

Master in Artificial Intelligence (2022)

91248 - Fundamentals of Artificial Intelligence and Knowledge Representation

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student knows the main knowledge representation techniques and reasoning methods that underlie artificial intelligence problem solving. The student is able to develop simple solvers for artificial intelligence systems.

Course contents

The course introduces the fundamental principles and methods used in Artificial Intelligence to solve problems, with a special focus on the search in the state space, planning, knowledge representation and reasoning, and on the methods for dealing with uncertain knowledge. The course will include hands-on labs and seminars on selected topics.

Prerequisites: user-level knowledge of a high-level programming language, in order to successfully understand case studies and applications presented during the lessons.

Part I (Module 1)

Module 1

Introduction to Artificial Intelligence: historical perspective, main application fields, introduction to knowledge-based systems and architectural organization.

Problem-solving in AI: representation through the notion of state, forward e backward reasoning, solving as a search and search strategies. Games. Constraint satisfaction problems.

Local Search methods, meta heuristics, solving through decomposition, constraint relaxation, branch-and-bound techniques

Introduction to Planning, Linear planning, partial order planning, graph-based methods (GraphPlan), Scheduling.

Part II (Modules 2, 3):

Module 2

Introduction to knowledge representation and reasoning

Representing Terminological Knowledge: semantic networks, description logics, foundation of ontologies

Representing actions, situations, and events.

Reasoning with Beliefs.

Nonmonotonic Reasoning and reasoning with default Information.

Rule-based systems: Prolog and extensions, meta-interpreters, DCG, planning in prolog, Prolog for temporal reasoning with the Event Calculus, LPAD.

Forward chaining and RETE, Drools

Limits of KR based on Logic: Abductive, Inductive, and case-based reasoning. Fuzzy reasoning.

Module 3

Uncertainty

Probabilistic Reasoning

91255 - Statistical and Mathematical Methods for Artificial Intelligence

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student masters the basic mathematical and statistical methods needed to acquire skills in artificial intelligence foundations, theories and applications.

Course contents

Unless new restrictions, the course will be held in blended modality. See the timetable for more details.

1.Elements of linear algebra.

vectorial calculus, linear mappings, normed spaces, orthogonal projections.

matrix calculus, matrix norms, special matrices.

Singular Values Decomposition, Principal Component Analysis.

Eigenvalues and eigenvectors.

Laboratory exercises in Matlab or Phyton.

1. Elements of multivariate analysis

Gradient, Jacobian, Hessian. Taylor theorem.

Convex functions and sets.

1. Multivariate optimization

Linear least squares.

Extrema of multivariate functions. Optimality conditions.

Descent methods. Gradient type methods and Newton type methods.

Regularization.

Basis concepts of stochastic optimization.

Laboratory exercises in Matlab or Phyton.

1. Elements of probability and statistics.

Probability and Bayes theorem.

Random variables. Continuous and discrete distributions of random variables. Normal and Poisson distributions.

Independent and dependent variables. Covariance and correlation.

Estimates: Maximum Likelihood and Maximum a Posteriori estimates.

Cross entropy and Kullback-Leibler divergence.

Laboratory exercises in Matlab or Phyton.

Reading

91247 - Cognition and Neuroscience

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student knows state-of-art human and animal research that uses neuroscience techniques to understand the cognitive and emotional aspects of the human mind and behavior. The student is able to critically read experimental and theoretical studies of cognitive and affective neuroscience, to evaluate their methods and results, explain their significance, and apply such notions in the study and development of artificial intelligence systems.

Course contents

How do neurons in the brain give rise to mind – to our abilities to sense and perceive the world, to act and think about it, to learn and remember it? This course will provide an accessible but highly

challenging survey of the empirical evidence, theories and methods in cognitive neuroscience exploring how cognition is instantiated in neural activity. Drawing on a wide variety of investigative tools available to cognitive neuroscience, the course explores the neural mechanisms underlying complex cognitive processes, including visual perception, selective attention, learning and emotion, cognitive control and decision-making. Importantly, the course examines the logic and assumptions that permit us to interpret brain activity in psychological terms.

At the end of the course the student is be able to:

get in-depth understanding of the neural substrates and functional mechanisms of mental processes;

get knowledge of the-state-of-the-art methodologies and novel approaches of current research in Cognitive Neuroscience;

critically review and discuss the theoretical and empirical contributions of the current literature, understand and analyze the methods employed, interpret their results and critically assess their conclusions;

exercise the ability to engage in creative thinking leading to formulations of new hypotheses and planning of their empirical testing.

The course Cognition & Neuroscience is divided into two teaching modules.

Module 1 (4 CFU):

Associative learning

Reinforcement learning

Decision-making

Module 2 (2 CFU):

Vision

Visual object recognition

Visual selective attention

Previous knowledge required:

Prerequisite involves high-school knowledge of the biology of the Central Nervous System.

In addition, it is highly recommended to attend to online video lectures on Neuroscience Core concepts, freely available at the Society for Neuroscience website: <https://www.brainfacts.org/core-concepts>

91762 - Combinatorial Decision Making and Optimization

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has an understanding of the most popular methods in operations research and artificial intelligence for modeling and solving complex combinatorial optimization problems such as constraint programming, integer linear programming, SAT and SMT.

Course contents

Prerequisites

Problem solving in artificial intelligence (covered by 91248 - Fundamentals of Artificial Intelligence and Knowledge Representation), logic for computer science and basic notions of algorithms and programming preferably in Python (covered by 91251 - Languages and Algorithms for Artificial Intelligence), basic algebra and analysis.

Course topics

The first half of the course (Module 1) will introduce the fundamental concepts in Constraint Programming (CP) and Boolean Satisfiability (SAT), which are two general-purpose approaches in Artificial Intelligence (AI) to combinatorial decision making and optimization. The topics include:

Introduction to combinatorial decision making and optimization

Modelling in CP

Constraint propagation and global constraints

Constructive tree search and propagation, branching heuristics, randomization and search restarts, Large Neighbourhood Search (LNS)

The SAT problem and encoding in SAT

Basic solving techniques (resolution, unit propagation, DPLL algorithm)

Conflict-driven clause learning algorithms

SAT encodings

The second half of the course (Module 2) will introduce the fundamentals of Satisfiability Modulo Theory (SMT), which is a generalization of SAT. The last part of the module will be about essential Operations Research results from Linear Programming (LP) and in particular Integer Linear Programming (ILP), by prioritizing the width of the covered area over the depth of the introduction. The topics include:

Satisfiability Modulo Theories (SMT)

Eager vs lazy SMT encoding, the DPLL(T) algorithm

SMT theory solvers and combinations of theories

Linear Programming (LP)

Simplex algorithm and duality

Integer Linear Programming (ILP)

Extensions and hybrid methods

The course will also host invited lectures by pioneer researchers in the field in order to provide a broader perspective on the AI-based approaches to combinatorial decision making and optimization.

91254 - Image Processing and Computer Vision

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student masters the basic algorithms, tools and systems for the management, processing and analysis of digital images. The student is able to design and develop simple systems oriented to real-world computer vision applications such as those requiring segmentation and classification of objects in digital images.

Course contents

Introduction: Basic definitions related to computer vision. An overview across major application domains.

Image Formation and Acquisition: geometry of image formation; lenses; field of view and depth of field; image sampling and quantization.

Spatial Filtering: Linear shift-invariant operators. Mean and Gaussian filtering. Median Filtering. Bilateral filtering.

Edge Detection: image gradient; non-maxima suppression; Laplacian of Gaussian; Canny edge detector.

Local Invariant Features: detectors and descriptors; Harris Corners; scale-invariant features; SIFT features;

Object Detection: pattern matching; shape-based matching; Hough transform; object detection by local invariant features; Hough-based voting.

Camera calibration: Projective coordinates and perspective projection matrix. Intrinsic and extrinsic camera parameters. Zhang's algorithm. Image rectification and stereo calibration.

Basics of Machine learning for Images: k-NN classifier and linear classifiers as image template matching.

Image representations: BoVW, VLAD, representation learning for images, localist versus distributed representations.

Convolutional Neural Networks: advantages of convolutions, vector-valued functions, LeNet5, normalization layers.

Successful architectures: AlexNet, ZFNet, VGG, Inception v1 and v3, ResNet, Inception-ResNet-v2.

The theory is complemented by lab sessions. Lab sessions cover selected topics such as spatial filtering, camera calibration, local invariant features, image classification and transfer learning.

91251 - Languages and Algorithms for Artificial Intelligence

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has an understanding of the main linguistic and algorithmic techniques used in the context of AI, including the main aspects related to functional, logic and constraint programming, with their theoretical foundations and practical applications. The student has an understanding of the basic principles of distributed models and algorithms, consensus problems, and large-scale algorithms for the analysis of big data.

Course contents

Module 1

Abstract machines. Python machine. Programming in Python: names and visibility, functions, immutable and mutable objects, basic data types (numbers, strings, tuples, lists, dictionaries) and their use in the solution of problems. Classes and objects. Methods and inheritance. Exceptions. Introduction to the library NumPy and its N-dimensional array objects.

Module 2

Introduction to mathematical logic. Unification. Resolution.

Introduction to logic programming. Prolog languages.

Constraint logic programming and concurrent constraint programming. Constraint programming. The language MiniZinc.

Module 3

Rudiments of Computability and Complexity Theory. Undecidability, polynomial time computability, and NP-completeness. Between the feasible and the unfeasible. A Formal Learning Model: PAC Learning. Learning via Uniform Convergence. The Bias-Complexity Tradeoff.

The VC-Dimension.

Module 4

Programming and learning paradigm, and the languages for inference. Introduction to the language Scheme/Racket and its probabilistic extensions Church/Gamble. Generative models as probabilistic programs. Bayesian programming in Church/Gamble. Learning as conditional inference. Inference algorithms (rejection sampling and Metropolis Hastings). Operational semantics and inference by evaluation. Time permitting, an overview of differentiable programming.

91258 - Natural Language Processing

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student masters the basic algorithms, tools and systems for the management, processing and analysis of natural language. The student is able to design and develop simple systems oriented to real-world natural language processing applications such as those requiring segmentation and classification of elements of text.

Course contents

Topics

Introduction to NLP

Elements of computational linguistics

Basic text processing

Language modeling with n-grams

Part-of-speech tagging

Constituency and dependency parsing

Text classification using linear models

Vector semantics and embeddings

Deep learning methods for NLP

Typical NLP applications, e.g., sentiment analysis, question answering, machine translation, and dialog systems

Advanced topics and applications, e.g., argument mining, legal analytics, neural-symbolic architectures

Prerequisites

The course is intended for 2nd-year students of the Masters' Degree in Artificial Intelligence, who have a solid background in foundational subjects and technological tools for tackling modern Natural Language Processing. There are no formal prerequisites in terms of exams taken, but students are expected to be familiar with machine learning and deep learning concepts, methodologies and tools, AI & KR frameworks and methods, statistical and mathematical methods for AI, and the Python programming language.

91257 - Ethics in Artificial Intelligence

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands the main ethical aspects related to the creation and use of artificial intelligence systems. The student is able to reason critically about the main machine ethics issues connected with the moral behavior of AI systems endowed with autonomous decision capabilities and is able to evaluate the impact of the ethical and legal aspects connected with AI systems.

Course contents

The course aims to provide a thorough overview of the many ethical and social issues raised by information and communication technologies (ICTs), with particular attention to Artificial Intelligence and its impact on society and individuals.

Students will learn about the most compelling social and ethical challenges posed by ICTs and how to approach them in a critical way. Conceptual analysis will be supported by discussion of practical case studies.

Erasmus students follow the same programme of regular students.

Course Contents:

Introduction to Ethics

Information technologies, purposeful behavior and intelligence;

Singularity and Superintelligence;

Artificial Agency, Free Will, Consciousness;

Artificial Agents and Responsibility;

Machine Ethics;

AI Ethics and Roboethics;

Machine Learning, Big Data and the issues of Bias and Discrimination;

Anthropomorphism, Human-Computer/Robot Interaction (HCI, HRI), and Human Dignity; AI and Trust;

Human in the loop, Security, and Accountability;

AI Explainability and Transparency;

Assessing AI use cases. Socio-Technical Scenarios

Assessing AI use cases. Ethical tensions, Trade offs.

Possible use cases:

Military robots and Autonomous Weapon Systems;

Self-driving cars;

Expert systems: COMPAS, Watson, ...;

Aviation and Air Traffic Management;

Machine artistic creativity (TheNextRembrandt, Obvious Art, Shimon, ...);

AI Ethics in Healthcare

Legal relevance of AI Ethics

Ethical Business

AI and privacy

Modelling of legal and ethical norms and reasoning (with exercises)

Rea

91259 - Architecture and Platforms for Artificial Intelligence

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has a deep understanding of the requirements of machine-learning workloads for computing systems, has an understanding of the main architectures for accelerating machine learning workloads and heterogeneous architectures for embedded machine learning, and of the most popular platforms made available by cloud providers to specifically

support machine/deep learning applications.

Course contents

Module 1:

Introduction to parallel programming.

Parallel programming patterns: embarrassingly parallel, decomposition, master/worker, scan, reduce, ...

Shared-Memory programming with OpenMP.

OpenMP programming model: the “omp parallel” construct, scoping constructs, other work-sharing constructs.

Some examples of applications.

Module 2:

From ML to DNNs - a computational perspective

Introduction to key computational kernels (dot-product, matrix multiply...)

Inference vs training - workload analysis characterization

The NN computational zoo: DNNs, CNNs, RNNs, GNNs, Attention-based Networks

Running ML workloads on programmable processors

recap of processor instruction set architecture (ISA) with focus on data processing

improving processor ISAs for ML: RISC-V and ARM use cases

fundamentals of parallel processor architecture and parallelization of ML workloads

Algorithmic optimizations for ML

Key bottlenecks taxonomy of optimization techniques

Algorithmic techniques: Strassen, Winograd, FFT

Topology optimization: efficient NN models - depthwise convolutions, inverse bottleneck, introduction to Neural Architectural Search

Re

91261 - Artificial Intelligence in Industry

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has a deep knowledge of industrial applications that benefit from the use of machine learning, optimization and simulation. The student has a domain-specific knowledge of practical use cases discussed in collaboration with industrial experts in a variety of domains such as manufacturing, automotive and multi-media.

Course contents

The course is primarily delivered as a series of simplified industrial use cases. The goal is provide examples of challenges that typically arise when solving industrial problems.

Use cases may cover topics such as:

Anomaly detection

Remaining Useful Life (RUL) estimation

RUL based maintenance policies

Resource management planning

Recommendation systems with fairness constraints

Power network management problems

Epidemic control

Production planning

The course will emphasize the ability to view problems in their entirety and adapt to their peculiarities. This will frequently require to combine heterogeneous solution techniques, using integration schemes both simple and advanced.

The employed method will include:

Mathematical modeling of industrial problems

Predictive and diagnostic models for time series

Combinatorial Optimization

Integration methods for Probabilistic Models and Machine Learning

Integration methods for constraints and Machine Learning

Integration methods for combinatorial optimization and Machine Learning

The course will include seminars on real-world use cases, from industry experts.

The course contents may (and typically will) be subject to changes, so as to adapt to some degree to the interests and characteristics of the attending students.

91262 - Data Mining, Text Mining and Big Data Analytics

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands how a possibly very large set of data can be analyzed to derive strategic information and to address "data-driven" decisions. The student has a knowledge of the main data-mining tasks such as data selection, data transformation, analysis and interpretation, with specific reference to unstructured text data, and with the issues related to analysis in "big data" environments.

Course contents

Presentation slides of the course

Module 1 - Data Mining (Claudio Sartori)

See 75194 - DATA MINING M Module 2 only

Module 2 - Text Mining (Gianluca Moro)

The text mining module focuses on several forms of knowledge discovery from large corpora of unstructured text by also means on the most important set of the state-of-the-art deep learning advancements.

These advances are more and more fundamental to cope with challenging and highly business demanded natural language processing and natural language understanding tasks, such as text representation learning, semantic parsing, similarity search and efficient large scale neural information retrieval, self-supervised methods for unsupervised data, deep metric learning in NLP, explaining behaviours and phenomenon of interest from texts, knowledge graph learning from documents, open domain question answering, generative summarisation, chatbot or digital assistant, multi-modal information retrieval of texts and images, etc, including some traditional tasks such as sentiment analysis and opinion mining, text indexing, classification and topic analysis.

The module - after a summary introduction to non-contextual language models based on word embeddings, such as GloVe and word2vec, and to memory based neural networks particularly effective for textual data, such as recurrent neural networks like LSTM, GRU and BiLSTM - deals with the following new state-of-the-art deep learning advances:

dimensionality reduction and feature selection methods for textual data; lab on a novel latent semantic analysis method to discover underlying explanations of phenomena; lab: discovering explanations of aircraft accidents from unstructured flight textual reports

the best SOTA language models transformers for each prominent business area that have largely outperformed BERT with its quadratic attention complexity (published in 2018)

representation learning with deep metric learning; lab: neural self-supervised information retrieval based on deep metric learning; transfer learning paradigm to exploit and fine tune existing models in target domains semantically different from their training source domains and particularly useful to overcome the lack of labeled data in the target domain

multi-modal learning; lab: search engines development for large fashion corpora of texts and images with quadratic and linear attention

Deep learning methods for low-resource regimes (i.e. models and training sets order of magnitudes smaller than SOTA largest solutions); extractive and abstractive text summarization of long docs in low-resource settings; lab: achieving SOTA summarization of legal cases with training on a few dozen instances and small models

Learning knowledge graphs from unstructured text with Graph Neural Networks and graph attention methods; lab: extracting n-ary relations from medical literature for modeling & explaining biological processes

Memory-based deep neural networks for text mining, new retrieval-memory-based transformers, question answering and chatbot

The module introduces new published SOTA results in top AI conferences, we also achieved on some of these topics.

Module 3 - Big Data Analytics (Stefano Lodi)

Maps and reductions in parallel programming. The MapReduce programming model.

The Hadoop implementation of MapReduce.

The Python API to the Spark system and examples of parallel programs.

The Machine Learning

Library (MLlib) of Spark.

Readin

91266 - Machine Learning for Computer Vision

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student masters the most popular modern machine-learning approaches to computer-vision tasks, with particular reference to specialized deep-learning architectures. The student has both a theoretical understanding and the necessary practical skills required to develop state-of-the-art image and video analysis systems for real-world applications.

Course contents

The course is held in the first semester, from September to December.

Topics:

Image classification/recognition: limits of hand-crafted methods; brief review of machine learning and data-driven methods; bag-of-words models applied to dense descriptors; introduction to Neural Networks (NNs) and Convolutional NNs; AlexNet and the deep learning revolution. Advanced and state-of-the-art architectures for image classification. Hands-on sessions on image classification: training from scratch and transfer learning.

Intro to PyTorch and practical aspects of model training. Discussion of notebooks provided by the instructor to get to know PyTorch and learn how to train CNNs for image classification; discussion on the impact of several design and optimization choices.

Object detection. Introduction to ensemble learning via boosting. The Viola-Jones detector and its applications. Specialized NN architectures for object detection. Two-stages, one-stage, and anchor-free detectors. RoI Pooling operator, Feature Pyramid Networks. Imbalanced learning and the focal loss. Hands-on session on training a state-of-the-art network on a small dataset.

Dense prediction problems: semantic/instance segmentation and depth from mono/stereo. Ensemble learning via bagging and random forests. The algorithm behind the Kinect body part segmentation. Fully Convolutional Networks. Transposed and dilated convolutions. RoI Align operator. Specialized NN architectures for semantic, instance, and panoptic segmentation. Deep networks for depth estimation: DispNet, GCNet, Monodepth.

Metric learning. Deep metric learning and its applications to face recognition/identification and beyond. Locally connected layers. Contrastive and triplet loss. Hands-on session on face recognition.

Attention and vision transformers. Image classification and object detection architectures based on Transformers.

Prerequisites:

Basic knowledge of the Python programming language

Basic knowledge of computer vision and image processing: image formation, image digitization, camera modelling and calibration, basic image manipulation and processing, local image features, basics of stereo vision.

If you attended Computer Vision and Image Processing, you already fulfill the first and second prerequisites. If you didn't, you can ask Prof. Di Stefano access to his slides and lab sessions.

Nice-to-have knowledge, but not prerequisites

Basic knowledge of machine learning: supervised versus unsupervised learning, classification versus regression, underfitting and overfitting, regularization; data split in training, validation and test sets; hyper-parameters and cross-validation.

Basic knowledge of PyTorch: a good intro is available at <https://pytorch.org/tutorials/beginner/basics/intro.html>

We will briefly revise the basic ML knowledge at the beginning of the course, and the instructor will provide notebooks to introduce PyTorch and discuss them in class.

R

91264 - User Experience Design

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to design, implement and evaluate the quality of the interaction with software systems with respect to the dimensions of practicality, experience, affection, meaning and value that they may have on the target audience. The student is able to focus the functional analysis of the software on the characteristics and needs of the target audience, to drive the development process so as to guarantee a constant connection between the technical and implementation features and the expectation of the audience, and to evaluate whether and according to which metrics a software system satisfies these expectations.

Course contents

Teaching will be mostly in English (some explanations will be repeated in Italian if needed and when asked). The exam will be allowed in both Italian and English. The course content is divided in distinct parts:

Background The evolution of the discipline from Human Computer Interaction to User Experience Design. A description of its scope: the human, the computer, and their interaction.

Usability analysis and design A systematic discussion of the techniques and standards for the management of the process of user experience design, with particular attention to the phases of usability analysis (with and without the participation of users) and the user- and goal-oriented usability design methodologies.

Guidelines, patterns and methods for usability design A discussion, with historical aspects, of the framework on which the concrete aspects of usability design is based. we will also give strong attention to the problem of usability for web applications and mobile apps.

81683 - Internet of Things

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands the architectures and methodologies at the basis

of the Internet of Things (IoT) following a data centric approach. The student has a deep understanding of network protocols, integration of mobile and pervasive end-devices, middleware platforms for M2M-based IoT systems, edge/distributed/gateway computing principles, new services, service platforms and innovative application scenarios.

Course contents

The classes will be in English (previous years' editions were in Italian). The class will introduce the enabling technologies, protocols, software architectures and applications for the development of the emerging Internet of Things (IoT) paradigm. After a short introduction to pervasive computing and the emerging fields of applications (e.g. Industry 4.0, domotics, intelligent transportation systems, wearable devices, etc.) the class will provide an exhaustive illustration of the enabling components of IoT systems. These will include wireless communications among end-devices and towards infrastructures, data processing, languages for the development of applications and prototypes (Arduino, STM32 Nucleo, etc.). Specifically, the components of typical IoT system will be illustrated by following a data-oriented approach: from data generation on behalf of sensors to data transmission in WSN, WPAN, WLAN, up to cloud streaming and storages, and including the data processing and integration in complex software systems.

A short summary of the course contents includes:

Introduction to IoT systems: definitions, applications, enabling technologies

IoT components: from sensors to gateway.

Sensors and actuators, sketchy illustration of design principles of electronic circuit boards, sensing and data acquisition policies.

Wireless communication and coding technologies for IoT

Machine-to-machine (M2M) communication among devices: wireless technologies for WPAN (BLE, IEEE 802.15.4, Z-Wave, etc.), WLAN and WSN (LoRA, Dash7, Spirit, etc.)

Network architectures and routing (6LoWPAN, RPL)

Microcontrollers and System on Chip (SoC) prototypes' programming (e.g. Arduino, STM32 Nucleo, ESP8266, etc.)

IoT components: from gateway to cloud.

Data acquisition protocols: XMPP, CoAP, MQTT, AMQP, WebSocket, etc.

Architectures for data storage and processing: cloud/fog/edge computing

IoT platforms: AWS IoT, Watson IoT, ThingSpeak, etc.

IoT & Big-data

IoT components: from cloud to applications.

Web of Things and Semantic Web 3.0 Standard for M2M and IoT

Machine learning principles (classification, clustering, regression, anomaly detection, etc.) and applications for sensor data analysis.

Frameworks for integrated development of IoT applications: AllJoyn, Google Thing, Apple HomeKit, etc.

IoT Open Issues and bottlenecks.

Privacy and Security management;

Energy efficiency for end-devices and power saving techniques;

Scalability.

R

91265 - Knowledge Engineering

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student knows some (semi-)automated methods for joint interpretation of data and content as sources of knowledge. The student masters the basics of knowledge extraction, engineering, and linking, making data suitable to machine querying and automated reasoning, typically on decentralized platforms such as the Web.

Course contents

Knowledge graphs and ontologies

OWL, RDF, SPARQL

Ontology Design methodologies, focus on Extreme Design

Intensional and Extensional modelling

Ontology design patterns

Applied reasoning

Knowledge extraction from text

Ontology evaluation and quality

91267 - Multi-Agent Systems

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands the conceptual and technical foundation of multi-agent systems, and is able to model complex computational and socio-technical systems in terms of agent-oriented abstractions; build complex computational and socio-technical systems using agent-oriented technologies and methodologies; design and engineer intelligent systems as multi-agent systems, by integrating techniques and methods from artificial intelligence in an effective and methodologically-sound way.

Course contents

- Autonomy in natural and artificial systems
- Agent-oriented middleware for (intelligent) distributed systems

ACL & Protocols; FIPA; JADE

- Models for intelligent agents and MAS

Game Theory; Intentional agents, BDI; A&A

- Technologies for intelligent agents and MAS

Jason, ASTRA, CArTAgO

- Agreement technologies

Semantics; Norms; Organisations & Institutions; Argumentation & Negotiation; Trust & Reputation

Coordination Models & Technologies; Linda & TuCSoN

- Agent-oriented Software Engineering

GAIA & SODA

- Other contents

Simulation and MAS; NetLogo, Alchemist

Self-organisation, Swarm intelligence and MAS

Logics for MAS

91252 - LABORATORY OF BUSINESS PLAN

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to account for and compare different theories that describe the starting up and organization of innovative ventures and understands the information

and the analyses needed to perform a feasibility analysis and write a draft business plan for an innovative venture. The student is able to collect and analyze relevant information for that purpose, and to communicate a feasibility analysis/business plan orally and in writing.

Course contents

The course focuses on the development of entrepreneurial behaviours in a variety of settings. In particular, it presents the fundamentals of developing an effective business plan, teaching the basic principles of developing plans for the sales, marketing, product or service operations, financials and management of a new venture. It assumes some prior knowledge of the subject area and will equip students with much of the knowledge required to launch and manage a high potential new venture.

91284 - Project Work in Data Mining, Text Mining and Big Data Analytics

PDF, ADOC.

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Data Mining, Text Mining And Big Data Analytics in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

See 91262 - Data Mining, Text Mining and Big Data Analytics [<http://www.unibo.it/en/teaching/course-unit-catalogue/course-unit/2020/446610>]

91277 - Project Work in Image Processing and Computer Vision

PDF, ADOC.

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Image Processing And Computer Vision in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

See "91254 - Image Processing and Computer Vision"

91287 - Project Work in Knowledge Engineering

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Knowledge Engineering in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

Students will develop a project, under the supervision of the teacher. The project can be both individual or developed in group (of max three students). The topic and the goals of the project will be identified and agreed case by case but they will be relevant for the knowledge engineering area. A non-complete list of application areas is the following:

knowledge graph creation from text

ontology engineering

question answering based on knowledge graphs

knowledge graph visualisation

knowledge graph embeddings

ontology / knowledge graph matching

development of Protégé plugins for ontology design

development of ontology and knowledge graph visualisation tools

A non-exhaustive list of available projects is the following:

Learning from and integrating common sense resources to enable common sense reasoning

Implementation of a method for question answering over knowledge graphs

Developing a tool for supporting semi-automatic ontology testing

Developing a framework for automatic ontology evaluation

Implementation of a visualisation tool for ontology and knowledge graph summarisation

Knowledge extraction from text for populating knowledge graphs

Knowledge extraction from audio and/or music notations for populating knowledge graphs

Music recommendation based on knowledge graphs

Populating and refining knowledge graphs with games

91273 - Project Work in Machine Learning

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Machine Learning in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

Three possible tracks:

A scientific paper is considered as the base for the development of a new data mining algorithm. The software engineering best practices have to be used to include the project into an open-source framework.

An innovative software/tool is chosen for experiments and tests on given datasets.

The professor assigns a dataset and a mining task to be developed on it, or the students proposes a dataset that can be accepted by the professor.

If you follow track 3 do the following:

visit the Kaggle website;

choose three competitions of your interest, either active or closed; the requirements are to have not less than some thousands of data units and not less than two data files of different structure/content type to be considered (train and test data count only once)

send an email to the teacher with the subject "[Project work] Kaggle competitions choice" and containing the link to the three competitions of your interest

wait for the answer and, when the competition is assigned, do as specified on the "Assessment methods section below"

91288 - Project Work in Machine Learning for Computer Vision

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Machine Learning For Computer Vision in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

See "91266 - Machine Learning for Computer Vision" contents.

93103 - Project Work in Multimedia Data Management M

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Multimedia data management M in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

A scientific paper is usually considered as the base for the development of a new multimedia data management algorithm and/or infrastructure and/or service.

91278 - Project Work in Natural Language Processing

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Natural Language Processing in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

This teaching activity is about developing a project in natural language processing or natural language understanding, following the course 91258 - Natural Language Processing. Upon request, the teaching staff can offer suggestions and help in defining the topic and objective of the project work. The student is, however, encouraged to take the initiative by reviewing the relevant literature and proposing a project work, for instance, about:

exploring a foundational aspect;

exploring NLP methods for low-resource languages or languages other than English;

working on a challenging application;

working on a research topic;

preparing a literature survey on a research topic;

significantly expanding a project developed within the NLP course.

The proposal has to be approved by the teaching staff before the student starts working on the project. Working as a group is encouraged.

Particularly innovative projects may receive extra support aimed to a possible submission of the work to a relevant scientific event.

81799 - Project Management and Soft Skills

M

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to acquire the necessary skills to be a successful engineer, with a particular focus on non-technical (also known as soft) skills such as inductive/deductive/analytical/design reasoning, communication skills and employability. The student is able to use the basic tools of project management, including Gantt and Pert diagrams, work package breakdown, project cost estimation, and relevant aspects.

Course contents

The syllabus is reported below. More information on the course approach can be found on www.robertoverdone.org.

I. Introduction (approx 3 hours).

II. A Pragmatic Approach to Project Management (approx 9 hours).

- Projects, Roles, WBS
- Gantt and PERT Diagrams
- Cost Assessment
- Submitting a Project Proposal

III. Soft Skills and Aptitudes (approx 12 hours).

- Introduction to Cognitive Processes

- Soft Skills at Time of Industry 4.0
- Success and Talent
- Aptitudes
- Team Work
- Leadership
- Learning Styles

IV. Communication Skills (approx 9 hours).

- Fundamentals of Human Communication
- Persuasion
- Inter-Personal Communication
- Presentation Skills

Additionally, the course will include a Madness session where all students will introduce themselves in one minute, a CV contest, and, possibly, one or two seminars from professionals.

91276 - Project Work in Combinatorial Optimization

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Combinatorial Optimization in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

This teaching activity is about developing a project in combinatorial optimization, following the course 91762 - Combinatorial Decision Making and Optimization. Upon request, the teaching staff can offer suggestions and help in defining the topic and objective of the project work. The student is, however, encouraged to take the initiative by reviewing the relevant literature and proposing a project work, for instance, about:

hybridization with machine learning;

a challenging application;

a research topic;

a literature survey on a research topic.

The proposal has to be approved by the teaching staff before the student starts working on the

project. Working as a group is a possibility depending on the project complexity.

91270 - Project Work in Fundamentals of Artificial Intelligence and Knowledge Representation

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Fundamentals of Artificial Intelligence and Knowledge Representation in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

To be discussed with the professor

91275 - Project Work in Languages and Algorithms for Artificial Intelligence

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to apply the knowledge acquired in Languages and algorithms for Artificial Intelligence in order to carry out autonomously a project focusing on a topic agreed upon with the teacher.

Course contents

The contents of the project will be agreed with the teacher in charge of the course.

91259 - Architecture and Platforms for Artificial Intelligence

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has a deep understanding of the requirements of machine-learning workloads for computing systems, has an understanding of the main architectures for

accelerating machine learning workloads and heterogeneous architectures for embedded machine learning, and of the most popular platforms made available by cloud providers to specifically support machine/deep learning applications.

Course contents

Module 1:

Introduction to parallel programming.

Parallel programming patterns: embarrassingly parallel, decomposition, master/worker, scan, reduce, ...

Shared-Memory programming with OpenMP.

OpenMP programming model: the “omp parallel” construct, scoping constructs, other work-sharing constructs.

Some examples of applications.

Module 2:

From ML to DNNs - a computational perspective

Introduction to key computational kernels (dot-product, matrix multiply...)

Inference vs training - workload analysis characterization

The NN computational zoo: DNNs, CNNs, RNNs, GNNs, Attention-based Networks

Running ML workloads on programmable processors

recap of processor instruction set architecture (ISA) with focus on data processing

improving processor ISAs for ML: RISC-V and ARM use cases

fundamentals of parallel processor architecture and parallelization of ML workloads

Algorithmic optimizations for ML

Key bottlenecks taxonomy of optimization techniques

Algorithmic techniques: Strassen, Winograd, FFT

Topology optimization: efficient NN models - depthwise convolutions, inverse bottleneck, introduction to Neural Architectural Search

Re

91261 - Artificial Intelligence in Industry

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has a deep knowledge of industrial applications that benefit from the use of machine learning, optimization and simulation. The student has a domain-specific knowledge of practical use cases discussed in collaboration with industrial experts in a variety of domains such as manufacturing, automotive and multi-media.

Course contents

The course is primarily delivered as a series of simplified industrial use cases. The goal is provide examples of challenges that typically arise when solving industrial problems.

Use cases may cover topics such as:

Anomaly detection

Remaining Useful Life (RUL) estimation

RUL based maintenance policies

Resource management planning

Recommendation systems with fairness constraints

Power network management problems

Epidemic control

Production planning

The course will emphasize the ability to view problems in their entirety and adapt to their peculiarities. This will frequently require to combine heterogeneous solution techniques, using integration schemes both simple and advanced.

The employed method will include:

Mathematical modeling of industrial problems

Predictive and diagnostic models for time series

Combinatorial Optimization

Integration methods for Probabilistic Models and Machine Learning

Integration methods for constraints and Machine Learning

Integration methods for combinatorial optimization and Machine Learning

The course will include seminars on real-world use cases, from industry experts.

The course contents may (and typically will) be subject to changes, so as to adapt to some degree to the interests and characteristics of the attending students.

90748 - Blockchain and Cryptocurrencies

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student knows the relevant themes related to blockchain technologies, cryptocurrencies, smart contracts and novel applications that can be built over the blockchain. The student is able to develop simple smart contracts that can be deployed on a blockchain.

Course contents

Bitcoin and novel cryptocurrencies gathered momentum in the last months. More and more investors look with interest at these technologies, while others label them as a dangerous speculative bubble. The truth is that the blockchain, and the alternative implementations of a distributed ledger, represent very interesting technologies, that can be exploited to build novel distributed applications. The underlying building blocks are related to many concepts and research areas of computer science in general. This course will illustrate the main principles and conceptual foundations of the blockchain and the Bitcoin network.

Program

Introduction to peer-to-peer systems

Overlay topologies and decentralization

Introduction to Crypto and Cryptocurrencies

The blockchain: how to achieve decentralization

Transactions and transaction scripting languages

Mining

Attacks to the blockchain

Anonymity

Smart contracts

81943 - Complex Systems & Network Science

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has basic notions of complexity and network sciences and is able to identify, formulate, model and analyze new problems that arise in modern computing systems that can be studied using them.

Course contents

Description: Modern information systems and services often rely on large numbers of independent interacting components to provide their functions. Under certain conditions, the behavior that results from these interactions can be unexpected and surprising. Complexity Science is an interdisciplinary field for studying global behaviors resulting from many simple local interactions in an effort to characterize and control them. Networks allow us to formalize the structure of interactions. They play a central role in the transmission of information, transportation of goods, spread of diseases, diffusion of innovation, formation of opinions and adoption of new technologies. Network Science is an interdisciplinary field for studying the interconnectedness of modern life by exploring fundamental properties that govern the structure and dynamic evolution of networks.

Contents: Complex systems: definitions, methodologies; Dynamical systems, Nonlinear dynamics; Chaos, Bifurcations and Feigenbaum constant, Predictability, Randomness and Chaos; Models of complex systems, Cellular automata, Wolfram's classification, Game of life; Autonomous agents, Flocking, Schooling, Synchronization, Formation creation; Cooperation and Competition, Game theory basics, Nash equilibrium; Game theory: Prisoner's Dilemma, Coordination games, Mixed strategy games; Adaptation, Evolution, Genetic algorithms, Evolutionary games; Network Science: Definitions and examples; Graph theory, Basic concepts and definitions; Diameter, Path length, Clustering, Centrality metrics; Structure of real networks, Degree distribution, Power-laws, Popularity; Models of network formation; The Erdos-Renyi random model; Clustered models; Models of network growth, Preferential attachment; Small-world networks, Network navigation; Peer-to-peer systems and overlay networks; Structured overlays, DHTs, Key-based routing, Chord; Distributed network formation: Newscast, Cyclon, T-Man; Processes on networks: Aggregation; Rational dynamics: Cooperation in selfish environments, Homophily, Segregation; Diffusion, Percolation, Tipping points, Peer-effects, Cascades.

Prerequisites: Basic notions of computer system architecture, computer networks, operating systems, and probability theory.

93467 - CRYPTOGRAPHY

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student knows the basics of modern cryptography and some techniques for the analysis of security protocols. He or she can verify the absence of security bugs on simple protocols.

Course contents

This course deals with some fundamental notions on modern cryptography and some techniques for the formal verification of cryptographic protocols. In addition to many new concepts, we discuss some cryptographic techniques already seen in previous courses, studying the most interesting properties. The main emphasis of this course is on the theoretical aspects of cryptography: it attempts to give a clear and accurate account on security, in order to give students

the tools to assess (and not just to use) cryptographic techniques . In other words, we focus a lot on the "why" maybe putting the "how" behind the scenes. More specifically, we seek to answer questions such as the following: when a cryptographic technique can be consider safe? Is it possible to formally prove that a cryptographic technique is secure ? Is it possible to analyze the security of a cryptographic protocol in an automatic or semi-automatic way?

The addressed topics will be the following ones:

Perfect Security

Private-key Encryption Schemes and Pseudorandomness

Message Authenticating Codes.

Block Ciphers: DES and AES.

One-Way Functions and Pseudorandomness

Assumptions from Number Theory and Algebra

The Public-Key Revolution

Public-key Encryption Schemes

The Symbolic Model.

93470 - CYBERSECURITY

[PDF](#), [ADOC](#).

Learning outcomes

The student will learn: historical perspective and motivations for cyber-security, principles of data security, principles of software security, component and connection security, network security.

Course contents

Scenarios and strategies

A top-down approach to cybersecurity

Information security

Cyber threats to industry

Analysis of cyber threats

Qualitative and quantitative evaluation

Management of cyber risks: methodologies and technologies

Assessment and validation

Machine Learning for cybersecurity

Disruptive solutions

Competences and jobs of cybersecurity

91262 - Data Mining, Text Mining and Big Data Analytics

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands how a possibly very large set of data can be analyzed to derive strategic information and to address "data-driven" decisions. The student has a knowledge of the main data-mining tasks such as data selection, data transformation, analysis and interpretation, with specific reference to unstructured text data, and with the issues related to analysis in "big data" environments.

Course contents

Presentation slides of the course

Module 1 - Data Mining (Claudio Sartori)

See 75194 - DATA MINING M Module 2 only

Module 2 - Text Mining (Gianluca Moro)

The text mining module focuses on several forms of knowledge discovery from large corpora of unstructured text by also means on the most important set of the state-of-the-art deep learning advancements.

These advances are more and more fundamental to cope with challenging and highly business demanded natural language processing and natural language understanding tasks, such as text representation learning, semantic parsing, similarity search and efficient large scale neural information retrieval, self-supervised methods for unsupervised data, deep metric learning in NLP, explaining behaviours and phenomenon of interest from texts, knowledge graph learning from documents, open domain question answering, generative summarisation, chatbot or digital assistant, multi-modal information retrieval of texts and images, etc, including some traditional tasks such as sentiment analysis and opinion mining, text indexing, classification and topic analysis.

The module - after a summary introduction to non-contextual language models based on word embeddings, such as GloVe and word2vec, and to memory based neural networks particularly effective for textual data, such as recurrent neural networks like LSTM, GRU and BiLSTM - deals

with the following new state-of-the-art deep learning advances:

dimensionality reduction and feature selection methods for textual data; lab on a novel latent semantic analysis method to discover underlying explanations of phenomena; lab: discovering explanations of aircraft accidents from unstructured flight textual reports

the best SOTA language models transformers for each prominent business area that have largely outperformed BERT with its quadratic attention complexity (published in 2018)

representation learning with deep metric learning; lab: neural self-supervised information retrieval based on deep metric learning; transfer learning paradigm to exploit and fine tune existing models in target domains semantically different from their training source domains and particularly useful to overcome the lack of labeled data in the target domain

multi-modal learning; lab: search engines development for large fashion corpora of texts and images with quadratic and linear attention

Deep learning methods for low-resource regimes (i.e. models and training sets order of magnitudes smaller than SOTA largest solutions); extractive and abstractive text summarization of long docs in low-resource settings; lab: achieving SOTA summarization of legal cases with training on a few dozen instances and small models

Learning knowledge graphs from unstructured text with Graph Neural Networks and graph attention methods; lab: extracting n-ary relations from medical literature for modeling & explaining biological processes

Memory-based deep neural networks for text mining, new retrieval-memory-based transformers, question answering and chatbot

The module introduces new published SOTA results in top AI conferences, we also achieved on some of these topics.

Module 3 - Big Data Analytics (Stefano Lodi)

Maps and reductions in parallel programming. The MapReduce programming model.

The Hadoop implementation of MapReduce.

The Python API to the Spark system and examples of parallel programs.

The Machine Learning

Library (MLlib) of Spark.

Readin

91266 - Machine Learning for Computer Vision

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student masters the most popular modern machine-learning approaches to computer-vision tasks, with particular reference to specialized deep-learning architectures. The student has both a theoretical understanding and the necessary practical skills required to develop state-of-the-art image and video analysis systems for real-world applications.

Course contents

The course is held in the first semester, from September to December.

Topics:

Image classification/recognition: limits of hand-crafted methods; brief review of machine learning and data-driven methods; bag-of-words models applied to dense descriptors; introduction to Neural Networks (NNs) and Convolutional NNs; AlexNet and the deep learning revolution. Advanced and state-of-the-art architectures for image classification. Hands-on sessions on image classification: training from scratch and transfer learning.

Intro to PyTorch and practical aspects of model training. Discussion of notebooks provided by the instructor to get to know PyTorch and learn how to train CNNs for image classification; discussion on the impact of several design and optimization choices.

Object detection. Introduction to ensemble learning via boosting. The Viola-Jones detector and its applications. Specialized NN architectures for object detection. Two-stages, one-stage, and anchor-free detectors. RoI Pooling operator, Feature Pyramid Networks. Imbalanced learning and the focal loss. Hands-on session on training a state-of-the-art network on a small dataset.

Dense prediction problems: semantic/instance segmentation and depth from mono/stereo. Ensemble learning via bagging and random forests. The algorithm behind the Kinect body part segmentation. Fully Convolutional Networks. Transposed and dilated convolutions. RoI Align operator. Specialized NN architectures for semantic, instance, and panoptic segmentation. Deep networks for depth estimation: DispNet, GCNet, Monodepth.

Metric learning. Deep metric learning and its applications to face recognition/identification and beyond. Locally connected layers. Contrastive and triplet loss. Hands-on session on face recognition.

Attention and vision transformers. Image classification and object detection architectures based on Transformers.

Prerequisites:

Basic knowledge of the Python programming language

Basic knowledge of computer vision and image processing: image formation, image digitization, camera modelling and calibration, basic image manipulation and processing, local image features, basics of stereo vision.

If you attended Computer Vision and Image Processing, you already fulfill the first and second prerequisites. If you didn't, you can ask Prof. Di Stefano access to his slides and lab sessions.

Nice-to-have knowledge, but not prerequisites

Basic knowledge of machine learning: supervised versus unsupervised learning, classification versus regression, underfitting and overfitting, regularization; data split in training, validation and test sets; hyper-parameters and cross-validation.

Basic knowledge of PyTorch: a good intro is available at <https://pytorch.org/tutorials/beginner/basics/intro.html>

We will briefly revise the basic ML knowledge at the beginning of the course, and the instructor will provide notebooks to introduce PyTorch and discuss them in class.

R

90730 - SOCIAL NETWORK ANALYSIS

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, students gain knowledge on the Web as a socio-technical system involving specific processes, entities, and behaviours, using interdisciplinary methods that blend computer science, sociology, ethnography, economics, linguistics, etc. The students are able to analyse the Web phenomena similarly to typical objects from natural sciences, distinguishing between data and applications, agents from computationally generated profiles, and addressing the characteristics of networks of entities emerging from the informationl, physical, social, and conceptual spaces constituting the Web.

Course contents

The programme is the same for both attending and non-attending students:

Introduction to Network Analysis, gentle introduction to the field of network analysis and its usages in other fields of research (e.g., computer science, forensics, archeology, literature, history, science of religion, etc.).

Research Design and How to Read a (Network Analysis) Research Paper: introduction to the scientific publication process, elements of research papers (on network analysis), research design, analysis of research papers.

Mathematics of Networks: networks and their representation, types of networks, graph representations, paths and components, adjacency matrices and matrix representations, ways and modes, operations on Matrices.

Data Collection and Data Management: network questions, data collection and reliability, data formats and transformation, algorithms and software for network analysis and visualisation.

Measures and Metrics, Nodes: kinds of measures, multi-category nominal scales, ordinal and scalar measures, centrality, degree and other kinds of centrality (e.g. Google's PageRank), hubs and authorities, closeness and betweenness centrality, groups of nodes (cliques, cores, components and k-components), clustering and clustering coefficients, reciprocity and similarity, structural and other types of equivalence, homophily and types of assortative mixing.

Testing Hypotheses: the role of hypotheses in the scientific method, testing hypotheses in network analysis, permutation tests, dyadic hypotheses.

Measures and Metrics, Networks: small-world effects, degree distribution, power laws and scale-free networks, visualisation and properties of power-law distributions, local-clustering coefficient, cohesion, reciprocity, transitivity and the clustering coefficient, triad census, centralisation and core-periphery indices, centrality, random graphs, means on edges and degree, degree distribution, giant and small component(s), locally tree-like networks.

Network Visualisation: the importance of network visualisation, graph-layout algorithms, embedding node attributes, node filtering, visualising ego networks, embedding tie characteristics, tie strengths, visualising network change.

Handling Large Networks: reducing the size of the problem, eliminating edges, pruning nodes, divide and conquer, aggregation, sampling, small-world and scale-free networks.

91264 - User Experience Design

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to design, implement and evaluate the quality of the interaction with software systems with respect to the dimensions of practicality, experience, affection, meaning and value that they may have on the target audience. The student is able to focus the functional analysis of the software on the characteristics and needs of the target audience, to drive the development process so as to guarantee a constant connection between the technical and implementation features and the expectation of the audience, and to evaluate whether and according to which metrics a software system satisfies these expectations.

Course contents

Teaching will be mostly in English (some explanations will be repeated in Italian if needed and when asked). The exam will be allowed in both Italian and English. The course content is divided in distinct parts:

Background The evolution of the discipline from Human Computer Interaction to User Experience Design. A description of its scope: the human, the computer, and their interaction.

Usability analysis and design A systematic discussion of the techniques and standards for the management of the process of user experience design, with particular attention to the phases of usability analysis (with and without the participation of users) and the user- and goal-oriented usability design methodologies.

Guidelines, patterns and methods for usability design A discussion, with historical aspects, of the framework on which the concrete aspects of usability design is based. we will also give strong attention to the problem of usability for web applications and mobile apps.

92858 - Autonomous and Adaptive Systems M

[PDF](#), [ADOC](#).

Learning outcomes

At the end of this course student will have a solid understanding of the state of the art and the key conceptual and practical aspects of the design, implementation and evaluation of intelligent machines and autonomous systems that learn by interacting with their environment. The course also takes into consideration ethical, societal and philosophical aspects related to these technologies.

Course contents

Introduction to the design of adaptive and autonomous systems, intelligent agents and intelligent machines.

Review of Machine Learning/Deep Learning fundamentals.

Introduction to Reinforcement Learning (RL): multi-armed bandits, Montecarlo methods, tabular methods, approximation function methods, and deep RL.

Applications of RL to games.

Applications of RL to classic control theory problems and robotics.

Applications of RL to networked and distributed systems design.

RL and cognitive sciences&neuroscience.

Intelligent machines and creativity: Generative Deep Learning, AI and the Arts.

Introduction to algorithmic game theory for autonomous systems: cooperation and coordination, social dilemmas, and Multi-Agent Reinforcement Learning.

Bio-inspired adaptive systems.

Implementing autonomous systems: mobile robots, driverless cars and autonomous transportation

systems.

Ethical implications of AI and autonomous systems.

Machine intelligence, super-intelligence, self-awareness, controllability and the future.

The course will include labs in which we will discuss implementation oriented aspects of the techniques and methodologies presented during the course.

B0708 - COMPUTATIONAL ETHICS

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is familiar with the main meta-ethical issues, normative ethical models, and theoretical contributions of cognitive science related to the translatability of moral judgments in computational terms. At the end of the course, the student is able to critically reason about alternative theoretical proposals for algorithmic formulation of ethics, as well as about theories that deny the possibility of computational ethics.

Course contents

The course aims to provide

an introduction to computational ethics;

an introduction, from a critical perspective, to the ethics of artificial intelligence.

There are no prerequisites. The course is designed for students who have never studied philosophy.

92995 - DISTRIBUTED AUTONOMOUS SYSTEMS M

[PDF](#), [ADOC](#).

Learning outcomes

The course focuses on a novel class of autonomous control and processing systems known as cyber-physical networks or distributed autonomous systems. These multi-agent systems consist of many cyber-physical agents that aim at performing global tasks via local communication and computation in a cooperative way. At the end of the course students will know how to design selected distributed control, learning and optimization methods to solve complex tasks in cooperative network systems. To bridge the gap between theory and application, laboratory activities will allow students to apply the proposed techniques to a number of application domains, including, e.g., planning and control of cooperative autonomous robots, machine learning and data analytics, decision making and optimal control of cooperative autonomous systems.

Course contents

Introduction to Distributed Autonomous Systems

New paradigms and applications of autonomous systems: smart cyber-physical network systems. Introduction to distributed systems: centralized versus distributed computing. Application domains of distributed autonomous systems: sensor networks, cooperative robots, social networks, decision systems for data analytics (e.g., recommender systems and localization), smart energy systems (e.g., microgrid control and synchronization in power networks). Key properties and main goals for distributed systems.

Modeling of distributed systems

Introduction to distributed systems models. Graph theory as a tool to model communication among network agents. Preliminaries on graph theory and examples. Distributed algorithms and distributed control laws. Introduction to Python programming with a focus on distributed computing and cooperative robotics via Robotic Operating System.

Basic distributed algorithms

Averaging protocols and linear consensus algorithms for multi-agent systems. Complex tasks (e.g., autonomous formation control, containment in leader-follower networks) based on linear consensus algorithms. Practical Python implementation of averaging and distributed control laws on case study examples borrowed from opinion dynamics and sensor networks.

Introduction to distributed optimization

Main problem set-ups and examples from estimation, machine learning, decision making, and control in cyber-physical networks.

Distributed decision making via consensus optimization

Consensus optimization algorithms based on averaging. Distributed gradient and gradient tracking methods. Practical Python implementation on case studies borrowed from distributed federated machine learning, e.g., logistic regression, support vector machine, training of neural networks for classification.

Cooperative robotics via distributed optimization

Constraint-coupled optimization algorithms based on decomposition schemes. Aggregative optimization. Practical Python implementation on case study examples borrowed from task allocation, formation control and surveillance in cooperative robotics. Introduction to the ChoiRbot toolbox for cooperative robotics.

Rea

B0265 - EXPERT SYSTEMS

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student acquires a comprehensive knowledge of methodologies and technologies for the development of expert systems. The student is able to design and develop simple expert systems

Course contents

Introduction to multidisciplinary research. Multidisciplinary goals and expert systems. Approaching complex multidisciplinary applications. Knowledge modelling as a tool for multidisciplinary research. Knowledge acquisition techniques. Pragmatic issues in Knowledge Acquisition. Exploiting hybrid AI technologies for building expert systems. (Symbolic) AI languages and hybrid artificial intelligence tools. Advanced problem solving methods for expert systems. Advanced methods for representing uncertain knowledge. Integrating Knowledge Representation and Reasoning and Machine Learning. Case studies: development of expert systems as laboratory tasks.

81683 - Internet of Things

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands the architectures and methodologies at the basis of the Internet of Things (IoT) following a data centric approach. The student has a deep understanding of network protocols, integration of mobile and pervasive end-devices, middleware platforms for M2M-based IoT systems, edge/distributed/gateway computing principles, new services, service platforms and innovative application scenarios.

Course contents

The classes will be in English (previous years' editions were in Italian). The class will introduce the enabling technologies, protocols, software architectures and applications for the development of the emerging Internet of Things (IoT) paradigm. After a short introduction to pervasive computing and the emerging fields of applications (e.g. Industry 4.0, domotics, intelligent transportation systems, wearable devices, etc.) the class will provide a exhaustive illustration of the enabling components of IoT systems. These will include wireless communications among end-devices and towards infrastructures, data processing, languages for the development of applications and prototypes (Arduino, STM32 Nucleo, etc.). Specifically, the components of typical IoT system will be illustrated by following a data-oriented approach: from data generation on behalf of sensors to data transmission in WSN, WPAN, WLAN, up to cloud streaming and storages, and including the data processing and integration in complex software systems.

A short summary of the course contents includes:

Introduction to IoT systems: definitions, applications, enabling technologies

IoT components: from sensors to gateway.

Sensors and actuators, sketchy illustration of design principles of electronic circuit boards, sensing and data acquisition policies.

Wireless communication and coding technologies for IoT

Machine-to-machine (M2M) communication among devices: wireless technologies for WPAN (BLE, IEEE 802.15.4, Z-Wave, etc.), WLAN and WSN (LoRA, Dash7, Spirit, etc.)

Network architectures and routing (6LoWPAN, RPL)

Microcontrollers and System on Chip (SoC) prototypes' programming (e.g. Arduino, STM32 Nucleo, ESP8266, etc.)

IoT components: from gateway to cloud.

Data acquisition protocols: XMPP, CoAP, MQTT, AMQP, WebSocket, etc.

Architectures for data storage and processing: cloud/fog/edge computing

IoT platforms: AWS IoT, Watson IoT, ThingSpeak, etc.

IoT & Big-data

IoT components: from cloud to applications.

Web of Things and Semantic Web 3.0 Standard for M2M and IoT

Machine learning principles (classification, clustering, regression, anomaly detection, etc.) and applications for sensor data analysis.

Frameworks for integrated development of IoT applications: AllJoyn, Google Thing, Apple HomeKit, etc.

IoT Open Issues and bottlenecks.

Privacy and Security management;

Energy efficiency for end-devices and power saving techniques;

Scalability.

R

91265 - Knowledge Engineering

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student knows some (semi-)automated methods for joint interpretation of data and content as sources of knowledge. The student masters the basics of knowledge

extraction, engineering, and linking, making data suitable to machine querying and automated reasoning, typically on decentralized platforms such as the Web.

Course contents

Knowledge graphs and ontologies

OWL, RDF, SPARQL

Ontology Design methodologies, focus on Extreme Design

Intensional and Extensional modelling

Ontology design patterns

Applied reasoning

Knowledge extraction from text

Ontology evaluation and quality

91267 - Multi-Agent Systems

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student understands the conceptual and technical foundation of multi-agent systems, and is able to model complex computational and socio-technical systems in terms of agent-oriented abstractions; build complex computational and socio-technical systems using agent-oriented technologies and methodologies; design and engineer intelligent systems as multi-agent systems, by integrating techniques and methods from artificial intelligence in an effective and methodologically-sound way.

Course contents

- Autonomy in natural and artificial systems
- Agent-oriented middleware for (intelligent) distributed systems

ACL & Protocols; FIPA; JADE

- Models for intelligent agents and MAS

Game Theory; Intentional agents, BDI; A&A

- Technologies for intelligent agents and MAS

Jason, ASTRA, CArTAgO

- Agreement technologies

Semantics; Norms; Organisations & Institutions; Argumentation & Negotiation; Trust & Reputation

Coordination Models & Technologies; Linda & TuCSoN

- Agent-oriented Software Engineering

GAIA & SODA

- Other contents

Simulation and MAS; NetLogo, Alchemist

Self-organisation, Swarm intelligence and MAS

Logics for MAS

77933 - Multimedia Data Management M

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student has the knowledge and skills required for an effective and efficient management of multimedia (MM) data, with particular attention to the problems of MM data representation, MM data retrieval models, and interaction paradigms between the user and the MM system both for purposes of data presentation and exploration. The student understands the architecture of traditional and advanced MM systems and services, search engines, social networks and recommendation systems.

Course contents

Basics on Multimedia Data Management

Multimedia data and content representations

MM data and applications

MM data coding

MM data content representation

How to find MM data of interest

Description models for complex MM objects

Similarity measures for MM data content

MM Data Base Management Systems

Efficient algorithms for MM data retrieval

MM query formulation paradigms

Sequential retrieval of MM data

Index-based retrieval of MM data

Automatic techniques for MM data semantic annotations

Browsing MM data collections

MM data presentation

User interfaces

Visualization paradigms

Dimensionality reduction techniques

Result accuracy, use cases and real applications

Quality of the results and relevance feedback techniques

Use cases and demos of some applications

Multimedia Data on the Web

Web search engines

Graph-based data: semantic Web and social networks

Web recommender systems

N.B. For students of the second cycle degree programmes (LM) in Artificial intelligence and Computer Science, the "Efficient algorithms for MM data retrieval" part of the program is not required.

81799 - Project Management and Soft Skills

M

[PDF](#), [ADOC](#).

Learning outcomes

At the end of the course, the student is able to acquire the necessary skills to be a successful engineer, with a particular focus on non-technical (also known as soft) skills such as inductive/deductive/analytical/design reasoning, communication skills and employability. The student is able to use the basic tools of project management, including Gantt and Pert diagrams, work package breakdown, project cost estimation, and relevant aspects.

Course contents

The syllabus is reported below. More information on the course approach can be found on www.robertoverdone.org.

I. Introduction (approx 3 hours).

II. A Pragmatic Approach to Project Management (approx 9 hours).

- Projects, Roles, WBS
- Gantt and PERT Diagrams
- Cost Assessment
- Submitting a Project Proposal

III. Soft Skills and Aptitudes (approx 12 hours).

- Introduction to Cognitive Processes
- Soft Skills at Time of Industry 4.0
- Success and Talent
- Aptitudes
- Team Work
- Leadership
- Learning Styles

IV. Communication Skills (approx 9 hours).

- Fundamentals of Human Communication
- Persuasion
- Inter-Personal Communication
- Presentation Skills

Additionally, the course will include a Madness session where all students will introduce themselves in one minute, a CV contest, and, possibly, one or two seminars from professionals.

97431 - SCALABLE AND RELIABLE SERVICES

M

[PDF](#), [ADOC](#).

Learning outcomes

Digital services have to meet various requirements in terms of performance, scalability and reliability. We analyze the whole lifecycle of high quality services, from design to deployment, testing and operation. In each phase, we aim to provide the students with the ability to gather and analyze data, to identify sources of outages and critical dependencies, and to avoid bottlenecks and

single points of failure. Finally, we expect students to be prepared to evaluate and discuss alternatives and justify investment in support of business services.

Course contents

Introduction

Cloud services

Cloud provider architectures

High-quality services: performance, scalability, reliability (and security)

Lifecycle of high-quality services

Part 1: Architectures for scalable and reliable services

Sources of outages and critical dependencies

Bottlenecks and single points of failure

Design: maturity assessment, architectural principles for reliability and scalability, AKF scale cube model

Migration to cloud

Critical data management: Strategy, Governance, Management

Implementation alternatives on cloud platforms and multi-clouds

Part 2: Service deployment

Secure and robust software production

Secure and robust DevOps

Labs on cloud platforms

Part 3: Tests

Benchmarking

Performance testing

Stress testing

Part 4: Scalable and reliable organizations

Roles and responsibilities

Leaders and managers

Putting all together

Capacity planning

Understanding and managing complexity

Managing changes

Part 5: Issues, incident and crises

Performance monitoring

Warning signs in operation

Incidents management

Crisis management