

Erasmus Mundus Joint Master Degree

Big Data Management and Analytics



Detailed Course Description

Academic Year 2021-2022

Contents

| | |
|---|-----------|
| Université libre de Bruxelles (ULB) | 3 |
| Advanced Databases (ADB) | 3 |
| Database Systems Architecture (DBSA) | 4 |
| Data Warehouses (DW) | 5 |
| Business Process Management (BPM) | 7 |
| Data Mining (DM) | 8 |
| Universitat Politècnica de Catalunya (UPC) | 10 |
| Big Data Management (BDM) | 10 |
| Semantic Data Management (SDM) | 12 |
| Machine Learning (ML) | 14 |
| Viability of Business Projects (VBP) | 16 |
| Big Data Seminar (BDS) | 18 |
| Debates on Ethics of Big Data (DEBD) | 19 |
| Università degli Studi di Padova (UniPD) | 20 |
| Technische Universiteit Eindhoven (TU/e) | 20 |
| Business Information Systems (BIS) | 20 |
| Introduction to Process Mining (IPM) | 21 |
| Visualization (VIS) | 23 |
| Statistics for Big Data (SBD) | 24 |
| Business Process Analytics Seminar (BPAS) | 25 |
| CentraleSupélec (CS) | 26 |
| Decision Modeling (DeM) | 26 |
| Advanced Machine Learning (ML) | 27 |
| Visual Analytics (VA) | 28 |
| Massive Graph Management & Analytics (MGMA) | 29 |
| Big Data Research Project (BDRP) | 30 |
| Business Innovation Management (BIM) | 31 |

| |
|---|
| <p>University: Université Libre de Bruxelles (ULB) Department: École polytechnique de Bruxelles Course ID: ADB (INFO-H-415) Course name: Advanced Databases Name and email address of the instructors: Esteban Zimányi (ezimanyi@ulb.ac.be) Web page of the course: http://cs.ulb.ac.be/public/teaching/infoh415 Semester: 1 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 24h. • Exercises: 24h. • Projects: 12h. |
| <p>Goals: Today, databases are moving away from typical management applications, and address new application areas. For this, databases must consider (1) recent developments in computer technology, as the object paradigm and distribution, and (2) management of new data types such as spatial or temporal data. This course introduces the concepts and techniques of some innovative database applications</p> |
| <p>Learning outcomes: At the end of the course students are able to</p> <ul style="list-style-type: none"> • Understand various different technologies related to database management system • Understand when to use these technologies according to the requirements of particular applications • Understand different alternative approaches proposed by extant database management systems for each of these technologies • Understand the optimization issues related to particular implementation of these technologies in extant database management systems. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • R.T. Snodgrass. <i>Developing Time-Oriented Database Applications in SQL</i>, Morgan Kaufmann, 2000 • Jim Melton, Alan R. Simon. <i>SQL: 1999 - Understanding Relational Language Components</i>, Morgan Kaufmann, 2001 • Jim Melton. <i>Advanced SQL: 1999 - Understanding Object-Relational and Other Advanced Features</i>, Morgan Kaufmann, 2002 • Shashi Shekhar, Sanjay Chawla. <i>Spatial Databases: A Tour</i>, Prentice Hall, 2003. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Knowledge of the basic principles of database management, in particular SQL |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Active Databases Taxonomy of concepts. Applications of active databases: integrity maintenance, derived data, replication. Design of active databases: termination, confluence, determinism, modularisation. • Temporal Databases Temporal data and applications. Time ontology. Conceptual modeling of temporal aspects. Manipulation of temporal data with standard SQL. New temporal extensions in SQL 2011. • Object-Oriented and Object-Relational Databases Object-oriented model. Object persistence. ODMG standard: Object Definition Language and Object Query Language. .NET Language-Integrated Query: Linq. Object-relational model. Built-in constructed types. User-defined types. Typed tables. Type and table hierarchies. SQL standard and Oracle implementation. • Spatial Databases Application domains of Geographical Information Systems (GIS), common GIS data types and analysis. Conceptual data models for spatial databases. Logical data models for spatial databases: raster model (map algebra), vector model (OGIS/SQL1999). Physical data models for spatial databases: Clustering methods (space filling curves), storage methods (R-tree, Grid files). |
| <p>Assessment breakdown: 75% written examination, 25% project evaluation</p> |

| |
|---|
| <p>University: Université Libre de Bruxelles (ULB) Department: École polytechnique de Bruxelles Course ID: DBSA (INFO-H-417) Course name: Database Systems Architecture Name and email address of the instructors: Mahmoud Sakr (Mahmoud.Sakr@ulb.be) Web page of the course: https://www.ulb.be/en/programme/info-h417 Semester: 1 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 24h. • Exercises: 12h. • Projects: 24h. |
| <p>Goals:</p> <p>In contrast to a typical introductory course in database systems where one learns to design and query relational databases, the goal of this course is to get a fundamental insight into the implementation aspects of database systems. In particular, we take a look under the hood of relational database management systems, with a focus on query and transaction processing. By having an in-depth understanding of the query-optimisation-and-execution pipeline, one becomes more proficient in administering DBMSs, and hand-optimising SQL queries for fast execution.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student:</p> <ul style="list-style-type: none"> • Understands the workflow by which a relational database management systems optimises and executes a query • Is capable of hand-optimising SQL queries for faster execution • Understands the I/O model of computation, and is capable of selecting and designing data structures and algorithms that are efficient in this model (both in the context of datababase systems, and in other contexts). • Understands the manner in which relational database management systems provide support for transaction processing, concurrency control, and fault tolerance |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Hector Garcia-Molina, Jeffrey D. Ullman, Jennifer Widom. <i>Database Systems: The Complete Book</i>, Prentice Hall, second edition, 2008. • Raghu Ramakrishnan, Johannes Gehrke. <i>Database Management Systems</i>. McGraw-Hill, third edition, 2002. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Introductory course on relational databases, including SQL and relational algebra • Course on algorithms and data structures • Knowledge of the Java programming language |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Query Processing With respect to query processing, we study the whole workflow of how a typical relational database management system optimises and executes SQL queries. This entails an in-depth study of: <ul style="list-style-type: none"> – translating the SQL query into a “logical query plan”; – optimising the logical query plan; – how each logical operator can be algorithmically implemented on the physical (disk) level, and how secondary-memory index structures can be used to speed up these algorithms; and – the translation of the logical query plan into a physical query plan using cost-based plan estimation. • Transaction Processing <ul style="list-style-type: none"> – Logging – Serializability – Concurrency control |
| <p>Assessment breakdown:</p> <p>70% written examination, 30% project evaluation</p> |

University: Université Libre de Bruxelles (ULB)
Department: École polytechnique de Bruxelles
Course ID: DW (INFO-H-419)
Course name: Data Warehousing
Name and email address of the instructors: Esteban Zimányi (ezimanyi@ulb.ac.be)
Web page of the course: <http://cs.ulb.ac.be/public/teaching/infoh419>
Semester: 1
Number of ECTS: 5

Course breakdown and hours:

- Lectures: 24h.
- Exercises: 24h.
- Projects: 12h.

Goals:

Relational and object-oriented databases are mainly suited for operational settings in which there are many small transactions querying and writing to the database. Consistency of the database (in the presence of potentially conflicting transactions) is of utmost importance. Much different is the situation in analytical processing where historical data is analyzed and aggregated in many different ways. Such queries differ significantly from the typical transactional queries in the relational model:

1. Typically analytical queries touch a larger part of the database and last longer than the transactional queries;
2. Analytical queries involve aggregations (min, max, avg, ...) over large subgroups of the data;
3. When analyzing data it is convenient to see it as multi-dimensional.

For these reasons, data to be analyzed is typically collected into a data warehouse with Online Analytical Processing support. Online here refers to the fact that the answers to the queries should not take too long to be computed. Collecting the data is often referred to as Extract-Transform-Load (ELT). The data in the data warehouse needs to be organized in a way to enable the analytical queries to be executed efficiently. For the relational model star and snowflake schemes are popular designs. Next to OLAP on top of a relational database (ROLAP), also native OLAP solutions based on multidimensional structures (MOLAP) exist. In order to further improve query answering efficiency, some query results can already be materialized in the database, and new indexing techniques have been developed.

The first and largest part of the course covers the traditional data warehousing techniques. The main concepts of multidimensional databases are illustrated using the SQL Server tools. The second part of the course consists of advanced topics such as data warehousing appliances, data stream processing, data mining, and spatial-temporal data warehousing. The coverage of these topics connects the data warehousing course with and serves as an introduction towards other related courses in the program. Several associated partners of the program contribute to the course in the form of invited lectures, case studies, and “proof of technology” sessions.

Learning outcomes:

At the end of the course students are able to

- Explain the difference between operational databases and data warehouses and the necessity of maintaining a dedicated data warehousing system
- Understand the principles of multidimensional modeling
- Design a formal conceptual multidimensional model based on an informal description of the available data and analysis needs
- Implement ETL-scripts for loading data from operational sources a the data warehouse
- Deploy data cubes and extract reports from the data warehouse
- Explain the main technological principles underlying data warehousing technology such as indexing and view materialization.

Readings and text books:

- Matteo Golfarelli, Stefano Rizzi. *Data Warehouse Design: Modern Principles and Methodologies*. McGraw-Hill, 2009
- Christian S. Jensen, Torben Bach Pedersen, Christian Thomsen. *Multidimensional Databases and Data Warehousing*. Morgan and Claypool Publishers, 2010
- Ralph Kimball, Margy Ross. *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, third edition, Wiley, 2013.
- Esteban Zimányi, Alejandro Vaisman. *Data Warehouse Systems: Design and Implementation*. Springer,

2014

Prerequisites:

- A first course on database systems covering the relational model, SQL, entity-relationship modelling, constraints such as functional dependencies and referential integrity, primary keys, foreign keys.
- Data structures such as binary search trees, linked lists, multidimensional arrays.

Table of contents:

There is a mandatory project to be executed in three steps in groups of 3 students, using the tools learned during the practical sessions, being SQL Server, SSIS, SSAS, and SSRS. Below is the succinct summary of the theoretical part of the course:

- Foundations of multidimensional modelling
- Dimensional Fact Model
- Querying and reporting a multidimensional database with OLAP
- Methodological aspects for data warehouse development
- Populating a data warehouse: The ETL process
- Using the data warehouse: data mining and reporting

Assessment breakdown:

75% written examination, 25% project evaluation

University: Université Libre de Bruxelles (ULB)
Department: École polytechnique de Bruxelles
Course ID: BPM (INFO-H-420)
Course name: Business Process Management
Name and email address of the instructors: Dimitris Sacharidis (dimitris.sacharidis@ulb.be)
Web page of the course: <https://www.ulb.be/en/programme/info-h420>
Semester: 1
Number of ECTS: 5

Course breakdown and hours:

- Lectures: 24h.
- Exercises: 24h.
- Assignments and project: 12h.

Goals:

This course introduces basic concepts for modeling and implementing business processes using contemporary information technologies. The first part of the course considers the *identification* of the business processes within an organization and the *modeling* of business processes, including the control flow, and the data and resource perspectives. The workflow language *Business Process Modeling and Notation* (BPMN) is introduced in detail. The second part of the course then goes into the analysis, simulation, and redesign of business processes.

During the course the students have to perform a few modeling assignments in BPMN. In the project, students will focus on a larger assignment capturing various aspect of business process management.

Learning outcomes:

At the end of the course students are able to

- Explain the business process management cycle;
- Recognize and report the value and benefit as well as the limitations of the automation and management of a concrete business process for an organization;
- Identify business processes amenable for automation in a concrete business context;
- Design a formal model of the business process based on an informal description;
- Implement the formal description of a model into a business process management tool, including resource management;
- Analyze an existing business process management solution, and, based on this analysis, propose optimizations to improve its performance.

Readings and text books:

- Dumas, La Rosa, Mendling & Reijers. *Fundamentals of Business Process Management* (second edition), Springer 2018.

Prerequisites:

- Basic programming understanding.
- Basic set theory (notions such as set, set operations, sequence, multiset, function) and logics (mathematical notation and argumentation; basic proofs)
- Basic graph theory (notions such as graphs, reachability, transitivity, ...)
- Experience with modeling languages such as UML and ER diagrams is recommended.

Table of contents:

There are three assignments and a project to be realized in groups.

The theoretical part of the course is dedicated to topics that allow the students to successfully carry out the project. Below is a high-level overview of the theoretical part of the course:

- Short overview of enterprise systems architecture and the place of business process management systems in it. The BPM life cycle.
- Modeling business processes: modeling the control flow, data and resource perspective.
- Analyzing the business process models.
- Redesigning the business process models.

Assessment breakdown:

50% written examination, 50% project+assignment evaluation

University: Université Libre de Bruxelles (ULB)

Department: École polytechnique de Bruxelles

Course ID: DM (INFO-H-423)

Course name: Data Mining

Name and email address of the instructors: Mahmoud Sakr (Mahmoud.Sakr@ulb.be)

Web page of the course: <http://cs.ulb.ac.be/public/teaching/infoh423>

Semester: 1

Number of ECTS: 5

Course breakdown and hours:

- Lectures: 24h.
- Exercises: 24h.
- Projects: 12h.

Goals:

Data Mining aims at finding useful regularities in large structured and unstructured data sets. The goal of this course is to get a fundamental understanding of popular data mining techniques strengths and limitations, as well as their associated computational complexity issues. It will also identify industry branches which most benefit from Data Mining such as health care and e-commerce. The course will focus on business solutions and results by presenting case studies using real (public domain) data. The students will use recent Data Mining software.

Learning outcomes:

At the end of the course students are able to

- Establish the main characteristics and limitations of algorithms for addressing data mining tasks.
- Select, based on the description of a data mining problem, the most appropriate combination of algorithms to solve it.
- Develop and execute a data mining workflow on a real-life dataset to solve a data-driven analysis problem.
- Use recent data mining software for solving practical problems.
- Identify promising business applications of data mining.

Readings and text books:

- David J. Hand, Heikki Mannila, Padhraic Smyth. *Principles of Data Mining*. MIT Press, 2001.
- Delmater Rhonda, Hancock Monte. *Data Mining Explained*. Digital Press, 2001.
- Pang-Ning Tan, Michael Steinbach, Vipin Kumar. *Introduction to Data Mining*. Pearson Education (Addison Wesley), 2006.

Prerequisites:

- Programming experience
- Data structures
- Algorithms, complexity

Table of contents:

- Data mining
 - Data mining and knowledge discovery
 - Data mining functionalities
 - Data mining primitives, languages, and system architectures
- Data preprocessing
 - Data cleaning
 - Data transformation
 - Data reduction
 - Discretization and generating concept hierarchies
- Data mining algorithms
 - Motivation and terminology
 - Different algorithm types
- Classification and Clustering
 - Classification: SVM classifier
 - Clustering: K-means, Latent Semantic Analysis

- Mining Associations and Correlations
 - Item sets
 - Association rules
 - Generating item sets and rules efficiently
 - Correlation analysis
- Advanced techniques, data mining software, and applications
 - Text mining: extracting attributes (keywords), structural approaches (parsing, soft parsing).
 - Bayesian approach to classifying text
 - Web mining: classifying web pages, extracting knowledge from the web
 - Recommendation systems
 - Data mining software and applications

Assessment breakdown:

75% written examination, 25% project evaluation

| |
|---|
| <p>University: Universitat Politècnica de Catalunya (UPC) Department: Department of Service and Information System Engineering Course ID: BDM Course name: Big Data Management Name and email address of the instructors: Alberto Abelló (aabello@essi.upc.edu) Web page of the course: https://www.fib.upc.edu/en/studies/masters/master-innovation-and-research-informatics/curriculum/syllabus/BDM-MIRI Semester: 2 Number of ECTS: 6</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 27 h. • Laboratories: 27 h. • Self-Study: 96 h. |
| <p>Goals:</p> <p>The main goal of this course is to analyze the technological and engineering needs of Big Data Management. The enabling technology for such a challenge is cloud services, which provide the elasticity needed to properly scale the infrastructure as the needs of the company grow. Thus, students will learn advanced data management techniques (i.e., NOSQL solutions) that also scale with the infrastructure. Being Big Data Management the evolution of Data Warehousing, such knowledge (see the corresponding subject in Data Science speciality for more details on its contents) is assumed in this course, which will specifically focus on the management of data Volume and Velocity.</p> <p>On the one hand, to deal with high volumes of data, we will see how a distributed file system can scale to as many machines as necessary. Then, we will study different physical structures we can use to store our data in it. Such structures can be in the form of a file format at the operating system level, or at a higher level of abstraction. In the latter case, they take the form of either sets of key-value pairs, collections of semi-structured documents or column-wise stored tables. We will see that, independently of the kind of storage we choose, current highly parallelizable processing systems using functional programming principles (typically based on Map and Reduce functions), whose processing framework can rely on temporal files (like Hadoop MapReduce) or in-memory structures (like Spark).</p> <p>On the other hand, to deal with high velocity of data, we need some low latency system which processes either streams or micro-batches. However, nowadays, data production is already beyond processing technologies capacity. More data is being generated than we can store or even process on the fly. Thus, we will recognize the need of (a) some techniques to select subsets of data (i.e., filter out or sample), (b) summarize them maximizing the valuable information retained, and (c) simplify our algorithms to reduce their computational complexity (i.e., doing one single pass over the data) and provide an approximate answer.</p> <p>Finally, the complexity of a Big Data project (combining all the necessary tools in a collaborative ecosystem), which typically involves several people with different backgrounds, requires the definition of a high level architecture that abstracts technological difficulties and focuses on functionalities provided and interactions between modules. Therefore, we will also analyse different software architectures for Big Data.</p> <p>This course participates in a joint project conducted during the second semester together with VBP, SDM and DEBD. In VBP, the students will come up with a business idea related to Big Data Management and Analytics, which will be evaluated from a business perspective. In DEBD, they should analyse the same idea from an ethical perspective. Finally, in BDM and SDM they will be introduced to specific data management techniques to deal with Volume and Velocity (BDM) and Variety (SDM). Therefore, as final outcome, a working prototype dealing with those challenges must be delivered meeting the business idea created.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Knowledge <ul style="list-style-type: none"> – Understand the main advanced methods of data management and design and implement non-relational database managers, with special emphasis on distributed systems. – Understand, design, explain and carry out parallel information processing in massively distributed systems. – Manage and process a continuous flow of data. – Design, implement and maintain system architectures that manage the data life cycle in analytical environments. • Skills <ul style="list-style-type: none"> – Design a distributed database using NoSQL tools. |

- Produce a functional program to process Big Data in a cloud environment.
- Manage and process a stream of data.
- Design the architecture of a Big Data management system.

Readings and text books:

- M. Tamer Özsu, Patrick Valduriez. *Principles of Distributed Database Systems*, Springer, 2011.
- Ling Liu, M. Tamer Özsu. *Encyclopedia of Database Systems*, Springer, 2009.
- Pramod J. Sadalage, Martin Fowler. *NoSQL Distilled*, Addison-Wesley, 2013.
- Mark Grover, Ted Malaska, Jonathan Seidman, Gwen Shapira. *Hadoop Application Architectures*, O’Rilley, 2015.
- Hasso Plattner, Alexander Zeier. *In-Memory Data Management*, Springer, 2011.
- Matei Zaharia. *An Architecture for Fast and General Data Processing on Large Clusters*, ACM Press, 2016.
- Jure Leskovec, Anand Rajaraman, Jeffrey D. Ullman. *Mining of Massive Datasets*, <http://www.mmds.org/>
- Charu C. Aggarwal (Ed.). *Data Streams*, Springer, 2007.

Prerequisites:

- Advanced Databases ([ADB](#))
- Database Systems Architecture ([DBSA](#))
- Data Warehousing ([DW](#))

Table of contents:

- I Fundamentals of distributed and in-memory databases
 - 1 Introduction
 - Cloud computing concepts
 - Scalability
 - 2 In-memory data management
 - NUMA architectures
 - 3 Distributed Databases
 - Kinds of distributed databases
 - Fragmentation
 - Replication and synchronization (eventual consistency)
 - Distributed query processing and Parallelism
 - Bottlenecks of relational systems
- II High volume management
 - 4 Data management
 - Physical structures (Key-values, Document-stores, and Column-stores)
 - Distribution and placement
 - 5 Data processing
 - Functional programming
- III High velocity management
 - 6 Stream management
 - Sliding window
 - 7 Stream processing
 - Sampling
 - Filtering
 - Sketching
 - One-pass algorithms
- IV Architectures
 - 8 Software architectures
 - Centralized and Distributed functional architectures of relational systems
 - Lambda architecture

Assessment breakdown:

60% written examination, 40% project, +10% class participation

| |
|--|
| <p>University: Universitat Politècnica de Catalunya Department: Department of Service and Information System Engineering Course ID: SDM Course name: Semantic Data Management Name and email address of the instructors: Oscar Romero (oromero@essi.upc.edu) Web page of the course: https://www.fib.upc.edu/en/studies/masters/erasmus-mundus-master-big-data-management-and-analytics/curriculum/syllabus/SDM-BDMA Semester: 2 Number of ECTS: 6</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 27 h. • Laboratories: 27 h. • Self-study: 96 h. |
| <p>Goals:</p> <p>Big Data is traditionally defined with the three V's: Volume, Velocity and Variety. Big Data has been traditionally associated with Volume (e.g., the Hadoop ecosystem) and recently Velocity has earned its momentum (especially, with the arrival of Stream processors such as Spark Streaming). However, even if Variety has been part of the Big Data definition, how to tackle Variety in real-world projects is yet not clear and there are no standardized solutions (such as Hadoop for Volume or Streaming for Velocity) for this challenge.</p> <p>In this course the student will be introduced to advanced database technologies, modeling techniques and methods for tackling Variety for decision making. We will also explore the difficulties that arise when combining Variety with Volume and / or Velocity. The focus of this course is on the need to enrich the available data (typically owned by the organization) with external repositories (special attention will be paid to Open Data), in order to gain further insights into the organization business domain. There is a vast amount of examples of external data to be considered as relevant in the decision making processes of any company. For example, data coming from social networks such as Facebook or Twitter; data released by governmental bodies (such as town councils or governments); data coming from sensor networks (such as those in the city services within the Smart Cities paradigm); etc.</p> <p>This is a new hot topic without a clear and well-established (mature enough) methodology. The student will learn about semantic-aware data management as the most promising solution for this problem. As such, it will be introduced to graph modeling, storage and processing. Special emphasis will be paid to semantic graph modeling (i.e., ontology languages such as RDF and OWL) and its specific storage and processing solutions.</p> <p>This course participates in a joint project conducted during the second semester together. In VBP, the students will come up with a business idea related to Big Data Management and Analytics, which will be evaluated from a business perspective. In the three other courses the students have to implement a prototype meeting the business idea created. In BDM and SDM they will be introduced to specific data management techniques to deal with Volume and Velocity (BDM) and Variety (SDM). As final outcome, a working prototype must be delivered.</p> |
| <p>Learning outcomes:</p> <p>The student will learn models, languages, methods and techniques to cope with Variety in the presence of Volume and Velocity. Specifically:</p> <ul style="list-style-type: none"> • Graph modeling, storage and processing, • Semantic Models and Ontologies (RDF, RDFS and OWL), • Logics-based foundations of ontology languages (Description Logics) and • Specific storage and processing techniques for semantic graphs. <p>The students will learn how to apply the above mentioned foundations to automate the data management lifecycle in front of Variety. We will focus on semantic data governance protocols and the role of semantic metadata artifacts to assist the end-user.</p> |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Serge Abiteboul, Ioana Manolescu, Philippe Rigaux, Marie-Christine Rousset, Pierre Senellart, <i>Web Data Management</i>, Cambridge University Press, 2011. • Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, Peter F. Patel-Schneider, <i>The Description Logic Handbook</i>, Cambridge University Press, 2010. • Sven Groppe, <i>Data Management and Query Processing in Semantic Databases</i>, Springer, 2011. |

Prerequisites:

- Advanced Databases ([ADB](#))
- Data Warehousing ([DW](#))
- Database Systems Architecture ([DBSA](#))

Table of contents:

- Introduction:
 - Variety and Variability. Definition. External (non-controlled) data sources. Semi-structured data models, non-structured data. Schema and data evolution.
 - The data management life-cycle. Challenges when incorporating external (to the organisation) semi-structured and non-structured data.
- Foundations:
 - Graph management: the graph data model, graph databases, graph processing.
 - Semantic-aware management: Ontology languages (RDF, RDFS, OWL). Logical foundations and Description Logics. Storage (triplestores). Processing (SPARQL).
- Applications:
 - Open Data. Linked Open Data.
 - The Load-First Model-Later paradigm. The Data Lake. Semantic annotations and the Semantic-Aware Data Lake paradigm. Semantic data governance.
 - Semantic-aware metadata artifacts to automate data integration, linkage and / or cross of data between heterogeneous data sources.

Assessment breakdown:

40% written examination, 50% exercises and laboratories, 10% project

| |
|--|
| <p>University: Universitat Politècnica de Catalunya Department: Department of Computer Science Course ID: ML Course name: Machine Learning Name and email address of the instructors: Marta Arias (marias@cs.upc.edu) Web page of the course: https://www.fib.upc.edu/en/estudis/masters/master-en-ciencia-de-dades/pla-destudis/assignatures/ML-MDS Semester: 2 Number of ECTS: 6</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 27 h. • Laboratories: 27 h. • Self-study: 96 h. |
| <p>Goals:</p> <p>The aim of machine learning is the development of theories, techniques and algorithms to allow a computer system to modify its behavior in a given environment through inductive inference. The goal is to infer practical solutions to difficult problems –for which a direct approach is not feasible– based on observed data about a phenomenon or process. Machine learning is a meeting point of different disciplines: statistics, optimization and algorithmics, among others.</p> <p>The course is divided into conceptual parts, corresponding to several kinds of fundamental tasks: supervised learning (classification and regression) and unsupervised learning (clustering, density estimation). Specific modelling techniques studied include artificial neural networks and support vector machines. An additional goal is getting acquainted with python and its powerful machine learning libraries.</p> |
| <p>Learning outcomes:</p> <ul style="list-style-type: none"> • Formulate the problem of (machine) learning from data, and know the different machine learning tasks, goals and tools. • Organize the workflow for solving a machine learning problem, analyzing the possible options and choosing the most appropriate to the problem at hand. • Ability to decide, defend and criticize a solution to a machine learning problem, arguing the strengths and weaknesses of the approach. Additionally, ability to compare, judge and interpret a set of results after making a hypothesis about a machine learning problem. • To be able to solve concrete machine learning problems with available open-source software. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • C.M. Bishop, <i>Pattern recognition and machine learning</i>, Springer, 2006. • V.S. Cherkassky, F. Mulier, <i>Learning from data: concepts, theory, and methods</i>, John Wiley, 2007. • E. Alpaydin, <i>Introduction to machine learning</i>, The MIT Press, 2014. • K.P. Murphy, <i>Machine learning: a probabilistic perspective</i>, MIT Press, 2012. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Elementary notions of probability and statistics. • Elementary linear algebra and real analysis. • Good programming skills in a high-level language. |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Introduction to Machine Learning. General information and basic concepts. Overview to the problems tackled by machine learning techniques. Supervised learning (classification and regression), unsupervised learning (clustering and density estimation) and semi-supervised learning (reinforcement and transductive). Examples. • Supervised machine learning theory. The supervised Machine Learning problem setup. Classification and regression problems. Bias-variance tradeoff. Regularization. Overfitting and underfitting. Model selection and resampling methods. • Linear methods for regression. Error functions for regression. Least squares: analytical and iterative methods. Regularized least squares. The Delta rule. Examples. • Linear methods for classification. Error functions for classification. The perceptron algorithm. Novikoff's theorem. Separations with maxi- |

mum margin. Generative learning algorithms and Gaussian discriminant analysis. Naive Bayes. Logistic regression. Multinomial regression.

- Artificial neural networks.

Artificial neural networks: multilayer perceptron and a peak into deep learning. Application to classification and to regression problems.

- Kernel functions and support vector machines.

Definition and properties of Kernel functions. Support vector machines for classification and regression problems.

- Unsupervised machine learning.

Unsupervised machine learning techniques. Clustering algorithms: EM algorithm and k-means algorithm.

- Ensemble methods.

Bagging and boosting methods, with an emphasis on Random Forests.

40% final exam, 40% practical work and 20% mid-term exam.

| |
|--|
| <p>University: Universitat Politècnica de Catalunya (UPC) Department: Department of Management Course ID: VBP Course name: Viability of Business Projects Name and email address of the instructors: Marc Eguiguren (marc.eguiguren@upc.edu) Web page of the course: https://www.fib.upc.edu/en/studies/masters/master-innovation-and-research-informatics/curriculum/syllabus/VBP-MIRI Semester: 3 Number of ECTS: 6</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 36 h. • Projects in the classroom: 18 h. • Projects: 56 h. • Self-Study: 40 h. |
| <p>Goals:</p> <p>University graduates can find themselves in the situation of having to analyse or take on the project of starting their own business. This is especially true in the case of computer scientists in any field related to Big Data Management (BDM) or more generally, in the world of services. There are moments in one's professional career at which one must be able to assess or judge the suitability of business ventures undertaken or promoted by third parties, or understand the possibilities of success of dealing with a Big Data based service. It is for this reason that this subject focuses on providing students with an understanding of the main techniques used in analysing the viability of new business ventures: business start-up or the implementation of new projects in the world of services based on BDM. This project-oriented, eminently practical subject is aimed at each student's being able to draft as realistic a business plan as possible.</p> <p>This course participates in a joint project conducted during the second semester together with VBP, BDM and CC. In VBP, the students will come up with a business idea related to Big Data Management and Analytics, which will be evaluated from a business perspective. In the three other courses the students have to implement a prototype meeting the business idea created. In CC they will be introduced to the main concepts behind large-scale distributed computing based on a service-based model and will have to choose the right infrastructure for their prototype. In BDM and SDM they will be introduced to specific data management techniques to deal with Volume and Velocity (BDM) and Variety (SDM). As final outcome, a working prototype must be delivered.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Being able to analyze the external situation to determine business innovative ideas in the field of BDM • Around an innovative BDM project, being able to build a reasonable and ethically solid business plan • Building a solid and convincing speech about a business idea and a business plan • Training the students to build a P&L forecast and a forecasted treasury plan for a starting company • Understanding and being able to apply the different instruments to finance the company, both debt instruments or private equity and venture capital sources • Understand and appreciate the role of the entrepreneur in modern society |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Rhonda Abrams, Eugène Kleiner. <i>The Successful Business Plan</i>. The Planning Shop, 2003. • Rob Cosgrove. <i>Online Backup Guide for Service Providers</i>, Cosgrove, Rob, 2010. • Peter Drucker. <i>Innovation and Entrepreneurs</i>. Butterworth-Heinemann, Classic Drucker Collection edition, 2007. • Robert D. Hisrich, Michael P. Peter, Dean A. Shepherd. <i>Entrepreneurship</i>. Mc Graw Hill, 6th Ed., 2005. • Mike McKeever. <i>How to Write a Business Plan</i>. Nolo, 2010. • Lawrence W. Tuller. <i>Finance for Non-Financial Managers and Small Business Owners</i>. Adams Business, 2008. <p>Other Literature:</p> <ul style="list-style-type: none"> • M. Eguiguren, E. Barroso. <i>Empresa 3.0: políticas y valores corporativos en una cultura empresarial sostenible</i>. Pirámide, 2011. • M. Eguiguren, E. Barroso. <i>Por qué fracasan las organizaciones: de los errores también se aprende</i>. Pirámide, 2013. |

Prerequisites:

Having some previous knowledge or experience in business administration is an additional asset.

Table of contents: This course focuses on developing a BDM service oriented business plan. So that, the students are expected to reuse and consolidate any previous knowledge on databases, software engineering and BDM obtained in previous courses to develop a comprehensible, sustainable and profitable business.

The course is structured in 14 well-defined stages:

- Introduction to the course and key aspects of a business idea
- The entrepreneur's role in society, characteristics and profile
- Innovation and benchmarking Axis 1) Identification of long-term market megatrends
- Innovation and benchmarking axis 2) Big Data evolution as a source of ideas. Technology applied to industry.
- Axis of innovation and benchmarking 3) ethical business models as a source of innovation and ideas
- From the idea to the company. Contents of the business plan. Market research.
- Competitive advantages. SWOT Analysis
- Marketing plan: strategic marketing for a BDM service company, distribution and product
- Marketing plan: price and promotion strategies
- The human team in a small innovative company
- Different kind of societies. Fiscal basics for entrepreneurs
- Need of resources. Building the balance sheet at the beginning of the company
- Building a forecasted P&L for the first two years. Cash-Flow
- Revising the initial balance sheet and building the forecasted balance sheet for year one
- Treasury plan, Identifying long and short term financial needs
- Conventional long and short term financial instruments
- Private equity: founders, fools, friends & family, venture capital. Their limitations. Cautions to be taken and how they work.
- Presenting the plan to possible simulated or real investors

The business plan is expected to be partially developed in internal activities (under supervision of the teacher), and in external activities, always as teamwork (with no supervision).

Assessment breakdown:

The assessment is based on student presentations and the defence of the business plan before a jury comprising course faculty members and - optionally - another member of the teaching staff or guest professionals such as business angels, investors and successful IT entrepreneurs.

Throughout the course there will be four evaluative milestones:

- presentation of the innovative business model,
- presentation of the marketing plan,
- presentation of the business plan as a whole, that will include an evaluation about ethics and sustainability of the project together with SEAIT,
- analysis of the financial plan and the proposal to investors.

The presentation simulates a professional setting. Accordingly, the following aspects will also be assessed: dress, formal, well-structured communication, etc.

In order to be able to publicly defend the business plan, students must have attended at least 70% of the classes and teams must have delivered on time the activities that have been planned during the course. The plan is the result of teamwork, which will be reflected in the grade given to the group as a whole. Each member of the group will be responsible for part of the project and will be graded individually on his or her contribution.

This approach is designed to foster teamwork, in which members share responsibility for attaining a common objective.

| |
|--|
| <p>University: Universitat Politècnica de Catalunya (UPC) Department: Department of Service and Information System Engineering Course ID: BDS Course name: Big Data Seminar Name and email address of the instructors: Oscar Romero (oromero@essi.upc.edu) Web page of the course: To be created https://www.fib.upc.edu/en/studies/masters/erasmus-mundus-master-big-data-management-and-analytics/curriculum/syllabus/BDS-BDMA Semester: 2 Number of ECTS: 2</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 21 h. • Autonomous work: 129 h. |
| <p>Goals:</p> <p>The students will be introduced to recent trends in Big Data. Seminars will be lectured by guest speakers, who will present business cases, research topics, internships and master's thesis subjects. Also, the second year specialisations will be presented and discussed with the students within the seminars umbrella. Students will also perform a state-of-the art research in one topic, which will be presented and jointly evaluated by all partners in the mandatory eBISS summer school.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Read and understand scientific papers • Develop critical thinking when assessing scientific papers • Write and explain a state-of-the-art in a rigorous manner • Elaborate on recent trends in Big Data |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Gordana Dodig-Crnkovic, <i>Theory of Science</i>, online resource: http://www.idt.mdh.se/kurser/ct3340/archives/ht04/theory_of_science_compendiu.pdf • Jennifer Widom, <i>Tips for Writing Technical Papers</i>, online resource: https://cs.stanford.edu/people/widom/paper-writing.html • Robert Siegel, <i>Reading Scientific Papers</i>, online resource: https://web.stanford.edu/~siegelr/readingsci.htm |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • No prerequisites |
| <p>Table of contents:</p> <p>The seminars content will vary from course to course as they will focus on current hot topics in Big Data.</p> |
| <p>Assessment breakdown:</p> <p>50% Written report on a chosen state-of-the-art, 50% Poster presentation.</p> <p>Note: Attendance to the eBISS summer school is mandatory. To be evaluated (see the formula above) students must guarantee >75% overall attendance to the summer school events and the semester seminars.</p> |

| |
|--|
| <p>University: Universitat Politècnica de Catalunya (UPC) Department: Department of Service and Information System Engineering Course ID: DEBD Course name: Debates on Ethics of Big Data Name and email address of the instructors: Alberto Abelló (aabello@essi.upc.edu) Web page of the course: https://www.fib.upc.edu/en/studies/masters/erasmus-mundus-master-big-data-management-and-analytics/curriculum/syllabus/DEBD-BDMA Semester: 2 Number of ECTS: 2</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 18 h. • Autonomous work: 32 h. |
| <p>Goals:</p> <p>In this course we debate the impact on society of new advances in Big Data. We focus on ethics and the impact of such approaches on society. This course fosters the social competences of students, by building on their acquired oral communication skills to debate about concrete problems involving ethical issues in Big Data. The aim is to start developing their critical attitude and reflection. A written summary of their position is meant to train their writing skills.</p> <p>During the course sessions the debates discussing innovative and visionary ideas on Big Data will take place. You must read the available material before the debate. Then, during the debate you will be assigned to a group: either to defend an idea, or go against it. You may also be asked to moderate the debate. Then, the debate takes place and afterwards, each group needs to write down a report with their conclusions.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student will develop:</p> <ul style="list-style-type: none"> • Ability to study and analyze problems in a critical mood • Ability to critically read texts • Develop critical reasoning with special focus on ethics and social impact • Develop soft skills to defend - criticize a predetermined position in public • Improve the writing skills |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Rudi Volti, <i>Society and technological change</i>, Worth, 2009. • Richard T. De George, <i>The Ethics of information technology and business</i>, Blackwell, 2003. • David Elliott, <i>Energy, society, and environment: technology for a sustainable future</i>, Routledge, 2003. • Kord Davis, <i>Ethics of big data</i>, O'Reilly, 2013. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • No prerequisites |
| <p>Table of contents:</p> <p>Debates on ethics and social impact of Big Data (some examples):</p> <ul style="list-style-type: none"> • Ethics codes • Right to privacy • Content piracy • Social responsibility of IT companies • Artificial Intelligence and their limits |
| <p>Assessment breakdown:</p> <p>80% Debate evaluations, 20% Ethical analysis of a data-based business idea</p> |

| |
|--|
| <p>University: Technische Universiteit Eindhoven (TU/e) Department: Department of Mathematics and Computer Science Course ID: BIS Course name: Business Information Systems Name and email address of the instructors: Dirk Fahland (d.fahland@tue.nl) Web page of the course: To be created Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 32 h. • Instructions: 32 h. |
| <p>Goals: In this course students will learn about the modelling, analysis, and enactment of business processes and the information systems to support these processes, understanding the relationship between systems and processes.</p> |
| <p>Learning outcomes: Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Model complex (business) processes, systems and web-service choreographies in terms of classical and Colored Petri nets. • Translate informal requirements into explicit models. • Analyze processes (and the corresponding systems) using state-space analysis and invariants. • Analyze and improve processes (and the corresponding systems) through simulation. • Suggest redesigns of existing processes (and the corresponding systems). |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Wil van der Aalst, Christian Stahl. <i>Modeling Business Processes: A Petri Net-Oriented Approach</i>, MIT Press, 2011. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Datamodelling and databases (recommended) • Programming (recommended) • Logic and set theory (recommended) • Automata and process theory (recommended) |
| <p>Table of contents: Process-aware information systems (e.g., workflow management systems, ERP systems, CRM systems, PDM systems) are generic information systems that are configured on the basis of process models. In some systems the process models are explicit and can be adapted (e.g., the control-flow in a workflow system) while in other systems they are implicit (e.g., the reference models in the context of SAP). In some systems they are hard-coded and in other systems truly configurable. However, it is clear that in any enterprise, business processes and information systems are strongly intertwined. Therefore, it is important that students understand the relationship between systems and processes and are able to model complex systems involving processes, humans, and organizations.</p> |
| <p>Assessment breakdown: 70% Written final exam, 20% Assignment(s), 10% Interim examination Participation in the intermediate exam and the CPN assignment are required for entrance to the written exam. The points for the intermediate exam and the assignment are retained for the retake exam.</p> |

| |
|--|
| <p>University: Technische Universiteit Eindhoven (TU/e) Department: Department of Mathematics and Computer Science Course ID: IPM Course name: Introduction to Process Mining Name and email address of the instructors: Wil van der Aalst (w.m.p.v.d.aalst@tue.nl) Web page of the course: To be created Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 12 h. (recorded on video) • Instructions: 14 h. • Videos and self-tests of the MOOC https://www.coursera.org/course/procmin are used <p>A form of blended learning is used: Students learn content online by watching video lectures and homework is discussed in class</p> |
| <p>Goals:</p> <p>In this course students will acquire the theoretical foundations of process mining and will be exposed to real-life data sets helping them understand the challenges related to process discovery, conformance checking, and model extension.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Have a good understanding of process mining. • Understand the role of data science in today's society. • Relate process mining techniques to other analysis techniques such as simulation, business intelligence, data mining, machine learning, and verification. • Apply basic process discovery techniques such as the alpha algorithm to learn a process model from an event log (both manually and using tools). • Apply basic conformance checking techniques (such as token-based replay) to compare event logs and process models (both manually and using tools). • Extend a process model with information extracted from the event log (e.g., show bottlenecks). • Have a good understanding of the data needed to start a process mining project. • Characterize the questions that can be answered based on such event data. • Explain how process mining can also be used for operational support (prediction and recommendation). • Use tools such as ProM and Disco. • Execute process mining projects in a structured manner using the L* life-cycle model. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Wil van der Aalst, <i>Process Mining: Data Science in Action</i>, Springer-Verlag, Berlin, 2016. • Video lectures based on https://www.coursera.org/course/procmin. • ProM, Disco, and RapidMiner are used for hands-on experiments. • Provided slides, event logs, exercises, and additional papers. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Basic computer science skills (e.g. bachelor in computer science, mathematics or industrial engineering). |
| <p>Table of contents:</p> <p>Data science is the profession of the future, because organizations that are unable to use (big) data in a smart way will not survive. It is not sufficient to focus on data storage and data analysis. The data scientist also needs to relate data to process analysis. Process mining bridges the gap between traditional model-based process analysis (e.g., simulation and other business process management techniques) and data-centric analysis techniques such as machine learning and data mining. Process mining seeks the confrontation between event data (i.e., observed behavior) and process models (hand-made or discovered automatically). This technology has become available only recently, but it can be applied to any type of operational processes (organizations and systems). Example applications include: analyzing treatment processes in hospitals, improving customer service processes in a multinational, understanding the browsing behavior of customers using a booking site, analyzing failures of a baggage handling system, and improving the user interface of an X-ray machine. All of these applications have in common that dynamic behavior needs to be related to process models. Hence, we refer to this as “data science in action”.</p> <p>The course covers the three main types of process mining. The first type of process mining is discovery. A</p> |

discovery technique takes an event log and produces a process model without using any a-priori information. An example is the a-algorithm that takes an event log and produces a Petri net explaining the behavior recorded in the log. The second type of process mining is conformance. Here, an existing process model is compared with an event log of the same process. Conformance checking can be used to check if reality, as recorded in the log, conforms to the model and vice versa. The third type of process mining is enhancement. Here, the idea is to extend or improve an existing process model using information about the actual process recorded in some event log. Whereas conformance checking measures the alignment between model and reality, this third type of process mining aims at changing or extending the a-priori model. An example is the extension of a process model with performance information, e.g., showing bottlenecks. Process mining techniques can be used in an offline, but also online setting. The latter is known as operational support. An example is the detection of non-conformance at the moment the deviation actually takes place. Another example is time prediction for running cases, i.e., given a partially executed case the remaining processing time is estimated based on historic information of similar cases.

Process mining provides not only a bridge between data mining and business process management; it also helps to address the classical divide between “business” and “IT”. Evidence-based business process management based on process mining helps to create a common ground for business process improvement and information systems development. The course uses many examples using real-life event logs to illustrate the concepts and algorithms. After taking this course, one is able to run process mining projects and have a good understanding of the data science field.

Assessment breakdown:

Written final exam, Assignments

- Final test Introduction to Process Mining
- Assignment 1 Introduction to Process Mining
- Assignment 2 Introduction to Process Mining

| |
|--|
| <p>University: Technische Universiteit Eindhoven (TU/e) Department: Department of Mathematics and Computer Science Course ID: VIS Course name: Visualization Name and email address of the instructors: Michel Westenberg (m.a.westenberg@tue.nl) Web page of the course: http://www.win.tue.nl/~mwestenb/edu/2IMV20_Visualization/ Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 14 h. • Practical work: 14 h. • Video recordings are available as additional teaching material and are not supposed to replace the lectures. |
| <p>Goals:</p> <p>In this course students will learn the theory and practice of data visualization, including topics such as data representation, grid types, data sampling, data interpolation, data reconstruction, datasets, and the visualization pipeline.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Have a good knowledge of the theory, principles, and methods which are frequently used in practice in the construction and use of data visualization applications. • Design, implement, and customize a data visualization application of average complexity in order to get insight in a real-world dataset from one of the application domains addressed during the lecture. • Understand (and apply) the various design and implementation trade-offs which are often encountered in the construction of visualization applications. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Lecture slides, selected papers, and assignments. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Computer graphics (basic) • Linear algebra • Programming (some experience in Java or a similar language) |
| <p>Table of contents:</p> <p>The course covers the theory and practice of data visualization. This addresses several technical topics, such as: data representation; different types of grids; data sampling, interpolation, and reconstruction; the concept of a dataset; the visualization pipeline. In terms of visualization application, several examples are treated, following the different types of visualization data: scalar visualization, vector visualization, tensor visualization. Besides these, several additional topics are treated, such as volume data visualization and a brief introduction to information visualization. The techniques treated in the course are illustrated by means of several practical, hands-on, examples.</p> |
| <p>Assessment breakdown:</p> <p>50% First assignment, 50% Second assignment Two assignments covering the core topics of the course.</p> |

| |
|--|
| <p>University: Technische Universiteit Eindhoven (TU/e) Department: Department of Mathematics and Computer Science Course ID: SBD Course name: Statistics for Big Data Name and email address of the instructors: Edwin van den Heuvel (e.r.v.d.heuvel@tue.nl) Web page of the course: To be created Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 32 h. • Practical work: 32 h. |
| <p>Goals:</p> <p>In this course students will learn various statistical methods for analysing BD, focusing on analysing temporal observational data, i.e., data that is collected over time without involving well-developed experimental designs.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Have a good knowledge of theoretical and practical aspects of mixed models for different type of outcomes (continuous, categorical, binary). • Understand how model features will affect the correlations in outcomes and how bias in estimates are affected by confounders. • Identify the difference between two classes of models: marginal and subject specific models. • Apply the different methods of estimation (generalized estimating equations and maximum likelihood). • Apply different approaches to handle missing data (like multiple imputation methods). • Apply the methods on real data sets, with real problems using statistical software. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Handouts • Tutorial and overview papers from statistical literature |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Knowledge on probability theory and on modeling data with both linear and generalized linear models (preferably 2WS40 and 2WS70). |
| <p>Table of contents:</p> <p>The course discusses statistical models for analyzing relative large data sets with complex structured correlated data. Both individual participant data analysis and aggregate data analysis will be discussed. These methods are useful when data sets from different sources are being pooled for analysis. One important topic is how different but similar variables can be harmonized and/or standardized for such a pooled analysis. The family of models that will be typically discussed are mixed models. These models are not just relevant for big data alone, but also useful for small data sets. Missing data problems in relation to mixed models will be discussed in detail, since missing data can really influence conclusions if it is not addressed appropriately. We will also discuss issues that are related to selection bias (e.g. propensity scores and confounding) and causal inference.</p> |
| <p>Assessment breakdown:</p> <p>Written final exam, Assignments Homework assignments. Two assignments on the analysis of data sets with complex correlated data. Another assignment is about a simulation study to investigate the performance of the discussed theory on statistical modeling. Both content and reporting will be evaluated.</p> |

| |
|---|
| <p>University: Technische Universiteit Eindhoven (TU/e) Department: Department of Mathematics and Computer Science Course ID: BPAS Course name: Business Process Analytics Seminar Name and email address of the instructors: H.A. Reijers (h.a.reijers@vu.nl) Web page of the course: http://www.win.tue.nl/is/doku.php?id=education:2ii96 Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 14 h. |
| <p>Goals:</p> <p>This seminar introduces students to research in business process analytics by following lectures by staff members and guest lecturers from industry, analysing master theses, study research papers, and execute a small research project.</p> |
| <p>Learning outcomes:</p> <p>Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • Have a good knowledge about the AIS research topics. • Analyze selected Master theses performed within the AIS group to learn about the research topics the group works on and to get a better idea about “patterns” and “anti-patterns” in the execution of a Master project and writing a thesis. The results of this work will be given in a review of a Master thesis and a presentation about it. • Use existing tools such as ProM, Disco, CPN Tools to work on their research project |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Various items (offered via the website or studyweb) |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Process modelling (recommended) • Process mining (recommended) • Business information systems (recommended) • Data mining and knowledge systems (recommended) • Business process management systems (recommended) • Visualization (recommended) |
| <p>Table of contents:</p> <p>Organizations are constantly trying to improve the way their businesses perform. To this end, managers have to take decisions on changes to the operational processes. However, for decision making, it is of the utmost importance to have a thorough understanding of the operational processes as they take place in the organization.</p> <p>Typically, such operational processes are supported by information systems that record events as they take place in real life in so-called event logs. Process Mining techniques allow for extracting information from event logs. For example, the audit trails of a workflow management system or the transaction logs of an enterprise resource planning system can be used to discover models describing processes, organizations, and products. Moreover, it is possible to use process mining to monitor deviations (e.g., comparing the observed events with predefined models or business rules in the context of SOX).</p> <p>Process mining is closely related to BAM (Business Activity Monitoring), BOM (Business Operations Management), BPI (Business Process Intelligence), and data/workflow mining. Unlike classical data mining techniques the focus is on processes and questions that transcend the simple performance-related queries supported by tools such as Business Objects, Cognos BI, and Hyperion.</p> <p>In this seminar, students are introduced to the research being conducted in the Architecture of Information Systems group of this university. Specifically, this seminar focuses on the state-of-the-art research on Process Mining. To emphasize that Process Mining is not only a research area, but has gained increasing interest from Industry, practical application will be considered.</p> |
| <p>Assessment breakdown:</p> <p>Assignment(s)</p> |

| |
|---|
| <p>University: CentraleSupélec, Paris-Saclay University Department: Computer Science Department Course ID: DeM Course name: Decision Modeling Name and email address of the instructors: Prof. Brice Mayag (brice.mayag@dauphine.fr) Web page of the course: (to be created) Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 24h • Laboratory: 24h • Project: 12h |
| <p>Goals: This course aims at presenting classical decision models with a special emphasis on decision making in uncertain situations, decision with multiple attribute, and decision with multiple stakeholders. During the course, various applications will be presented, emphasizing the practical interest and applicability of the models in real-world decision situations.</p> |
| <p>Learning outcomes: Upon successful completion of this course, the student will acquire knowledge and skills about:</p> <ul style="list-style-type: none"> • decision models, validity of the decision models • the three levels of decision analysis: representation of observed decision behavior (descriptive analysis), decision aiding and recommendation (prescriptive analysis), and the design of artificial decision agents (normative analysis). |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Denis Bouyssou, Thierry Marchant, Marc Pirlot, Alexis Tsoukiás, Philippe Vincke, <i>Evaluation and decision models with multiple criteria: Stepping stones for the analyst</i>, Springer, International Series in Operations Research and Management Science Volume 86, 2006. • William W. Cooper, Lawrence M. Seiford, and Kaoru Tone, <i>Introduction to Data Envelopment Analysis and Its Uses</i>, Springer, 2006. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Operational research algorithms foundations. |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Data envelopment Analysis: analysis of the efficiency of the production units. • Decision under uncertainty and decision trees: theory, modeling and applications. • Behavioural decision analysis: empirical analysis of decision behaviour, cognitive decision biases, prospect theory. • Outranking methods (theory and applications): presentation of the Electre methods (Electre I, Electre 3, Electre Tri), reference based ranking. • Applications on a generic Decision platform: decision Deck. Analysis of some use cases and use of an open source platform for decision aid. • Group decision: group decision, elicitation of a group decision model. • Preference learning: eliciting preference model for a decision maker, for several decision makers. • Decision making using Multiple Objective Optimisation: epsilon constraint method, applications, approximation algorithms, evolutionary algorithms, and NSGA II. |
| <p>Assessment breakdown: Homework and class participation (10%), Written exam (50%), Project (40%)</p> |

| |
|--|
| <p>University: CentraleSupélec, Paris-Saclay University Department: Computer Science Department Course ID: ML Course name: Machine Learning Name and email address of the instructors: Prof. Antoine Cornuéjols (Antoine.Cornuejols@Iri.fr) Web page of the course: (to be created) Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 24h • Laboratory: 24h • Project: 20h |
| <p>Goals: The goal of this course is to provide the student with knowledge about supervised, unsupervised and reinforcement learning paradigms; the mathematical foundations and practices of different variants of machine learning methods.</p> |
| <p>Learning outcomes: Upon successful completion of this course, the student will be able to:</p> <ul style="list-style-type: none"> • choose the best techniques to solve a given machine learning task; • tune the parameters of the chosen method; • interpret the results and compare different learning methods. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • David J. Hand, Heikki Mannila and Padhraic Smyth, <i>Principles of Data Mining</i>, The MIT Press, 2001. • Trevor Hastie, Robert Tibshirani, Jerome Friedman, <i>The Elements of Statistical Learning: Data Mining, Inference, and Prediction</i>, Second Edition, Springer, 2009. • Christopher Bishop, <i>Pattern Recognition and Machine Learning</i>, Springer, 2006. • Koller and Friedman, <i>Probabilistic Graphical Models</i>, MIT Press, 2009. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Familiarity with the basic probability and linear algebra theory |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Supervised, unsupervised, and reinforcement learning paradigms. • Linear and logistic regression: gradient descent, locally weighted regression, exponential families, generalized linear models. • Generative learning algorithms, Gaussian discriminant analysis, Kernel methods, neural networks as functions, deep neural networks. • Feature selection, curse of dimensionality, dimensionality reduction • Probabilistic graphical models (HMM, MRF, Bayesian Networks, Inference, Kalman filters). • Problems: on-line learning, multi-task learning. |
| <p>Assessment breakdown: Homework and class participation (10%), Written exam (50%), Project (40%)</p> |

| |
|--|
| <p>University: CentraleSupélec, Paris-Saclay University Department: Computer Science Department Course ID: VA Course name: Visual Analytics Name and email address of the instructors: Petra Isenberg (petra.isenberg@inria.fr) Web page of the course: (to be created) Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 24h • Laboratory: 24h • Project: 12h |
| <p>Goals: In this course students learn how to bring together automated and human-driven data analysis approaches; including innovative aspects such as: data collection, data cleaning, basic statistics, exploratory data analysis, perception and cognition, storytelling, text analysis, and multi-dimensional data representation.</p> |
| <p>Learning outcomes: Upon completing the course, the student will be able to:</p> <ul style="list-style-type: none"> • understand basic concepts, theories, and methodologies of Visual Analytics; • analyse data using appropriate visual analytics thinking and techniques; • present data using appropriate visual communication and graphical methods; • design and implement a Visual Analytics system for supporting decision making. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Edward Tufte, <i>Envisioning Information</i>, Graphics Press, 1990. • Robert Spence, <i>Information Visualisation: Design for Interaction</i>, Second Edition, Prentice Hall, 2007. • Colin Ware, <i>Information Visualisation: Perception for Design</i>, Second Edition, Morgan Kaufmann, 2004. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Database systems • Data mining foundations |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • VA fundamentals: theories, methodologies and techniques; • Designing interactive graphics; • Appropriate methods for different data types: Graphs, Hierarchies, Spatio-temporal data, High Multi-dimensional, Text; • Data Analysis under Uncertainty, Visualizing and exposing uncertainty; • VA system design practices, Dashboard design; • Exploratory Data Analysis with Tableau/Microsoft Power BI. Data Scraping in R/Python. Data Cleaning & Wrangling using OpenRefine. |
| <p>Assessment breakdown: Homework and class participation (10%), Written exam (50%), Project (40%)</p> |

| |
|---|
| <p>University: CentraleSupélec, Paris-Saclay University Department: Computer Science Department Course ID: MGMA Course name: Massive Graph Management & Analytics Name and email address of the instructors: Prof. Binh-Minh Bui-Xuan (buixuan@lip6.fr) Web page of the course: (to be created) Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 24h • Laboratory: 24h • Projects: 12h |
| <p>Goals: The objectives of this course is to provide the student with knowledge about designing high-performance and scalable algorithms for massive graph analytics. The course focuses on modeling and querying massive graph data in a distributed environment, designing algorithms, complexity analysis and optimization, for massive data graph problem analytics.</p> |
| <p>Learning outcomes: Upon successful completion of this course, the student is able to:</p> <ul style="list-style-type: none"> • model and query massive graph data in a distributed environment • design and analyse efficient graph algorithms in real-world data-intensive applications; • develop efficient applications using the best practices in a distributed environment (Spark, MapReduce, Neo4J, GraphX, etc.). |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Richard Brath, David Jonker, <i>Graph Analysis and Visualization: Discovering Business Opportunity in Linked Data</i>, Wiley, 2015 • Ian Robinson, Jim Webber, Emil Eifrem <i>Graph Databases</i>, O’Reilly, 2013 |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Graph theory foundations • NoSql databases in distributed environment |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Modeling massive graph data in a distributed NoSql databases • Querying massive graph data in a distributed environment • Graph Search, Spanning Tree, Betweenness Centrality, Community Detection, Connected Components, Minimum Spanning Tree, Anomaly Detection • Streaming Data Analysis, Data Structures for Streaming Data • Analyzing data Streams (e.g. Twitter) and Bioinformatics data |
| <p>Assessment breakdown: Homework and class participation (10%), Written exam (50%), Project (40%)</p> |

| |
|---|
| <p>University: CentraleSupélec, Paris-Saclay University Department: Computer Science Department Course ID: BDRP Course name: Big Data Research Project Name and email address of the instructors: Profs. Nacéra Seghouani Bennacer & Francesca Bugiotti (nacera.bennacer@centralesupelec.fr, francesca.bugiotti@centralesupelec.fr) Web page of the course: (to be created) Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 18 h. • Seminars: 12 h. • Projects: 24 h. |
| <p>Goals: This course aims at preparing the students for the master thesis of the 4th semester. The students will learn how to manage a research project related to massive and heterogeneous data management and analytics from scratch, working in a team, and using all the steps required in a scientific methodology. During this course the students will attend seminars in order to have a better understanding of research methodologies and to be aware of some ongoing research projects presented by researchers.</p> |
| <p>Learning outcomes: Upon successful completion of this course, the student is able to manage a scientific project from scratch in a team and provide scale-up algorithms and a program prototype for massive data management and analysis.</p> |
| <p>Readings and text books: Scientific papers will be distributed by the lecturer and the invited speakers according to the covered topics.</p> |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Bachelor program in computer science • Basic programming languages such as Python, Java and C++ • Basic knowledge in distributed environments (SPARK, NoSql databases) |
| <p>Table of contents: The students will choose a project proposal and follow the following steps:</p> <ul style="list-style-type: none"> • Selection of relevant related works, writing citations and bibliography according to author rights (plagiarism issues); • Summarize and classify a selection of related works; • Formalize a research problem using the appropriate notation; • Provide a formalized solution studying its complexity and proving its properties, to compare with existing approaches; • Implement a prototype program (with documentation) provide experiments with detailed evaluation; • Write a final report, and final presentation. |
| <p>Assessment breakdown: 20% team project progress, 10% intermediate presentations, 10%Final defense, 30% Final report, 30% Final prototype & results</p> |

| |
|--|
| <p>University: CentraleSupélec, Paris-Saclay University Department: Computer Science Department Course ID: BIM Course name: Business Innovation Management Name and email address of the instructors: Mr Karim Tadrist (legal affairs of Paris-Saclay) and Invited speakers from companies (karim.tadrist@universite-paris-saclay.fr) Web page of the course: (to be created) Semester: 3 Number of ECTS: 5</p> |
| <p>Course breakdown and hours:</p> <ul style="list-style-type: none"> • Lectures: 30h • Homework: 20h |
| <p>Goals: The objectives of this course are to provide the student: (i) knowledge about intellectual and industrial properties, data protection and security in European research context, (ii) an overview about current and innovative company projects and technology needs for real data analytics and machine learning.</p> |
| <p>Learning outcomes: Upon completing the course, the student will acquire knowledge and skills about:</p> <ul style="list-style-type: none"> • intellectual and industrial properties, data protection in European research context; • corporate and entrepreneurship culture; • innovative projects and technologies related to massive and real data management, analytics and machine learning. |
| <p>Readings and text books:</p> <ul style="list-style-type: none"> • Scientific papers will be distributed by the course lecturer and invited speakers according to the topics covered. |
| <p>Prerequisites:</p> <ul style="list-style-type: none"> • Bachelor program in computer science |
| <p>Table of contents:</p> <ul style="list-style-type: none"> • Lectures: Intellectual property, Industrial property, Patent law, Author rights, Data protection (European General Data Protection Regulation) and security, Sensitive data, Legal tools for startups (incubators, valorisation). • Various seminars led by invited speakers from companies, key BI actors, and startups. |
| <p>Assessment breakdown: Oral exam (50%), Written report (50%)</p> |