Ph.D. in Information Technology
Thesis Defense

February 18th, 2025
at 15:00
Room BIO1 – building 21

Alessandro FALCETTA – XXXVII Cycle

Privacy-Preserving Deep Learning: Algorithms, Technologies, and Ethics

Supervisor: Prof. Manuel Roveri

Abstract:

In recent years, the field of Machine Learning (ML) and Deep Learning (DL) has witnessed exponential growth, thanks to significant advances in hardware and the availability of large datasets. ML and DL algorithms, which form a crucial subset of Artificial Intelligence (AI), have demonstrated remarkable success across a range of domains such as computer vision, natural language processing, and time-series analysis. This widespread adoption has led to the rise of the "as-a-service" model, where
powerful ML and DL capabilities are offered to users without requiring technical expertise. However, this model introduces serious concerns regarding privacy, especially when sensitive user data is involved, such as in healthcare or personal information
processing. The need to balance the utility of these algorithms with stringent privacy requirements is becoming increasingly pressing. A promising solution to this challenge lies in the application of Homomorphic Encryption (HE), a cryptographic technique that enables computations on encrypted data. This allows third parties to process data without ever accessing the underlying plaintext, preserving user privacy. However, the integration of HE with ML and DL introduces significant obstacles. Current HE schemes are limited in the types of operations they can support and the depth of the computation pipelines they can handle. Additionally, the overhead introduced by performing computations on encrypted data is substantial, both in terms of time and memory requirements, which necessitates the development of optimized and scalable solutions. Beyond these technical challenges, the use of privacy-preserving techniques such as HE raises important ethical questions. Privacy is a critical ethical value, but it is not the only one. The design of AI systems must also consider other values like transparency, fairness, and beneficence. The process of learning from encrypted data presents unique ethical dilemmas, where the prioritization of privacy may come into conflict with these other values. This thesis proposes a comprehensive methodology that addresses the algorithmic, technological, and ethical challenges of privacy-preserving ML and DL. It begins by exploring the algorithmic limitations posed by HE and proposing solutions for adapting or redesigning existing ML and DL architectures to operate within these constraints. The thesis then tackles the technological challenge, introducing cloud-based platforms desIgned to manage the heavy computational demands of HE-based algorithms. Finally, it extends the ethical analysis beyond privacy, highlighting the trade-offs necessary for the responsible design of privacy-preserving AI systems. The thesis is organized into four main parts.

The introductory section provides the background on HE and the formulation of the research problem, setting the foundation for the subsequent sections. The second part addresses the algorithmic challenges, presenting new methodologies for building HE-compatible ML and DL models. The third part focuses on the technological aspects, proposing platforms that optimize the performance of these models in real-world environments. The fourth part delves into the ethical implications, discussing the intersection of privacy with other ethical considerations in AI system design. Finally, the thesis concludes with reflections on the current state of privacy-preserving ML and DL, and outlines directions for future research.

Massimo PAVAN – XXXVII Cycle

Effective and Adaptive Tiny Machine Learning

Supervisor: Prof. Manuel Roveri

Abstract:

In the last few years, Tiny Machine Learning (TinyML) has emerged as the branch of Machine Learning research that studies the execution of Machine and Deep Learning (MDL) models on devices extremely constrained in terms of memory, computational power, and power consumption, such as embedded and Internet-of-things devices. TinyML was deemed interesting both by academic and industrial communities because it enables the continuous execution of Machine Learning tasks on battery-operated devices for long amounts of time, opening up a wide variety of previously impossible use cases. Despite the promising results obtained so far, TinyML solutions are still limited in two main ways: i) they can not work effectively with complex types of data, such as radar and multi-frame video data, and ii) they are not meant to execute the learning phase of MDL algorithms on-device, assuming that this phase will be carried out on more powerful computers before the deployment. This thesis aims to introduce a methodology, named TinyWorks, as well as solutions implementing it to overcome the current limitations of the TinyML environment, with the objective of enabling a variety of new tasks and use cases on resource-constrained devices. Achieving this goal required re-designing the way in which TinyML solutions are organized and executed, giving high importance not only to the MDL algorithms but also to all the other components in the solutions that enable their on-device inference and learning. TinyWorks operates at two distinct levels. First, it addresses the design of inference-based Effective TinyML solutions for complex types of data. The application of TinyWorks has led to the first TinyML solutions working on Ultrawide-band Radar data and has enabled the effective analysis of multi-frame video data at a TinyML level. The feasibility of the proposed approach is demonstrated by porting the solutions directly on target Tiny devices. Second, the methodology has been used to design Adaptive TinyML solutions that are translated into real-world scenarios and use cases specific to the TinyML environment. Adaptive solutions are presented both for environments in which the distribution of the data-generating process changes just during deployment, named on-device single-change environments, and in environments in which it can change at any time during the operating life of the device, named non-stationary environments. Two novel adaptive TinyML solutions, designed using TinyWorks for these two environments, are presented. Their improved performance compared to the state-of-the-art in On-device Learning is

demonstrated through extensive experimentation. Furthermore, a toolbox for the design and generation of effective and adaptive TinyML solutions following the methodology is presented, to enable the seamless adoption of the methodology by the TinyML community. Finally, the ethical implications of applying effective and adaptive solutions to safety-critical scenarios, such as healthcare ones, are studied in the thesis, along with TinyWorks solutions specifically designed to address the ethical concerns.

PhD Committee

Prof. Giacomo Boracchi, Politecnico di Milano

Prof. Eiman Kanjo, Imperial College, London

Dr. Manish Gupta, MICROSOFT