

COMPUTER ENGINEERING (LM55)

(Lecce - Università degli Studi)

Insegnamento HIGH PERFORMANCE COMPUTING

GenCod A003136

Insegnamento HIGH PERFORMANCE COMPUTING

Insegnamento in inglese HIGH PERFORMANCE COMPUTING

Settore disciplinare ING-INF/05

Corso di studi di riferimento COMPUTER ENGINEERING

Tipo corso di studi Laurea Magistrale

Crediti 9.0

Ripartizione oraria Ore Attività frontale: 81.0

Per immatricolati nel 2019/2020

Erogato nel 2020/2021

Anno di corso 2

Lingua INGLESE

Percorso PERCORSO COMUNE

Docente Giovanni ALOISIO

Sede Lecce

Periodo Secondo Semestre

Tipo esame Orale

Valutazione Voto Finale

Orario dell'insegnamento

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

BREVE DESCRIZIONE DEL CORSO

The course discusses the cache-based processor architectures, with a special focus on the issues related to performance. Multicores solutions and those bases on simultaneous multithreading are considered. The general strategies for optimizing serial coded are also presented. Moreover, 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.

PREREQUISITI

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.

OBIETTIVI FORMATIVI

KNOWLEDGE AND UNDERSTANDING:

By the end of the course, students will gain knowledge of the key theoretical and practical notions to solve the main issues related to the use of high-performance computers.

In addition, students will have a basic theoretical knowledge of Machine Learning.

APPLYING KNOWLEDGE AND UNDERSTANDING:

By the end of the course, students will know how to use the acquired notions for the efficient use of a high-performance computer. In addition, students will be able to develop Machine Learning applications.

MAKING JUDGEMENTS

By the end of the course, students will know how to evaluate the different approaches to the optimization of a program running on a high-performance computer, being able to choose the best solutions for performance optimization. Students will also be able to critically assess the different Machine Learning approaches.

COMMUNICATION SKILLS:

By the end of the course, students will be able to use a clear language and an adequate scientific terminology to argue on the topics dealt with in the course.

LEARNING SKILLS:

Students will be able to classify, schematize, summarize and process the acquired knowledge. By the end of the course, students will have the appropriate skills to develop and widen their knowledge of high-performance computing, with particular regard to the use of reference documentation and other information available online. Students will be able to widen their knowledge about the best strategies for the use and performance optimization of a high-performance computer.

Students will also be able to further widen their knowledge in the field of artificial intelligence.

METODI DIDATTICI

The credits of this course are obtained by means of an oral exam, that will assess the overall learning results achieved by the student.

Students (both attending and non-attending) will be asked three questions, one of which is aimed at checking problem solving skills and the student's ability to use the theoretical and practical notions acquired.

The final score will consider:

- the level of theoretical/practical notions acquired (50%);
- the ability to use the theoretical/practical notions acquired (30%);
- the ability to make autonomous assessments (10%);
- the communication skills (10%).

Honors are awarded to students mastering the topics covered in the course.

MODALITA' D'ESAME

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.

APPELLI D'ESAME

Dates for the exams will be published on the official University web site.

ALTRE INFORMAZIONI UTILI

Lecture slides presented during the course will be provided and available on-line.

Introduction of the course and review of previous topics covered in the first level course on Computer Architectures.

General-purpose cache-based microprocessor architecture: Performance metrics and benchmarks. Moore's Law. Pipelining. Superscalarity. SIMD.

Memory hierarchies: Cache. Cache mapping. Prefetch.

Advanced Solutions: Multicore processors. Multithreaded processors. Vector processors.

Basic optimization techniques for serial code: Scalar profiling (Function and line-based runtime profiling, Hardware performance counters, Manual instrumentation). Common sense optimizations.

Data access optimization: Balance analysis and lightspeed estimates (Bandwidth-based performance modeling, The STREAM benchmarks).

Machine Learning: What we will learn on Machine Learning.

Supervised Learning Algorithms: - Linear Regression - Classification.

Hypothesis Function for Linear Regression: - Representation of the hypothesis function for regression problems.

Linear Regression with one variable - Cost Function: - Analysis of the hypothesis function and the Cost Function - Use of graph for their interpretation (3 hours).

Contour Plots: Use of the Contour plots to have a better intuition of the Cost Function (3 hours).

Minimization of the Cost Function-Gradient Descent Algorithm: Use of the Gradient Descent to minimize the Cost Function (3 hours).

Linear Regression with multiple features: Hypothesis function for Multivariate Linear regression (3 hours).

Gradient Descent for Multivariate Linear Regression: Cost Function and Gradient Descent Algorithm for Multivariate Linear Regression (3 hours).

Feature Scaling: Practical tricks for making Gradient Descent work well.

Learning Rate: Another practical tricks for making Gradient Descent work well.

Features and Polynomial Regression: How to use the machinery of Linear Regression to fit very complicated, even very non linear functions.

Normal Equation: - Normal Equation as a better way to solve for the optimal value of the Linear Regression parameters - Advantages and disadvantages of Gradient Descent and Normal Equation methods.

Normal Equation and Non-Invertibility: Cases of Non-Invertibility.

Logistic Regression: How do we develop a Classification Algorithm?

Logistic Regression - Hypothesis Representation: Mathematical formula defining the hypothesis for Logistic Regression and its probabilistic interpretation.

Decision Boundaries: - What Decision Boundaries mean? - Example for non-linear Decision Boundaries.

Logistic Regression-Cost Function: - Cost Function for Logistic Regression - Simplified Cost Function and Gradient Descent.

Multi-class Classification: One-versus-all Classification.

Overfitting and Regularization: - The problem of Overfitting - Regularization - Regularized Cost Function.

Neural Networks: Non Linear hypothesis, Neurons and the brain, Model representation, Multi-class classification, Cost Function & Backpropagation algorithm

Advice for Applying ML: Model selection & Training_Validation_Test sets

Machine learning system design_Error analysis, Trading off precision and recall

TESTI DI RIFERIMENTO

Georg Hager and Gerhard Wellein " Introduction to High Performance Computing for Scientists and Engineers ", CRC Press, © 2011 by Taylor and Francis Group, LLC - ISBN 978- 1-4398-1192-4.
Andrew Ng : "Video Course on Machine Learning"