METHODOLOGIES FOR THE FUTURE

9.30 Pietro Crovari, **Multi-Modal Conversational Agents: A New Era for Chatbots**

Conversational agents are becoming pervasive. More and more applications are using this technology to support users. Yet, they are not fully integrated with the underlying applications: users must communicate separately with the interface and with the conversational agent. In my research, I am investigating how create a new generation of conversational agents that can "observe" people's behavior and provide more effective, personalized, and multimodal support.

9:36 Alessio Bernardo, **Motus.ml Gives Data an Edge**

We live in a digital revolution that will generate more and more data over time. Due to memory and latency constraints, such data streams pose new challenges to ML systems that have to learn them in real-time. motus.ml introduces the first cloud-detachable AI. It harnesses the full potential of data delivering the frontier of advanced real-time analytics to all devices positioned at the edge of the network.

9:42 Stefano Samele, **Challenges of Unsupervised Machine Learning Techniques in Visual Industrial Anomaly Detection**

Detecting anomalies is of utmost importance in the industrial quality control process. In many industries, this is still achieved through an operator's visual and manual inspection. In this presentation, we will talk about recent advances in the automation of this task, discussing what artificial intelligence can achieve and current challenges that prevent state-of-the-art solutions from scaling in real-world scenarios.

9:48 Andrea Tocchetti, **AI Explainability Through Human-in-the-Loop Approaches**

The increasingly widespread use of Machine Learning in companies led to the development of models that achieved high performance on various tasks. Hence, companies have begun to face a novel problem. "Is the system working properly? Did it learn the right features? Did it learn the right features for the right reasons?". Explainability and Explainable AI strive to answer these questions and many others to improve humans' understanding of the logic and behaviour of models.

9:54 Diego Stucchi, **Change and Anomaly Detection in Multivariate Datastreams and Neural Networks**

Designing algorithms for automatically detecting changes is fundamental to identifying faults or anomalies in high-end manufacturing, where expert inspection is time-consuming, expensive, or poorly repeatable. In my research, I have designed new histogram-based CD algorithms for multivariate datastreams: Kernel QuantTree (KQT) that outperforms state-of-the-art solutions and MultiModal QuantTree (MMQT), which operates in a new setup where data comes from multiple distributions. Peculiarity of these algorithms is to provide a theoretically sound control of false alarm rate when the distribution of normal data is unknown. I have also investigated the application of detection techniques to design new inexpensive solutions for detecting faults in Neural Networks and in a concept-drift scenario.

10:00 Andrea Damiani, **Large Forests and Where to "Partially" Fit Them**

Data Scientists cannot ignore the irruption of Machine Learning (ML) into their research field. Indeed, models that can learn based on data examples have been gaining increasing attention, fueled by Big Data, because "More Data Beats a Cleverer Algorithm" in ML. Nevertheless, ML models have the terrible name of not being explainable nor efficient. Fortunately, balancing these benefits and drawbacks in ML modeling has recently become possible thanks to the renaissance of specialization that Information Technology is living, with Domain Specific Architectures (DSAs) increasingly often replacing general-purpose software solutions based on CPUs or gpGPUs. Unfortunately, designing, developing, programming, and deploying such DSAs on Application-Specific Integrated Circuit (ASIC) accelerators or, even worse, Field Programmable Gate Arrays' (FPGAs) programmable logic requires hardware design skills that rarely intersect with the area of expertise of Data Scientists and ML specialists. This talk tackles exactly this trade-off and how to build a Cloud-Fog-Edge toolchain to facilitate ML experts in the exploitation of the benefits of DSAs for hardware acceleration, without them having to master hardware design and development.

10:06 Alberto Zeni, **Paving the Way for Fast and Accurate Genome Analysis**

While typical workloads become more and more challenging, the semiconductor industry is battling with performance scaling and power efficiency on next-generation technology nodes. As a result, the industry is turning towards more customized compute architectures to help meet the latest performance requirements. In this talk, we'll explore the application of heterogeneous architectures such as GPUs and FPGAs, on genomics problems, showing how these architectures are exploited to solve problems previously considered impossible for general-purpose architectures.

10:12 Riccardo Nembrini, **Quantum Annealing Feature Selection for Ranking and Classification**

The growing availability of Quantum Computers in recent years has sparked research interest for practical and industrial applications. In particular, Quantum Annealing, one of the possible Quantum Computing paradigms, leverages quantum phenomena to rapidly search for solutions to optimization problems. As of today, we can use this technology to solve Machine-Learning-related problems, such as Feature Selection for Ranking and Classification, with results comparable to classical methods.

10:18 Simone Perriello, **Concrete Quantum Cryptanalysis of Code-Based Cryptosystems**

Code-based cryptosystems are among the most promising candidates for Post-Quantum Cryptography, as witnessed by NIST decision to advance all the code-based proposals to the last round of its Post-Quantum standardization initiative. By adapting Grover's framework, we designed concrete quantum circuits to attack such cryptoschemes, providing the first quantitative analysis of the required computational complexity. The proposed circuits solve a general combinatorial problem, since the problem addressed can be seen as a binary constraint satisfaction problem. In the process, we developed quantum circuits to permute matrix columns, to compute the Hamming weight of a given bitvector, and to perform Gauss-Jordan Elimination on a matrix.

10:24 Livia Lestingi, **Model-Driven Development of Formally Verified Human-Robot Interactions**

Soon, service robots will operate in unconstrained environments due to the significant presence of humans. This Ph.D. research introduces a model-driven framework based on formal methods to develop interactive robotic applications designed to handle the uncertainty of human behavior. Users formally model the human-robot interaction scenario, estimate the most likely outcome, and deploy the application. Collected traces constitute the data pool for a novel active automata learning algorithm to update the formal model of human behavior based on the accumulated knowledge. We validate the framework on realistic use cases from the healthcare setting.

10:30 Lorenzo Jr. Sabug, **A New Global Optimization Method for Complex Design Problems**

Have you ever been faced with a design problem, with too many knobs to tune? How would it be when your machine is too complex, you don't even know if your setting would even arrive at a good performance? Are you tired of doing tons of trial-and-error in design? Let us remove some headache from the design process, with our fast and fully automatic method called SMGO! Trial-and-error becomes trial-and-strategize, speeding up the search for the best version of your design, and most of all, without getting your hands too dirty!

10:36 Davide Croci, **Advanced Optimization Algorithms for Last Mile Logistics**

Due to the rapidly expanding e-commerce business, the Italian parcel delivery market has witnessed a boom in the last decade. This exponential volume growth, however, has posed unprecedented problems for both traditional last-mile carriers and shippers. Our research is concerned with the development of optimization algorithms to support last-mile parcel delivery companies in solving these new challenges. In the first step, we propose an efficient algorithm to solve the problem of districting for last-mile delivery. Secondly, we develop an effective algorithm for the pallet loading problem, a variant of the three-dimensional bin packing problem which includes practical constraints related to loading boxes onto pallets occupying a minimal amount of volume. Finally, we combine the previous projects to develop a new algorithm for the combined optimization of truck routing and loading.

10:42 Davide Salvi, **A Multimodal Approach to Face Multimedia Forensics Challenges**

Nowadays, creating and manipulating digital multimedia objects is within everyone's reach. Thanks to recent advances in deep learning and computer vision, anybody can perform stunningly realistic and convincing manipulations for each kind of media with almost no effort. On one side, this can open the door to new exciting possibilities. However, on the other hand, when these technologies are misused, they can have a negative social impact. This is the case of deepfakes, which have led to numerous cases of fraud, blackmail and fake news spreading. For this reason, the automatic analysis of the authenticity of multimedia material is becoming of paramount importance to prevent dangerous consequences. We propose to tackle this problem by considering an innovative approach from two points of view: multimodality and semanticity.

10:48 Riccardo Giampiccolo, **Multiphysics Modeling of Audio Systems in the Wave Digital Domain**

The performance of audio systems, such as audio transducers and sensors, strictly depends on the nonlinearities affecting the involved physical domains as well as on how those domains interact. In order to derive signal processing algorithms able to virtualize, compensate, or linearize the transducer behavior, multiphysics models must be derived. Our research thus focuses on the multiphysics modeling of audio systems, for both Virtual Analog (VA) and consumer electronics applications, with the purpose of deriving on the one hand efficient methods for the emulation of analog audio gear, while on the other hand efficient digital signal processing techniques for enhancing the acoustic performance of audio transducers, paving the way towards algorithms that leverage the physics of the device.