N.	Denominazione dell'insegnamento	Numero di ore totali	Distribuzione durante il ciclo di dottorato	Descrizione del corso
1	Approximate Computing & digital systems	4	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.
2	Control and Coordination in Cyber-Physical Systems	12	primo anno	The increasing reliance on cyber-physical systems (CPSs) in applications such as autonomous robotics, intelligent transportation, and industrial automation has raised critical concerns about their resilience to cyber and operational disruptions. Ensuring reliable coordination and control in these distributed systems requires advanced strategies that can detect anomalies, mitigate disturbances, and maintain stability under adverse conditions. This course provides a comprehensive introduction to resilient control strategies for CPSs, with a focus on leader-follower (LF) formations and distributed model predictive control (DMPC). Participants will explore potential vulnerabilities in CPSs, advanced anomaly detection methods, and mitigation strategies, including altruistic recovery mechanisms and adaptive formation reconfiguration. Additionally, the course will cover flexible leader-follower formations in dynamic environments, enabling robust decision-making in the presence of uncertainties and structural changes. Theoretical concepts will be reinforced through simulation-based case studies, demonstrating real-world applicability in networked robotic systems and cooperative vehicle coordination.
3	Advanced cybersecurity	12	primo anno	The course aims to discuss modern cybersecurity issues and provide an overview of research topics that have recently received significant attention. By taking the course, students will acquire knowledge of these topics, as well as related research issues. The first part of the course will cover methodological aspects related to the description and analysis of security scenarios. Next, the course will explore key topics in IoT security, highlighting the main differences between IoT security and traditional security paradigms, as well as attack surfaces and common vulnerabilities. Students will then delve into firmware analysis techniques, focusing on firmware structure, extraction methods, cross-compilation, and re-hosting for security testing. The course will also provide insights into device exploitation, covering attacks via UART and SPI interfaces, protocol-based vulnerabilities, and firmware exploitation strategies. In addition, through theoretical discussions and hands-on tutorials, participants will gain practical expertise in analyzing and securing embedded systems, equipping them with essential skills for both research and real-world cybersecurity challenges. Finally, the course will introduce tools and methodologies used in embedded systems security research, providing the foundation for further studies in the field.
4	Advanced Machine Learning	12	primo anno	Machine Learning (ML) 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 Machine 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).
5	Federated Learning: Challenges, Trends and Emerging Applications	12	primo anno	This course introduces students to the fundamentals in federated learning, with a focus on model aggregation techniques and different types of model aggregations. Federated learning is a distributed machine learning paradigm that enables model training on a large body of decentralized data. Its goal is to make full use of data across organizations or devices while meeting regulatory, privacy, and security requirements. By critically analyzing the current state of the field, this course will provide insight on the current trends, challenge and emerging applications. For example a special focus will be devoted to green federated learning, which is a new paradigm for researchers interested in working on energy-efficient AI systems or solutions.
6	Data Analysis & Signal Processing for Electrical Engineering Applications	8	primo anno	The course introduces some data analysis and signal processing tools for electrical engineering 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 also be introduced focusing on signal processing applications, all developed for dealing with optimal deconvolution and linear system characterization in noisy environment. Numerical exercises and experimental demonstrations will be done during the course.
7	Privacy in Information Engineering - Digital technologies and artificial intelligence law	4	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. 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).
8	Mathematical Optimization for Machine Learning - Foundations	8	primo anno	Behind every successful machine learning model lies an optimization engine. In this course we introduce the fundamental mathematical tools and techniques to understand optimization algorithms crucial for training and deploying effective machine learning models. We focus on nonlinear optimization problems, i.e., optimization problems where the objective function or constraints, or both, are nonlinear. We cover topics like unconstrained and constrained optimization, convex optimization, gradient-based methods, Newton-type methods, and more.
9	Advanced methods to design hardware accelerators	8	secondo anno	Modern applications, like Computer Vision (CV), Internet of Things (IoT), Deep Learning (DL) oriented to image classification, object detection and segmentation tasks, intelligent autonomous vehicles, smart manufacturing, and many others, demand high computational speed, significant energy efficiency and flexibility. For this reason, designing hardware accelerators currently receives a great deal of attention. This course overviews the main advanced methodologies to design both Application Specific Integrated Circuits (ASICs) and Field-Programmable-Gate-Array (FPGA)-based hardware accelerators suitable to support the above cited applications. In particular, design and algorithmic strategies exploitable to reduce the computational complexity of deep learning models without compromising the achievable accuracy are examined. Moreover, several design techniques oriented to edge computing systems are explored. The course also provides 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 ASIC implementations will be also presented. At the end of this course, students will have a comprehensive knowledge of the main design techniques applicable at circuit-, architecture- and system-level to hardware accelerate computationally intensive elaborations. They also will get an understating of methodologies and tools that can be used in several artificial intelligence applications also related to their research topics.
10	Microwaves for Sensing: Circuits and Methods for Biomedical Data Analysis	8	secondo anno	The course is devoted to illustrate the basic principles of sensing procedures adopting microwaves and millimeter-waves. In the first part, the sensors/circuits architectures to be adopted for sensing applications are studied, while in the second part a set of super-resolution signal processing techniques are illustrated, which are able to estimate relevant biomedical parameters from a limited set of data acquired with microwave sensors. In particular, a spectral estimation approach with a specific statistical analysis is presented to guarantee the accurate identification of the biomedical parameter of interest, with high robustness against noise. Numerical applications on selected biomedical contexts will be considered.
11	Emerging networking paradigms for 5G/6G systems	8	secondo 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 can 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, AI driven solutions for 5G/6G autonomous network management and optimization. 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
12	Mastering Large Language Models (LLMs) – From Foundations to Real-World Applications	12	secondo anno	Large Language Models (LLMs) represent a major breakthrough in artificial intelligence, exhibiting unprecedented capabilities in natural language understanding and generation. As these models continue to shape industries and research fields, understanding their underlying principles is essential to fully harness their potential. This course provides a comprehensive exploration of LLMs, from their theoretical foundations to real-world use cases. Students will explore the core architecture behind LLMs, including encoders (e.g., BERT) and decoders (e.g., GPT) models, gaining a deep understanding of the Transformer framework. The course will provide an overview of the full development cycle of LLMs, covering data preparation, training procedures, and post-training refinement techniques (e.g., model alignment with human expectations, evaluation, safety). Beyond theoretical comprehension, students will be introduced to the state-of-the-art frameworks for the training, deployment, and usage of LLMs (e.g., HuggingFace Transformers' library, Unsloth), with a particular focus on open-weight and small models, suitable for computationally efficient applications. Finally, the course will contextualize LLMs within a broader scope, discussing their limitations, ethical and societal considerations, and emerging frontiers–including retrieval augmented generation and multimodal models. At the end of this course, students will have acquired a comprehensive understanding of the inner workings of LLMs the state-of-the-art frameworks and tools. They will develop a critical awareness of both the capabilities and limitations of these models, enabling more effective and responsible use, empowering them to exploit LLMs to support their research topics.

13	Digital Forensics: Cybersecurity and Privacy Challenges	12	secondo anno	The increasing prevalence of cybercrime, data breaches, and digital espionage has placed Digital Forensics at the forefront of research in cybersecurity and digital privacy protection. This course, designed for PhD students in ICT, offers an in-depth exploration of forensic methodologies tailored to addressing modern cybersecurity challenges and safeguarding digital privacy. Beginning with the fundamental principles of digital forensics—such as data acquisition, preservation, and analysis—this course will lead students through the technical and legal frameworks governing digital evidence handling. The focus will then shift to advanced forensic techniques, emphasizing their application in the context of network security, encrypted data recovery, and mobile and cloud computing environments. Students will engage with cutting-edge research topics such as anti-forensic techniques, forensic analysis of blockchain technologies, and the challenges posed by privacy-preserving technologies (e.g., homomorphic encryption, differential privacy) in forensic investigations. Special attention will be given to the forensic implications of GDPR and other legal frameworks related to data privacy and cybersecurity. A core component of the course is its research orientation. Students will critically analyze contemporary forensic tools and methods, including the integration of artificial intelligence in automating forensic processes and detecting cyber threats in real-time. Practical case studies will highlight high-profile incidents involving sophisticated cyberattacks, offering insights into both technical solutions and legal ramifications. By the end of the course, students will have gained a thorough understanding of how digital forensics intersects with the broader fields of cybersecurity and digital privacy. They will be equipped to apply forensic methodologies in their research, contribute to the development of more secure digital systems, and propose solutions to ongoing challenges in the protection of sensitive data in the ICT landscape.
14	Advanced Metrology, Measurement and Data Acquisition Systems	12	secondo anno	The rapid advancement of AI is being driven by increasingly large and computationally intensive machine learning models and dataset thes. As a result, the amount of computing power used to train state-of-the-art models is growing exponentially, doubling approximately every 10 months between 2015 and 2022. This trend is leading to a significant carbon footprint. FL can also be resource-intensive and contribute to a substantial carbon footprint, especially when deployed at scale. Unlike centralized AI, which can access renewable energy at strategically located data centers, cross-device FL may utilize hundreds of millions of globally distributed end-user devices with diverse energy sources
15	Advanced modeling for engineering	4	secondo anno	This PhD-level course provides an in-depth exploration of advanced modeling techniques for engineering systems, with a focus on transport phenomena, multi-scale modeling, and hybrid modeling approaches. Students will gain a solid understanding of how to develop and analyze models that describe momentum, heat, and mass transfer processes across different spatial and temporal scales. Special emphasis is placed on hybrid modeling, integrating artificial neural networks (ANNs) with first-principles models to leverage both data-driven insights and physical laws. The course includes applications relevant to process, chemical, and food engineering, such as drying processes, reactive flows, filtration, fermentation, and heat/mass transfer in complex media. Through theoretical lectures and practical modeling exercises, students will learn how to construct, validate, and implement models that support decision-making and optimization in advanced engineering contexts.
16	Mathematical Optimization for Machine Learning - Foundations	8	secondo anno	Behind every successful machine learning model lies an optimization engine. In this course we introduce the fundamental mathematical tools and techniques to understand optimization algorithms crucial for training and deploying effective machine learning models. We focus on nonlinear optimization problems, i.e., optimization problems where the objective function or constraints, or both, are nonlinear. We cover topics like unconstrained and constrained optimization, convex optimization, gradient-based methods, Newton-type methods, and more.
17	Enabling Hardware Technologies For Sustainable AI Computing	4	terzo anno	Transformer-based neural networks have revolutionized artificial intelligence, driving innovation across business, robotics, transportation, education, and numerous other sectors. At their core, these models rely on self-attention mechanisms with quadratic computational complexity, requiring extensive matrix-vector operations that demand substantial computing resources. Large language models like GPT-3 exemplify these challenges, necessitating trillions of operations and enormous datasets that generate significant carbon footprints. As models continue to grow in size and complexity, they increasingly strain data center infrastructures and exacerbate the "memory wall" problem—a fundamental limitation in traditional von Neumann architectures where memory bandwidth and access speed constraints create bottlenecks, while simultaneously consuming vast amounts of energy during data transfer. This course explores innovative circuit and computer architecture solutions to these computational barriers. It focuses on Processing-in-Memory (PIM) architectures as a promising alternative to conventional computing approaches. By integrating computation directly within memory units, PIM-based HW solutions offer dramatic improvements in energy efficiency and sustainability for AI systems. Students will develop comprehensive understanding of these cutting-edge approaches to address the computational challenges facing modern AI systems.
18	Networking issues and perspectives for future generation of IoT systems	8	terzo anno	The recent advances in micro-electronics and in the computational capabilities of micro-devices opened a new perspective in the design of intelligent systems where things can work together to perform also complex tasks. The new era of Internet of Things (IoT) is already arrived and new networks and protocols design issues are arising in this evolving context. The classical protocol stack and network design criteria need to be changed to meet new network and application requirements and novel issues and threats need to be considered. The following course will present the main difference existing between a classical network and the emerging IoT technologies focusing, among others, also on QoS issues, distributed intelligence related issues, security threats, cyberphysical/physical actions issues that can be observed in these novel systems. Moreover, some references to novel application domains that can benefit by the IoT application will be illustrated, such as, for example, Cyber Physical Systems, networked robotics and swarm robotics, IoT-enabled automation, autonomous systems, v2x (vehicle to x) systems, e-health platforms, to offer a novel perspective in the IoT design and in the IoT market development.
19	Radio Systems for Aerospace	8	terzo anno	This PhD course explores advanced radio systems for aerospace applications, focusing on the design and development of phased arrays. The curriculum covers both antenna architectures and monolithic integrated circuits, addressing key challenges in system miniaturization, power efficiency, and beamforming techniques. Students will gain expertise in RFIC design, electromagnetic simulation, and array synthesis, with an emphasis on modern semiconductor technologies suitable for spaceborne and airborne communication and sensing systems.
20	Big Data Privacy	4	terzo anno	The course provides an introduction to big data and their processing with MapReduce. First, foundations of big data and NoSQL databases are provided. Here, big data characteristics and real-life applications are described, along with NoSQL databases and the reference Cloud Computing computational framework. Then, the course focuses the attention on the MapReduce processing model and its application to several big-data-processing-related application scenarios, such as database management. Case studies and examples, as well as state-of-the art systems and tools, are provided and discussed in details.
21	Advanced Strategies for Energy-Efficient Large Language Models	12	terzo anno	Large Language Models (LLMs) have revolutionized natural language processing (NLP) and understanding (NLU), enabling advancements in applications such as conversational AI, text classification, and machine translation. However, deploying these models efficiently remains a critical challenge due to their computational, memory and energy requirements. This course provides a comprehensive overview of cutting-edge techniques for optimizing LLMs, focusing on both training and inference efficiency, and their implications in energy-aware scenarios. The course begins with an introduction to LLM architectures and strategies for efficient model training and inference. It then explores Knowledge Distillation (KD), covering both white-box (logit-based, hint-based) and black-box (in-context distillation, chain-of-thought distillation) approaches. Advanced KD techniques such as Mini-LLM, Patient KD, Explainable AI (XAI)-driven KD, and KD via meta-learning are also examined. Next, transfer learning strategies are discussed, including pre-trained model adaptation and transferability estimation for LLMs. A key focus of the course is Parameter-efficient Fine-tuning (PEFT), a crucial technique for adapting LLMs to specific tasks with minimal computational cost. Various PEFT strategies are covered, including additive, selective, re-parameterized, and hybrid approaches, alongside practical applications of adapters and LoRA (Low-Rank Adaptation). By the end of the course, students will gain a deep understanding of advanced LLM optimization techniques, enabling them to train and deploy models efficiently for real-world applications, such as edge AI and efficient NLP and NLU. They will explore the trade-offs between energy efficiency and performance while developing the skills to balance computational constraints with model accuracy.
22	Knowledge Organization and Representation	4	terzo anno	This course aims to teach ways of formally representing knowledge using structured knowledge organization systems and, above all, ontologies, to enable the students to understand and effectively apply principles and techniques of representation and organization of knowledge currently used to index, classify and provide access to information resources. Lectures will be focused on both theory and methodology in the modelling of specialized domains, and in the use of representations for automatic classification and handling of information, and automated reasoning. During the course, emphasis will be placed on concepts and on the main features of ontologies, with appropriate attention to their application.
23	La responsabilità civile nell'uso delle nuove tecnologie ICT	4	terzo anno	La responsabilità civile, sia essa da inadempimento di un rapporto obbligatorio o da illecito extracontrattuale, e' oggetto di rivisitazione. Il ricorso alle tecnologie di ICT e di AI ha richiesto un intervento dell'Unione Europea. E' necessario indagare su nuovi criteri di imputazione del danno di tali problematiche.