN 9780387310732.
- Murphy, Kevin P. Machine learning : a probabilistic perspective. Cambridge, Mass: MIT Press, cop. 2012. ISBN 9780262018029.
- Alpaydin, Ethem. Introduction to machine learning. 4th edition. Cambridge, Massachusetts ; London: The MIT Press, [2020]. ISBN 9780262043793.

Complementary:

- Burkov, Andriy. The Hundred-page machine learning book. Leipzig: Andriy Burkov, 2019. ISBN 9781999579500.



RESOURCES

Hyperlink:

- <https://openlearninglibrary.mit.edu/courses/course-v1:MITx+6.036+1T2019/course/>. The MITx 6.036 Introduction to Machine Learning course on the MIT Open Learning Library is a self-paced, archived version of MIT's undergraduate machine learning class, designed to introduce core principles, algorithms, and applications of machine learning from the standpoint of modeling and prediction. It includes lecture notes, exercises, labs, and homework covering supervised learning, representation, over-fitting, generalization, and reinforcement learning, with practical examples and sequence learning applications. Users can browse content freely, though no certificate is awarded through the Open Learning Library version.