Denominazione dell'insegnamento	Numero di ore totali	Distribuzione durante il ciclo di dottorato	Descrizione del corso
Functional Programming for Big Data Processing	12	primo anno	Through the pervasive use of computers, smartphones and other digital objects, huge amounts of digital data are generated and collected. This data, commonly referred as Big Data, represents a challenge for the current storage, process, and analysis capabilities. To extract value from such data, novel architectures, programming models and frameworks have been developed in recent years for capturing and analyzing complex and/or high velocity data. The usage of high performance computers, such as Clouds and clusters, paired with parallel and distributed algorithms, are commonly used by data analysts to solve big data problems and obtain valuable information and knowledge in a reasonable time. This course introduces the most effective programming approaches for Big Data processing. In particular, the functional programming paradigm is introduced and it is discussed how Big Data processing frameworks use it to define scalable and distributed applications. Although functional programming has a long tradition, in the last few years it is becoming very popular, driven by the success of some languages, such as Scala, but also by the adoption of functional programming principles in mainstream languages like Java and Python. The course includes practical examples and real case studies aiming at highlighting how Apache Spark, one of the most popular frameworks for Big Data analysis, exploits functional programming for implementing scalable Big Data analysis applications.
From Big Data to Big Multidimensional Data: Models, Issues, Challenges	12	primo anno	Big data is gaining momentum in the research community, due to the several challenges posed by the management of such kind of data. Big data are relevant not only in the academic context, but also in the industrial context, where they play the major role. Indeed, several kinds of application are now exploiting big data, such as: Web advertisement, social network intelligence, e-science applications, smart city applications, and so forth. Among big data, big multidimensional data are a special case of big data that fully expose the "famous" 3V (volume, velocity, variety) and are of relevant interest at now. In this course, tailored to PhD students in Computer Science and Computer Engineering, we first investigate foundations of big data, critical state-of-the-art analysis, research challenges and industrial applications. After that, we move the attention on special lectures focused on big multidimensional data.
ACTIVE LEARNING	12	primo anno	In this course we start by introducing the general framework of active learning and then go into the details of the most relevant techniques defined in the literature emphasizing their advantages and drawbacks. The student will be involved in both a theoretical study of the various active learning techniques and invited to apply the theory to practical case studies.
Approximate Computing for low-energy, highspeed and area-efficient digital circuits and systems: when "good enough" is better than "good"	12	primo anno	The design optimization of digital circuits and systems typically consists in a three-dimensional trade-off among energy dissipation, area occupancy and computational speed. In the last decade, breaking such a trade-off has become very challenging since the breakdown of Dennard's scaling, and as Moore's and Koomey's laws approaching their end. Therefore, newer devices, architectures, and design techniques have become extremely urgent, also due to the strict energy, speed and area- efficiency requirements dictated by the emerging Internet-of-Things applications. Approximate Computing (AC) is a recent design paradigm aiming to fill the gap between requirements and capabilities of current platforms. It consists in introducing a new dimension in the optimization space, accuracy, to significantly reduce the hardware complexity, energy consumption and computational time. AC can be applied in several application areas that are intrinsically resilient to computational errors, e.g., machine learning, sensor signal processing, data mining and multimedia. The degree of accuracy can span across all the vertical computing stack, starting from the algorithm level and going down to the circuit and device levels. Preferably, a cross layer interaction should be enabled to optimize the energy-area-speed-accuracy trade-off. In this course, we will describe the most recent techniques based on AC, focusing in particular on arithmetic circuits at transistor and logic level and on memory architectures. Moreover, we will present a general overview about how approximation can be leveraged at software and device levels. Finally, we will discuss about several application examples and computing platforms where AC can be applied, thus underlining the interdisciplinarity of such a design paradigm.
Binary Analysis with Applications to Machine and Deep Learning	12	primo anno	Extracting usable information from executable files is a crucial task in nowadays computing, since it is the basis for a plethora of analyses performed in various fields of computer science and electronics. Analyzing binaries, thus figuring out the true properties of binary programs, however, is not a straightforward task. This is because during the analysis we have to deal with missing symbolic and type information, location-dependent code and data, missing high-level abstraction, mixed code and data and so on. The course aims at giving a complete overview of the techniques, challenges, tools and applications of binary analysis, introducing fundamentals of both static and dynamic analysis. We will talk about techniques such as disassembling, debugging, binary instrumentation, taint analysis, symbolic execution and we will discuss approaches in security oriented scenarios and in other fields of interest as well. An important role in the course will be played by machine/deep learning and how these emerging topics can contribute in the context of binary analysis. Specifically, we will talk about (i) various machine/deep learning techniques as versatile approaches to solve binary analysis task, such as similarity detection and vulnerability discovery and (ii) applications of binary analysis as approaches for solving challenging learning tasks, such as authorship attribution and plagiarism detection.
Digital technologies and artificial intelligence law	12	primo anno	The course focuses on the main legal aspects of digital technologies and artificial intelligence. The first part of the course (modules 1 and 2) is dedicated to general aspects, relating to the development of a legal notion of cyberspace and the analysis of different regulatory approaches. Subsequently, having acknowledged the disciplinary inadequacy of formal legislation, the issue of non-regulatory based regulatory systems is addressed. In this sense, the need for an "interdisciplinary" approach is recognized that can lead to the definition of legal rules "modeled on the nature of things" to be regulated. Specifically, the disciplinary potential of design is analyzed, i.e. the suitability of design standards in regulating digital phenomena in a much more incisive way than the laws in a formal sense. From the examination of the advantages and criticalities of such an approach - especially with regard to respect for fundamental rights and freedoms – we come to envisage a new model of regulation, based on the interaction between legal and technical factors, the result of which is to to arrive at the creation of a so-called norm "techno-juridical" which derives its binding character, in the light of the principle of horizontal subsidiarity. The second part of the course (modules 3 and 4) addresses the problem of the so-called cybersecurity, with regard to both digital activities and more specifically artificial intelligence. We come to the definition - also in the light of the main European regulatory interventions - of a "proceduralized" and "multiphase" security model, based on the subsidiary interaction between legal, technical and organizational-managerial factors. Subsequently, the issue of responsibility deriving from the production and management of artificial intelligence systems is dealt with, both in terms of existing legislation and in terms of European regulatory proposals. Module 5 will address some sectoral issues (e.g. copyright; personal data; smartcontracts, predictive justice).
Deep Learning and Statistical Learning	6	primo anno	The course is aimed at reviewing a set of statistical and feature learning methods and tools which can be exploited in data science. We will start by studying probabilistic modeling, the problems of inference and parameter estimation, generative vs. discriminative learning, Bayesian learning. Within this framework, we will focus on the most prominent machine learning methods: Neural networks and Feature learning, latent variable modeling. We shall apply these mathematical tools to the problems of supervised, semi-supervised and unsupervised learning and to scenarios such as text and document modeling, image and video analysis, social network analysis, recommendation.
Numerical Computations on the Infinity Computer	6	primo anno	The course is dedicated to numerical calculus using infinite and infinitesimal numbers. The Infinity Computer is considered here as the main framework for this purpose. Examples of different challenging problems, where it can be applied, are presented during this course. In particular, it is demonstrated that it is possible in several cases to obtain high precision results dealing with infinite and infinitesimal numbers. Examples of a software implementation of such the framework are also presented as well as exercises on the respective arithmetics.
Quantum Computing and its Application to Machine Learning	12	primo anno	The course will introduce the basic elements of quantum information and quantum computation. Starting from the physical fundaments (principle of superposition, unitary evolution, principle of measurement), the course will introduce the students first to basic quantum gates, and then to quantum algorithms and quantum circuits. The course will discuss why, how, and in which contexts, a significant "quantum speedup" can be obtained by using quantum instead of classical computation. The course will expose the two most renowned quantum algorithms (Grover and Shor) and then will focus on the most recent applications of quantum computing to machine learning, with an eye to the applications fields for which both public and private companies are investing a huge amount of money: e-health, finance, etc.

			The theoretical presentation will be enriched by a seminar given by Francesco Plastina, Professor of Quantum Mechanics at the University of Calabria, who will review the Physics fundaments of quantum computing. Furthermore, the lectures will be complemented by lab sessions that will exploit the IBM Quantum Experience platform and the Python Qiskit library. The course will be oriented not only to computer engineers, but also to engineers that work in the fields of telecommunications and optimization, since the basics of quantum communications and the basic quantum algorithms for optimization will be introduced. To take full benefit from the course, the students will be invited to brush up their basic knowledge on complex numbers and linear algebra.
Lipschitz Global Optimization	6	primo anno	The course is dedicated to numerical solution of global optimization problems, where theobjective function is multiextremal, non differentiable, given as a "black box", and hard to evaluate, i.e., each its evaluation even at one point requires large amounts of computational or temporal resources. Moreover, it is supposed that the objective function is Lipschitz continuous. This condition has a very practical meaning: in physical systems it means that the energy changes in the system described by the function are always limited.
Statistical data analysis and signal processing techniques for imaging and non-destructive testing	12	primo anno	The course introduces some statistical data analysis and signal processing techniques for imaging and non- destructive testing applications which all rely on the least-squares theory. By starting with solving the linear regression problem with noisy data, the least-squares theory will be introduced and its relationship with the pseudo-inverse matrix illustrated. The pseudo-inverse approach will be then used to solve generic polynomial fitting, regression and non-linear fitting through iterative algorithms. The basic principles of inverse theory and deconvolution will be suddenly introduced and then the course will focus on signal processing applications, all developed for dealing with optimal deconvolution and linear system characterization in noisy environment. Wiener filter theory, matched filter theory, pulse compression and ARMA modelling will be presented and numerical exercises will be done by the students during the course.
Biomedical Systems for Remote Patient Monitoring	12	primo anno	Biomedical systems based on radar adoption represent an emerging technology in telemedicine and Internet of Things (IoT). They can remotely measure heartbeat and respiration, speed, and position of a subject even through clothes, blankets and many barriers (e.g., glass, doors and some walls), also in totally dark environments and in smoke-filled areas. Moreover, measurements can be made unobtrusively over an extended period in a supportive home environment offering excellent long-term care benefits. These features can effectively be used to measure contactlessly biomedical parameters of strong interest (heartbeat and respiration) as well as to detect emergency situations such as sleep disorders, apnea, dyspnea, arrhythmia, Sudden Infant Death Syndrome (SIDS), and fall incidents. This opens a multitude of applications in home and clinical environments. This course reviews recent advances in biomedical and healthcare applications of radars, by recalling the operational principles and focusing on advanced signal processing techniques. Validations of this technology will be also discussed.
Fault-tolerance in the Internet of Things	12	primo anno	The course is aimed at reviewing and analyzing redundancy-based fault-tolerant techniques for the Internet of Things as a paradigm to support two of the main goals of computer security: availability and integrity. First, the autonomic computing paradigm will be introduced according to which autonomous systems are able to manage errors and failures, and dynamically adapt themselves to changes without the need of human intervention. Then, various failure classes identified in the literature will be presented, focusing in particular on the fail-stop (a.k.a. crash) and the malicious (a.k.a. byzantine) failures, which to date, have been the most addressed failure classes. Finally, we will discuss the main approaches that have been used to address failures in three different areas, namely sensing, routing, and control, that are the three main tasks performed by the nodes of an IoT network. More in detail, the following contents will be addressed: data aggregation, majority voting, and optimal sensor placement algorithms, for the sensing task; probabilistic (a.k.a. gossip-based) vs. non-probabilistic multi-path-based protocols, and hierarchical-based protocols, for the routing task; state machine replication paradigm, distributed consensus protocols for both byzantine and crash faults, and consensus in synchronous vs. asynchronous systems, for the control task. To conclude, some references to alternative approaches like detection- and prevention-based techniques, and on the type of attacks IoT systems are subject to. Although the main argument of the course is the fault-tolerance in IoT-based systems, related topics of different nature come into play. This will allow PhD students to deepen their knowledge of concepts of general interest concerning different branches of computer science, such as cyber security, routing, optimization techniques, state machine replication, and distributed consensus protocols, particularly popular since the advent of the distributed ledgers technologies (like blockchain), which nowadays are receiving great interest in the industry and in various research fields.
Data Stream Mining	12	primo anno	In this course, we will discuss the problem of learning from data streams generated by evolving nonstationary processes. It will overview the advances of techniques, methods, and tools that are dedicated to managing, exploit and interpreting data streams generated from time-evolving environments. In particular, the tutorial will examine the problems of learning classification and regression models from high-speed streams of non-stationary data. How to design the experimental setup and evaluate those models. We will also discuss issues related to concept drift, change detection, and novelty detection. Auto-ML for data streams.
Drones for measurements and measurements for drones: current applicative scenarios and future prospective	12	primo anno	Drones are becoming popular as carrier for several sensors and measurement systems, due to their low weight, small size, low cost and easy handling, which make them flexible and suitable in many measurement applications, mainly when the quantity to be measured is spread over a wide area or it lies in human-hostile environments. However, the drone itself can interact with both the measurand and the sensors, thus influencing the measurement results, and as consequence the performance of the whole application. For this reason, the drone equipped with the sensors must be thought as a measurement platform and must be characterized as a whole. The course program consists of the following topics: - A Description of the general architecture of a drone, by highlighting its subsystems and the parameters that can influence the on-board sensors and measurement systems. - An overview of the sensors and measurement systems that can be embedded on the drone, by presenting their operating principle and applications. - A description of some measurement applications that use drones. For such applications, the measurement chain is analyzed, and the influence of the flight parameters is taken into account to assess the measurement uncertainty. - A review of test benches used for the electrical and mechanical testing of drones and their components. - A description of compressed sensing algorithm applied to images acquired from drones.
Hardware for Deep Learning	6	primo anno	Ever since the the breakthrough image recognition achievement, shown by AlexNet in the ImageNet challenge of 2012, deep learning has risen to unprecedented levels of popularity and adoption in virtually every aspect of today's world. However, the sheer number of computations required for both training and inference of state-of-the-art networks is infeasible for execution upon traditional general-purpose computing platforms. This has led to an explosion of both research and industrial implementations of acceleration platforms and hardware-aware algorithmic approaches for fast and efficient computation of deep learning tasks. This course briefly introduces the basic concepts of deep learning and the mechanics of the fundamental computations required for inference and training. This is followed by a deep dive into popular methods for building hardware to efficiently execute these tasks, as well as an overview of leading hardware-aware algorithmic methods to reduce the complexity of the computation. At the end of this six-hour course, the students will have an understanding of the deep learning hardware landscape that will enable them to approach further research hardware design for artificial intelligence applications, as well as provide basic knowledge to software designers about the implications of their code on the underlying hardware.
Modelling and mining multilayer networks	12	primo anno	A multilayer network is a network (or graph) where nodes can be organised into sets, called layers, and the same node can belong to one or more of the layers. This allows us to model a wide range of systems of interconnected entities, for example social networks where different types of actors are connected through different types of ties ("working together", "being friend", etc.). After presenting the multilayer network model, this course focuses on the analysis of a specific type of multilayer network, known as multiplex, where there is only one type of actors but multiple types of connections. The course covers a selection of topics related to community detection, layer comparison methods, actor measures, data exploration, and network generation. For each topic, a quick presentation of the relevant theory and methods will be followed by a practical application on a real pedagogical dataset. The practical tasks can be performed in Python or R; basic knowledge of Python or R is useful but not a strict requirement. The course is organised over four consecutive days. The first two days, the lecturer presents the theory and the attendees try some practical exercises to get familiar with the terminology and the main methods (one 3- hour seminar per day). The third day the attendees work independently, by reading and preparing a presentation of some selected papers on multilayer networks. On the fourth day, the papers are presented and discussed.
IoT and edge computing: stories of a necessary relationship	8	primo anno	In this technical seminar, organized in 4 sessions, two mainly theoretical and two mainly practical, we will try to trace the approach path from IoT to edge computing, presenting, as a main result, the concept of TinyML, i.e. the use of machine learning in 1 mW devices. We will start with a description of what the IoT is and what are its various components, focusing on two relevant communication protocols: LoRaWAN, as a low power communication standard and MQTT as a widely used protocol to connect devices and services in a more flexible way than the classical REST (HTTP) based solutions. The next step, in the first practical session, will be to provide an overview of

			the data path, from the sensors through the different network components, gateways, networks servers, visualization platforms up to the "TIG stack" used to facilitate the collection, storage, visualization and generation of alerts from time series data. The next two sessions will justify the need for edge computing in IoT and describe TinyML, which is a rapidly growing area where machine learning technologies and applications, including hardware, algorithms and software, are capable of performing sensor data analysis on devices with extremely low power consumption, typically in the mW range and below, and thus enabling a variety of broader use cases targeting battery-powered devices. The theory session will describe the process of using the TensorFlow Lite libraries and devices such as the Arduino Nano 33 BLE Sense and the hands-on session will use the Edge Impulse platform, currently the leading development platform for ML on edge devices, to build a real-time model using the accelerometer, microphone or camera of a smartphone, how to collect data and train machine learning algorithms, and observe what happens live on the platform.
Advanced Deep Learning Models	12	primo anno	Deep Learning (DL) is used in different fields of engineering and is proving exceptionally useful for solving problems of extraordinary complexity that until a few years ago could only be addressed by human beings. The translation of texts from one language to another, autonomous driving, image analysis are just some of the problems that can be solved by Deep Learning models with performances comparable or superior to those of human operators. The topic of this course will be the presentation of two advanced Deep Learning models: Graph Neural Networks (GNNs) and Deep Reinforcement Learning (DRL). Graph Neural Networks are Deep Learning models used when the input data does not have a sequential (e.g. texts) or matrix (e.g. images) structure but can be modeled with a graph. The GNNs allow to associate to each node a data structure that summarizes its properties (embedding) and which is calculated by aggregating and processing the information of the node with that of the neighboring nodes that are at most a certain number of hops away from the node itself. This process can be performed with a message-passing mechanism which is very similar to how Convolutional Neural Networks (CNNs) extract features from images. The embeddings provide an easy way to do node-level, edge-level, and graph-level prediction tasks. GNNs are very powerful tools and are successfully applied in many different domains such as drug research, fraud detection, route planning and network optimization. Deep Reinforcement Learning models combine the Reinforcement Learning (RL) paradigm with Deep Learning. Reinforcement Learning requires an agent to operate in an environment by performing actions that change the state of the environment. The agent receives rewards and penalties and has the goal of maximizing his earnings. The agent's decisions are returned by a Deep Learning model that is trained in order to learn the most convenient actions. Deep Reinforcement Learning models have achieved super-human performance; as an example, they excel in robotic tasks and can beat human players in competitive games (e.g., Atari, StarCraft, Dota, and Go).
Novel device and architecture concepts for CMOS Logic scaling: a reliability engineer perspective	6	primo anno	Despite the slowdown of the traditional transistor gate length scaling from one technology node to the next, Moore’s law is alive and well and keeps dictating the performance enhancement and cost scaling pace of CMOS technology. This has been sustained in recent years by the introduction of several disruptive process technology innovations (e.g., strain engineering for mobility enhancement as of the 90nm node, high-k/metal gate introduction as of the 45nm node, multi-Vth offering by gate metal stack engineering enabled by the so-called Replacement Gate integration introduced as of the 32nm node), of a novel device architecture (transition from the traditional planar transistor structure to the finFET as of the 22nm node), and of a novel channel material (SiGe for enhanced hole mobility as of the 5nm node).We will review these recent innovations and discuss device architectures and novel integration concepts which are currently considered for possible introduction in next CMOS technology nodes, such as: Gate-All-Around nanowires or nanosheets Complementary-FETs, Sequential 3D stacking of multiple device tiers within a single chip fabrication flow, Buried Power Delivery Network. While technology innovation is primarily driven by performance enhancement requirements, novel devices need to guarantee traditional standards in terms of reliable operation, despite the inherently increased fabrication complexity. Device reliability is therefore a crucial aspect that needs to be carefully assessed early on when down selecting options for future technology nodes. We will review the physics of the main CMOS device degradation mechanisms (Bias Temperature Instabilities, Hot Carrier Degradation, Time-Dependent Dielectric Breakdown, Self-Heating Effects), and discuss the challenges and opportunities that novel device and architecture concepts present in terms of device reliability.
Binary Analysis with Applications to Machine and Deep Learning	12	secondo anno	Extracting usable information from executable files is a crucial task in nowadays computing, since it is the basis for a plethora of analyses performed in various fields of computer science and electronics. Analyzing binaries, thus figuring out the true properties of binary programs, however, is not a straightforward task. This is because during the analysis we have to deal with missing symbolic and type information, location-dependent code and data, missing high-level abstraction, mixed code and data and so on. The course aims at giving a complete overview of the techniques, challenges, tools and applications of binary analysis, introducing fundamentals of both static and dynamic analysis. We will talk about techniques such as disassembling, debugging, binary instrumentation, taint analysis, symbolic execution and we will discuss approaches in security oriented scenarios and in other fields of interest as well. An important role in the course will be played by machine/deep learning and how these emerging topics can contribute in the context of binary analysis. Specifically, we will talk about (i) various machine/deep learning techniques as versatile approaches to solve binary analysis task, such as similarity detection and vulnerability discovery and (ii) applications of binary analysis as approaches for solving challenging learning tasks, such as authorship attribution and plagiarism detection.
ACTIVE LEARNING	12	secondo anno	Nowadays the large availability of digital data from heterogeneous sources opens up new application scenarios for Artificial Intelligence techniques aiming at devising classifiers and predictors, using previously labeled (classified) instances to train classifiers for new instances. Unfortunately, there are many contexts characterized by the presence of a huge amount of unlabeled data or no labeled data or a scarce number of labeled data thus making it impossible to traingood classifiers. Active learning is a machine learning framework in which a learning algorithm interacts with a human expert (or another source of information) to classify fresh data points in order to add them to the pool of labeled instances. Indeed,since attaching class labels to data is generally costly and time consuming, Active Learning techniques aim at selecting the unlabeled data that give the best improvement in training good classifiers. Several active learning techniques have been defined in the literature that can usually be grouped into the following two main categories. ● Uncertainty based: Techniques falling in this category assume the presence of a current model. Unlabeled instances are ranked w.r.t the uncertainty of the model in assigning them a class. The idea is that the more uncertain the model isin assigning a class to an instance, the more the learning strategy will benefit from having the possibility to know the real class. ● Density based: Techniques belonging to this category rank instances according to the density of their neighborhood.The idea is that this way the more representative instances are selected avoiding outliers. Several other techniques have been defined that do not fall into the above mentioned categories such as techniques measuring the expected model change reduction or techniques that use machine learning to understand which are thebest instances to be selected for labeling. In this course we start by introducing the general framework of active learning and then go into the details of the most relevant techniques defined in the literature emphasizing their advantages and drawbacks. The student will be involvedin both a theoretical study of the various active learning techniques and invited to apply the theory to practical case studies.
Statistical data analysis and signal processing techniques for imaging andnon-destructive testing	12	secondo anno	The course introduces some statistical data analysis and signal processing techniques for imaging and non-destructive testing applications which all rely on the least-squares theory. By starting with solving the linear regression problem with noisy data, the least-squares theory will be introduced and its relationship with the pseudo-inverse matrix illustrated. The pseudo-inverse approach will be then used to solve generic polynomial fitting, regression and non-linearfitting through iterative algorithms. The basic principles of inverse theory and deconvolution will be suddenly introduced and then the course will focus on signal processing applications, all developed for dealing with optimal deconvolution and linear system characterization in noisy environment. Wiener filter theory, matched filter theory, pulse-compressionand ARMA modelling will be presented and numerical exercises will be done by the students during the course.
Biomedical Systems for RemotePatient Monitoring	12	secondo anno	Biomedical systems based on radar adoption represent an emerging technology in telemedicine and Internet of Things (IoT). They can remotely measure heartbeat and respiration, speed, and position of a subject even through clothes, blankets and many barriers (e.g., glass, doors and some walls), also in totally dark environments and in smoke-filled areas. Moreover, measurements can be made unobtrusively over an extended period in a supportive home environmentoffering excellent long-term care benefits. These features can effectively be used to measure contactlessly biomedical parameters of strong interest (heartbeat and respiration) as well as to detect emergency situations such as sleep disorders, apnea, dyspnea, arrhythmia, Sudden Infant Death Syndrome (SIDS), and fall incidents. This opens a multitude of applications in home and clinical environments. This course reviews recent advances in biomedical and healthcare applications of radars, by recalling the operationalprinciples and focusing on advanced signal processing techniques. Validations of this technology will be also discussed.

Advanced Deep Learning Models	12	secondo anno	Deep Learning (DL) is used in different fields of engineering and is proving exceptionally useful for solving problems ofextraordinary complexity that until a few years ago could only be addressed by human beings. The translation of texts from one language to another, autonomous driving, image analysis are just some of the problems that can be solved by Deep Learning models with performances comparable or superior to those of human operators. The topic of this course will be the presentation of two advanced Deep Learning models: Graph Neural Networks (GNNs) and Deep Reinforcement Learning (DRL). Graph Neural Networks are Deep Learning models used when the input data does not have a sequential (e.g. texts) or matrix (e.g. images) structure but can be modeled with a graph. The GNNs allow to associate to each node a data structurethat summarizes its properties (embedding) and which is calculated by aggregating and processing the information of thenode with that of the neighboring nodes that are at most a certain number of hops away from the node itself. This processcan be performed with a message-passing mechanism which is very similar to how Convolutional Neural Networks (CNNs) extract features from images. The embeddings provide an easy way to do node-level, edge-level, and graph-levelprediction tasks. GNNs are very powerful tools and are successfully applied in many different domains such as drug research, fraud detection, route planning and network optimization. Deep Reinforcement Learning models combine the Reinforcement Learning (RL) paradigm with Deep Learning. Reinforcement Learning requires an agent to operate in an environment by performing actions that change the state of theenvironment. The agent receives rewards and penalties and has the goal of maximizing his earnings. The agent's decisionsare returned by a Deep Learning model that is trained in order to learn the most convenient actions. Deep ReinforcementLearning models have achieved super-human performance; as an example, they excel in robotic tasks and can beat humanplayers in competitive games (e.g., Atari, StarCraft, Dota, and Go).
Circuits, Architectures and Systems for Intelligent Computing at the Edge	12	secondo anno	Deep learning oriented to image classification, object detection and segmentation tasks has reached unprecedented popularity in several real-world applications, ranging from intelligent autonomous vehicles to smart manufacturing processes. In these contexts, ensuring real-time performances also on power-constrained devices is crucial. However, state-of-the-art deep learning models, and in particular Convolutional Neural Networks (CNNs), owe their high accuracyto massive computational and memory requirements, which hinder their deployment into edge computing systems. This issue has opened the door to an exciting study field that aims at researching new design and algorithmic strategies in orderto reduce the computational complexity of deep learning models without compromising the achievable accuracy. In this course, we will introduce the basics of CNNs with reference to different computer vision tasks. After that, we willprovide an overview on techniques currently used to efficiently implement CNN inference on low-power edge devices through data-level approximations, such as quantization and pruning. Some noteworthy state-of-the-art FPGA and ASICimplementations will be also presented. Finally, we will describe more recent advances in the field of layer-level approximate computing techniques and their application to state-of-the-art models, discussing how well error-resilient tasks like CNNs can really tolerate aggressive approximations and quantifying the benefits of such strategies on power consumptions. At the end of this course, students will have a comprehensive knowledge of the main design techniques applicable at both circuit-, architecture- and system-level to enable intelligent computing on edge devices. They also willget an understating of methodologies and tools that can be used in varying artificial intelligence applications also relatedto their research topics.
Distributed Simulation of Complex Engineered Systems	12	secondo anno	The ever-growing advances in science and technology have led to a rapid increase in the complexity of most engineeredsystems. Complex Engineered Systems (CES) are the result of this technological advancement that involves new paradigms, architectures, and functionalities derived from different engineering domains. Generally, such systems are composed of many heterogeneous components, often designed and manufactured by organizations belonging to several industrial domains such as mechanical, electrical, and software. Components are characterized by independent behaviorsand interact with each other to pursue the objectives of the whole CES. Understanding, studying, and designing CES is going to be a great challenge in the next few years. In this context, Distributed Simulation (DS) represents a powerful methodology to support the analysis and design of such systems by enabling the evaluation and comparison of different design alternatives against requirements through virtual testing. This opportunity becomes even more crucial when complete and actual tests are too expensive to be performed in terms of cost, time, and other resources. This course introduces Distributed Simulation as a methodology for analyzing, designing, and operating CES. Key aspectswill concern the definition and implementation of distributed simulators according to the IEEE 1516.2010 standard and the analysis of the simulation results. The course relies on a problem-solving approach: starting from real case studies, students will be required to concretely apply the introduced Distributed Simulation methodology, tools, and techniques.
Mathematical Optimization for Machine Learning	8	terzo anno	The course is aimed at providing basic Numerical Optimization tools to handle some classes of Machine Learning problems, particularly focusing on supervised classification. Optimality conditions for functions of several variables both in the constrained and the unconstrained case will be briefly recalled, along with some fundamental notions in convex analysis. The problem of separating sets in n-dimensional spaces by appropriate separation surfaces will be put in the form of an optimization problem. The Support Vector Machine (SVM) approach, where separating hyperplanes are adopted for classification purposes, will be discussed, together with possible alternative separation methods based on piecewise affine, ellipsoidal, and spherical surfaces. Multiple-Instance classification models will be discussed as well, together with the results of some applications.
Advanced Cyber Security	8	terzo anno	The course aims at discussing the main methodological aspects of cybersecurity and at giving an overview of research issues that have recently received huge attention. By taking the course, the students will acquire knowledge about how security scenarios should be appropriately described and analyzed, and about interesting related research issues. The course topics are as follows. Motivations of attacks and attack vectors. Security policies. Security objectives: confidentiality, integrity, availability, authenticity, accountability. Threat models. Reference user models: network users, snooping users, co-located users. Security mechanisms. Main kinds of mechanisms: authentication, authorization, audit. Classes of attacks: eavesdropping, modification, interruption, forging. Typical issues with the design of security policies, threat models, and security mechanisms. Design principles and design rules. Security architectures. Artificial intelligence techniques for cybersecurity. Intellectual property protection, fraud detection in reviewing systems, adversarial defense of enterprise systems. Data management techniques for cybersecurity. Activity detection.
Advanced Edge-Cloud computing systems	8	terzo anno	The applications used today for processing Big Data repositories are highly centralized and leverage cloud platforms to perform all major operations involving data collection, storing, processing, analysis and machine learning. However, using only the cloud can result in serious inefficiencies in terms of network traffic, latency times and optimization of energy consumption. In the latest years researchers and IT companies have proposed the adoption of the edge computing paradigm for processing data closer to where they are generated, for achieving low latency, privacy preserving and scalability. These benefits can be complemented by those provided by the cloud, which allows for aggregating large amounts of data persistently and performing compute-intensive analysis using a large amount of computing resources. This scenario defines the so-called edge-cloud compute continuum. In this course we will introduce the basic concepts of edge-cloud compute continuum and describe models, architectures, and frameworks based on this paradigm. Applications of cloud-edge computing systems in the area of Big Data analysis and machine learning will also be presented and discussed.
Digital technologies and artificial intelligence law	8	terzo anno	The course focuses on the main legal aspects of digital technologies and artificial intelligence. The first part of the course (modules 1 and 2) is dedicated to general aspects, relating to the development of a legal notion of cyberspace and the analysis of different regulatory approaches. Subsequently, having acknowledged the disciplinary inadequacy of formal legislation, the issue of non-regulatory based regulatory systems is addressed. Inthissense, the need for an "interdisciplinary" approach is recognized that can lead to the definition of legal rules "modeledon the nature of things" to be regulated. Specifically, the disciplinary potential of design is analyzed, i.e. the suitability of design standards in regulating digital phenomena in a much more incisive way than the laws in a formal sense. From the examination of the advantages and criticalities of such an approach - especially with regard to respect forfundamental rights and freedoms – we come to envisage a new model of regulation, based on the interaction between legal and technical factors, the result of which is to to arrive at the creation of a so-called norm "techno- juridical" which derives its binding character, in the light of the principle of horizontal subsidiarity. The second part of the course (modules 3 and 4) addresses the problem of the so-called cybersecurity, with regard toboth digital activities and more specifically artificial intelligence. We come to the definition - also in the light of the main European regulatory interventions - of a "proceduralized" and "multiphase" security model, based on the subsidiary interaction between legal, technical and organizational-managerial factors. Subsequently, the issue of responsibility deriving from the production and management of artificial intelligence systems is dealt with, both in terms of existing legislation and in terms of European regulatory proposals. Module 5 will address some sectoral issues (e.g. copyright; personal data; smartcontracts, predictive justice).

Approximate Computing & digital systems	8	terzo anno	The design optimization of digital circuits and systems typically consists in a three-dimensional trade-off among energy dissipation, area occupancy and computational speed. In the last decade, breaking such a trade-off has become very challenging since the breakdown of Dennard's scaling, and as Moore's and Koomey's laws approaching their end. Therefore, newer devices, architectures, and design techniques have become extremely urgent, also due to the strict energy, speed and area- efficiency requirements dictated by the emerging Internet-of-Things applications. Approximate Computing (AC) is a recent design paradigm aiming to fill the gap between requirements and capabilities of current platforms. It consists in introducing a new dimension in the optimization space, accuracy, to significantly reduce the hardware complexity, energy consumption and computational time. AC can be applied in several application areas that are intrinsically resilient to computational errors, e.g., machine learning, sensor signal processing, data mining and multimedia. The degree of accuracy can span across all the vertical computing stack, starting from the algorithm level and going down to the circuit and device levels. Preferably, a cross layer interaction should be enabled to optimize the energy-area-speed-accuracy trade-off. In this course, we will describe the most recent techniques based on AC, focusing in particular on arithmetic circuits at transistor and logic level and on memory architectures. Moreover, we will present a general overview about how approximation can be leveraged at software and device levels. Finally, we will discuss about several application examples and computing platforms where AC can be applied, thus underlining the interdisciplinarity of such a design paradigm.
A Primer on Resilient Control Methodologies for Cyber-Physical Systems	8	terzo anno	Recent progress on high-speed networks, wireless communication technologies and the development of novel control strategies for embedded systems gave rise to a boost in the deployment of the cyber-physical systems (CPSs) paradigm within a wide range of applications. The heterogeneous nature of CPS components may give intruders the chance of launching severe attacks. Therefore, control solutions capable of ensuring CPS safety and performance under cyber-attacks are extremely important for security issues. How to keep system operations at a satisfactory level in the presence of attacks still remains an open challenge. Usually, the assumption is that the intruder is removed from the system architecture once detected. On the other hand, in most of the operating scenarios the plant must operate even if the attack is running: this poses the key question to develop a joint design for the detector and the controller in order to maintain suitable plant performance. According to these premises, the so-called resilient control problem for constrained cyber-physical systems subject to false data injections is addressed. The core of the proposed course consists in defining an ad-hoc versatile framework whose main feature consists in the ability of being geared to different classes of attacks. This is formally achieved by resorting to the receding horizon philosophy that is fully exploited for detection, countermeasures and control purposes. In particular, set-theoretic model predictive arguments are combined with the perturbation analysis and sequential quadratic programming to reduce as much as possible the occurrence of refresh procedures on the communication network when resilient command actions are no longer available. Further, one of its main merits consists in the dismissal of constructive assumptions existing in recent competitors. In this respect, the framework is then customized for replay and covert attacks by specifying actuation/detection phases and proving feasibility and closed-loop stability properties.
Emerging networking paradigms for 5G/6G systems	8	terzo anno	Future telecommunication networks will definitely be the key enablers for emerging critical services such as autonomous driving, smart industry, AR/VR, and remote medicine, that require low latency and high reliability, along with massive connectivity and data availability. In view of this, they will have to evolve and overcome their current limits. This course aims precisely to present new network paradigms being defined for future 5G and 6G telecommunications systems. Advanced technological solutions will be presented that support the creation of: edge-device software platforms that are modular, open, and scalable, and able to meet the requirements of novel emerging services; open software frameworks for network data plane and control plane programmability; programmable access networks leveraging virtualization and flexible (re-)configuration; and breakthrough algorithmic solutions for network resource orchestration. At the end of this course, students will be able to apply the theoretical knowledge acquired for solving problems related to the design, implementation, and management of 5G/6G network architectures based on the new virtualization paradigms in order to guarantee a more adequate response to user requests.
Numerical Computations on the Infinity Computer	8	terzo anno	The course is dedicated to numerical calculus using infinite and infinitesimal numbers. The Infinity Computer is considered here as the main framework for this purpose. Examples of different challenging problems, where it can be applied, are presented during this course. In particular, it is demonstrated that it is possible in several cases to obtain high precision results dealing with infinite and infinitesimal numbers. Examples of a software implementation of such the framework are also presented as well as exercises on the respective arithmetics
Functional Programming for Big Data Processing	8	terzo anno	Through the pervasive use of computers, smartphones and other digital objects, huge amounts of digital data are generated and collected. This data, commonly referred as Big Data, represents a challenge for the current storage, process, and analysis capabilities. To extract value from such data, novel architectures, programming models and frameworks have been developed in recent years for capturing and analyzing complex and/or high velocity data. The usage of high performance computers, such as Clouds and clusters, paired with parallel and distributed algorithms, are commonly used by data analysts to solve big data problems and obtain valuable information and knowledge in a reasonable time. This course introduces the most effective programming approaches for Big Data processing. In particular, the functional programming paradigm is introduced and it is discussed how Big Data processing frameworks use it to define scalable and distributed applications. Although functional programming has a long tradition, in the last few years it is becoming very popular, driven by the success of some languages, such as Scala, but also by the adoption of functional programming principles in mainstream languages like Java and Python. The course includes practical examples and real case studies aiming at highlighting how Apache Spark, one of the most popular frameworks for Big Data analysis, exploits functional programming for implementing scalable Big Data analysis applications.
Advanced Deep Learning Models	8	terzo anno	The topic of this course will be the presentation of two advanced Deep Learning models: Graph Neural Networks (GNNs) and Deep Reinforcement Learning (DRL). Graph Neural Networks are Deep Learning models used when the input data does not have a sequential (e.g. texts) or matrix (e.g. images) structure but can be modeled with a graph. The GNNs allow to associate to each node a data structure that summarizes its properties (embedding) and which is calculated by aggregating and processing the information of the node with that of the neighboring nodes that are at most a certain number of hops away from the node itself. This process can be performed with a message-passing mechanism which is very similar to how Convolutional Neural Networks (CNNs) extract features from images. The embeddings provide an easy way to do node-level, edge-level, and graph-level prediction tasks. GNNs are very powerful tools and are successfully applied in many different domains such as drug research, fraud detection, route planning and network optimization. Deep Reinforcement Learning models combine the Reinforcement Learning (RL) paradigm with Deep Learning. Reinforcement Learning requires an agent to operate in an environment by performing actions that change the state of the environment. The agent receives rewards and penalties and has the goal of maximizing his earnings. The agent's decisions are returned by a Deep Learning model that is trained in order to learn the most convenient actions. Deep Reinforcement Learning models have achieved super-human performance; as an example, they excel in robotic tasks and can beat human players in competitive games (e.g., Atari, StarCraft, Dota, and Go).
Binary Analysis with Applications to Machine and Deep Learning	8	terzo anno	Extracting usable information from executable files is a crucial task in nowadays computing, since it is the basis for a plethora of analyses performed in various fields of computer science and electronics. Analyzing binaries, thus figuring out the true properties of binary programs, however, is not a straightforward task. This is because during the analysis we have to deal with missing symbolic and type information, location-dependent code and data, missing high-level abstraction, mixed code and data and so on. The course aims at giving a complete overview of the techniques, challenges, tools and applications of binary analysis, introducing fundamentals of both static and dynamic analysis. We will talk about techniques such as disassembling, debugging, binary instrumentation, taint analysis, symbolic execution and we will discuss approaches in security oriented scenarios and in other fields of interest as well. An important role in the course will be played by machine/deep learning and how these emerging topics can contribute in the context of binary analysis. Specifically, we will talk about (i) various machine/deep learning techniques as versatile approaches to solve binary analysis task, such as similarity detection and vulnerability discovery and (ii) applications of binary analysis as approaches for solving challenging learning tasks, such as authorship attribution and plagiarism detection
Data Stream Mining	8	terzo anno	In this course, we will discuss the problem of learning from data streams generated by evolving nonstationary processes. It will overview the advances of techniques, methods, and tools that are dedicated to managing, exploit and interpreting data streams generated from time-evolving environments. In particular, the tutorial will examine the problems of learning classification and regression models from high-speed streams of non-stationary data. How to design the experimental setup and evaluate those models. We will also discuss issues related to concept drift, change detection, and novelty detection. Auto-ML for data streams

Modelling and mining multilayer networks	8	terzo anno	A multilayer network is a network (or graph) wherenodes can be organised into sets, called layers, and the same node can belong to one or more of the layers. This allows us to model a wide range of systems of interconnected entities, for example social networks where different types of actors are connected through different types of ties ("working together", "being friend", etc.). After presenting the multilayer network model, this course focuses on the analysis of a specific type of multilayer network, known as multiplex, where there is only one type of actors but multiple types of connections. The course covers a selection of topics related to community detection, layer comparison methods, actor measures, data exploration, and network generation. For each topic, a quick presentation of the relevant theory and methods will be followed by a practical application on a real pedagogical dataset. The practical tasks can be performed in Python or R; basic knowledge of Python or R is useful but not a strict requirement. The course is organised over four consecutive days. The first two days, the lecturer presents the theory and the attendees try some practical exercises to get familiar with the terminology and the main methods (one 3-hour seminar per day). The third day the attendees work independently, by reading and preparing a presentation of some selected papers on multilayer networks. On the fourth day, the papers are presented and discussed.
Quantum Computing and its Application to Machine Learning	8	terzo anno	The course will introduce the basic elements of quantum information and quantum computation. Starting from the physical fundaments (principle of superposition, unitary evolution, principle of measurement), the course will introduce the students first to basic quantum gates, and then to quantum algorithms and quantum circuits. The course will discuss why, how, and in which contexts, a significant "quantum speedup" can be obtained by using quantum instead of classical computation. The course will expose the two most renowned quantum algorithms (Grover and Shor) and then will focus on the most recent applications of quantum computing to machine learning, with an eye to the applications fields for which both public and private companies are investing a huge amount of money: e-health, finance, etc.