

COMPUTER ENGINEERING (LM55)

(Lecce - Università degli Studi)

Teaching HIGH PERFORMANCE COMPUTING

GenCod A003136

Owner professor Giovanni ALOISIO

Teaching in italian HIGH PERFORMANCE COMPUTING
Teaching HIGH PERFORMANCE COMPUTING
SSD code ING-INF/05

Reference course COMPUTER ENGINEERING

Course type Laurea Magistrale

Credits 9.0

Teaching hours Ore-Attività-frontale: 81.0

For enrolled in 2018/2019

Taught in 2019/2020

Course year 2

Language INGLESE

Curriculum PERCORSO COMUNE

Location Lecce

Semester Secondo-Semestre

Exam type Orale

Assessment Voto-Finale

Course timetable
<https://easyroom.unisalento.it/Orario>

BRIEF COURSE DESCRIPTION

The course provides a broad introduction to Machine Learning and Neural Networks Topics include: Linear regression, Linear regression with multiple variables, Regularization, Logistic regression, Neural Networks, Advice for Machine Learning, Machine Learning system design, Support Vector Machines. Hands-off are also organized to provide students the capacity to develop specific use cases, using the Python language and Jupiter Notebooks.

REQUIREMENTS

Students are expected to have the following background: Knowledge of basic computer science principles and skills, at a level sufficient to write a reasonably non-trivial computer program. Familiarity with the basic probability theory. Familiarity with the basic linear algebra. Good knowledge of the contents of first level courses on Informatics ("Fondamenti di Informatica" and "Calcolatori Elettronici") and of the courses on "Parallel Algorithms" - 2nd Year (1st semester) of the 2nd Level Degree in Computer Engineering.

COURSE AIMS

The course focuses on the basic of Machine Learning and Neural Networks. The theory behind Machine Learning and Neural Networks is presented and applications are developed using the Python language and Jupiter Notebooks.

TEACHING METHODOLOGY

Compulsory attendance is requested, since attendance at lectures and laboratory is mandatory because the course is based on the "learning by doing" approach on the advanced computing resources provided. The course will be held in the HPC Lab of the Department of Innovation Engineering. Lectures will be also given exploiting communication and collaboration platforms, like Microsoft Teams.

ASSESSMENT TYPE

The oral exam is aimed at verifying to what extent the student has gained knowledge and understanding of the selected topics of the course and is able to communicate about his understanding. Students, divided into small groups (max two students), will also get hands-on experience, developing small projects on specific topics of the course. The max final vote is expressed as 30/30 with the possibility to get the laude.

FULL SYLLABUS

Introduction to Machine Learning (2 hours).

Linear Regression: Model representation, Cost Function, Cost Function Intuition, Gradient Descent, Gradient Descent Intuition, Gradient Descent for Linear Regression (6 hours).

Linear Regression with multiple variables: Multiple features, Gradient Descent for Multiple variables, Gradient Descent in practice I: feature scaling, Gradient Descent in practice II: Learning rate, Features and polynomial regression, Normal Equation, Normal Equation and non-invertibility (6 hours).

Regularization: The problem of Overfitting, Cost Function, Regularized Linear Regression (6 hours).

Logistic Regression: Classification, Hypothesis Representation, Decision Boundary, Cost Function, Simplified cost function and gradient descent, Advanced Optimization, Multi-class classification: One-vs-all, Overfitting, Cost Function, Regularized Logistic Regression (6 hours).

Neural Networks-Representation: Non-linear hypothesis, Neurons and brains, Model Representation, Examples and intuitions, Multi-class classification (6 hours).

Neural Networks-Learning: Cost Function and Backpropagation algorithm, Backpropagation intuition. Implementation note: Unrolling parameters, Gradient checking, Random initialization, Putting it together (6 hours).

Advice for applying Machine Learning: Deciding what try next, Evaluating a hypothesis, Model selection and training/validation/test sets, Diagnosing bias vs variance, Regularization & bias/variance, Learning curves (6 hours).

Machine Learning system design: Prioritizing what to work on: Spam classification example, Error analysis, Error metrics for skewed classes, Trading off precision and recall, Data for machine learning (6 hours).

Support Vector Machines: Optimization objective, Large Margin Intuition, The mathematics behind large margin classification, Kernels, Using an SVM (6 hours).

Hands-on Machine Learning: Use cases developed using Python and Jupiter Notebook (25 hours).

REFERENCE TEXT BOOKS

- *Andrea Ng*, Machine Learning – Online book (<http://cnx.org/content/col11500/1.4/>) and in particular the Lectures notes (reported at pag. 223-320 of the online book).
- Teacher's Slides presented during the course.
- Jupiter Notebooks provided during the course.