

Big Data Processing and Analytics

Teachers

- Paolo Garza
 - paolo.garza@polito.it
 - 011-090-7022
- Lorenzo Vaiani
- Etibar Vazirov

Office hours

- Class-time (break, end of lesson)
- Or send an e-mail for an appointment

Weekly schedule

- Lectures (45 hours)
 - Monday - 14:30-17:30 – Room R3
- Practices (15 hours)
 - Friday - 14:30-16:00 – LAIB1 - Team 1 (A-L)
 - Friday - 16:00-17:30 – LAIB1 - Team 2 (M-Z)
 - There are **no lab activities** during the **first week**

Practices

- We will provide you a specific account on the BigData@Polito cluster
 - <https://jupyter.polito.it>
- Detailed information will be provided next week
 - You will receive an email from the administrator of the cluster with username and password

Topics

- Lectures
 - Introduction to Big data
 - Hadoop
 - Architecture
 - MapReduce programming paradigm
 - Spark
 - Architecture
 - Spark programs based on RDDs (Resilient Distributed Data sets) and Spark SQL (DataFrames and Datasets)

Topics

- Data mining and Machine learning libraries for Big Data
 - **MLlib** (Apache Spark's scalable machine learning library)
- Streaming data analysis
 - **Spark Streaming**

Topics

- Laboratory activities
 - Application development on Hadoop and Spark

Prerequisites / prior knowledge

- Basic programming skills
 - **Java language (MapReduce for Hadoop)**
 - **Python language (Spark)**
- Basic knowledge of traditional database concepts (recommended)
 - Relational data model
 - SQL language

Material

- Web page
 - https://dbdmg.polito.it/dbdmg_web/2024/big-data-processing-and-analytics-2025-26
 - Slides, exercises, lab activities, ..
- Video lectures/Virtual classrooms
 - On the Teaching portal
 - <https://didattica.polito.it>

Books and Readings

- Reference books:
 - Matei Zaharia, Bill Chambers. Spark: The Definitive Guide (Big Data Processing Made Simple). O'Reilly Media, 2018.
 - Advanced Analytics and Real-Time Data Processing in Apache Spark. Packt Publishing, 2018.
 - Matei Zaharia, Holden Karau, Andy Konwinski, Patrick Wendell. Learning Spark (Lightning-Fast Big Data Analytics). O'Reilly, 2015.
 - Tom White. Hadoop, The Definitive Guide. (Third edition). O'Reilly Media, 2015.
 - Donald Miner, Adam Shook . "MapReduce Design Patterns: Building Effective Algorithms and Analytics for Hadoop and Other Systems." O'Reilly, 2012

Exam rules

- Written exam
 - 2 programming exercises (max 27 points)
 - Design and develop programs based on the Hadoop MapReduce programming paradigm and/or Spark APIs
 - 2 questions/theoretical exercises (max 4 points)
 - Topics
 - Technological characteristics and architecture of Hadoop and Spark
 - HDFS
 - MapReduce programming paradigm
 - Spark RDDs, transformations and actions
 - Spark SQL
 - Spark Streaming
 - Spark MLlib

Exam rules

- On-site written exam on the **Exam platform with Lockdown browser** – **You must bring your own PC**
 - 1.5 hours
 - The exam is **open book**
 - Books, notes, and any other paper material are allowed
 - Electronic devices of any kind (PC, mobile phone, calculators, etc.) are not allowed, except the PC used for the exam itself
- Exam examples will be available on the web page of the course