Ph.D. in Information Technology
Thesis Defense

May 8th, 2026
at 2:30 pm
Room PT1 – building 20A

Luca COLOMBO – XXXVIII Cycle

From a Real Present to a Binary Future: Algorithmic and System Co-Design for Privacy-Preserving Deep Learning Training

Supervisor: Prof. Manuel Roveri

Abstract:

In recent years, the widespread adoption of Deep Learning (DL) as-a-service has revealed a fundamental trade-off between the utility of advanced Artificial Intelligence (AI) solutions offered by third-party providers and the increasingly stringent requirements for data privacy. A promising approach to resolve this tension lies in Homomorphic Encryption (HE), a cryptographic technique that enables computations to be performed directly on encrypted data, thereby preserving the confidentiality of the underlying information throughout its entire lifecycle. However, integrating HE into DL poses significant challenges. Most HE schemes impose strict constraints on supported operations and computational depth, while the substantial processing and memory overhead make training deep Neural Networks (NNs) prohibitively expensive. Consequently, prior research has primarily focused on encrypted inference, leaving the end-to-end encrypted training of DL models as a major open research challenge.

This thesis introduces ForgHE, a principled methodology for designing effective and efficient privacy-preserving solutions that enable the training of DL models directly on encrypted data while overcoming the technological limitations of HE. Rather than adapting cryptographic schemes to fit existing DL paradigms, ForgHE guides the systematic redesign of NN architectures and learning algorithms to be natively compatible with HE. The contributions of this thesis, structured around the proposed ForgHE methodology, address the core challenges from three complementary perspectives.

First, the foundational feasibility of end-to-end encrypted training is established through an algorithmic-cryptographic co-design approach. By reformulating standard learning procedures, novel integer-based and real-valued training pipelines are proposed to operate with the TFHE and CKKS schemes, respectively. These pipelines demonstrate that ForgHE enables the construction of DL solutions inherently compliant with the operational and numerical constraints of specific HE schemes.

Second, to mitigate the computational overhead of HE and minimize algorithmic-cryptographic mismatch from the outset, a paradigm shift toward binary-native learning is introduced. This approach aligns the training process with the most efficient Boolean primitives of the TFHE scheme. Specifically, a gradient-free learning rule and a fully binary error propagation mechanism are proposed, demonstrating that DL models can be trained effectively using only simple bitwise operations.
Finally, ForgHE addresses system-level scalability through a novel form of intra-model parallelization. By designing distributed learning frameworks equipped with HE-compatible local learning algorithms, the training workload of a single NN can be partitioned across multiple computing units, each responsible for training only a portion of the entire model. This strategy mitigates the prohibitive resource demands of monolithic execution and enables practical deployment.
The contributions of this thesis are validated through extensive experimental evaluations, demonstrating state-of-the-art performance and establishing a comprehensive, full-stack methodology for end-to-end privacy-preserving training of deep NNs on encrypted data. Collectively, these results pave the way for the development of advanced, privacy-by-design AI solutions in as-a-service settings, laying the foundation for a new generation of secure and trustworthy DL technologies

Matteo GAMBELLA – XXXVIII Cycle

Neural Architecture Search in Constrained and Dynamic Environments

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.
Despite their success, deploying DL solutions in real-world scenarios poses several challenges. These include hardware constraints (e.g., memory demand and computational requirements), functional restrictions (only a set of operations is permitted), robustness to perturbations, and the ability to adapt to dynamic environments. All these aspects make the design of efficient and highly accurate DL solutions a complex, time-consuming process that requires significant expertise.
A promising approach to automate and simplify this process is Neural Architecture Search (NAS), which designs efficient DNN architectures tailored to specific tasks. NAS significantly reduces human effort and accelerates development. While early NAS methods based on Reinforcement Learning required thousands of GPU hours, the field has since produced far more efficient approaches — including one-shot, zero-shot, training-free, and surrogate-assisted NAS — reducing search cost to under one

GPU hour in many settings. Despite this progress, current NAS methodologies still suffer from notable limitations: they often focus primarily on accuracy while neglecting other critical requirements, and fail to enforce the hard functional and technological constraints imposed by real deployment targets — such as strict memory budgets or restrictions on permissible operations. Crucially, achieving search efficiency simultaneously with constraint satisfaction is far from trivial: enforcing hard constraints during search fundamentally alters the structure of the optimisation problem, as the feasible architecture space may be sparse or disconnected, rendering standard gradient-based or sampling strategies unreliable.

Moreover, existing NAS methods often operate under static assumptions and lack the ability to deal with real-world dynamics such as distributional shifts, environmental changes, or input perturbations. While transfer learning, domain adaptation, and knowledge distillation partially address robustness, they operate on fixed architectures without reconsidering structural design choices. Incorporating robustness objectives into NAS introduces additional, often conflicting, optimisation criteria that cannot simply be appended to an existing efficient NAS pipeline without compromising search quality. In such contexts, architectures capable of dynamically adjusting their structure and computational behaviour in response to environmental changes become essential for sustained performance and reliability across changing conditions.

This thesis proposes a unified NAS-based framework designed to meet these emerging challenges by progressively addressing the multiple layers of complexity involved in real-world DL deployment. It begins by focusing on the technological and functional constraints typical of resource-constrained environments, proposing NAS strategies that remain effective even under strict hardware limitations. It then tackles the issues of generalization and robustness by enriching NAS with mechanisms that exploit the geometric and structural properties of neural architectures to improve resilience against data and hardware perturbations. Finally, it extends the NAS paradigm by introducing dynamic and adaptive strategies, enabling the evolution of DL architectures. What sets this work apart is its integrative approach: while efficiency, robustness, and adaptability are often optimized in isolation, this thesis addresses them in a unified and coherent framework. By gradually incorporating these layers into a single NAS methodology, the work lays the foundation for neural architectures that are not only accurate, but also practical, resilient, and truly deployable in dynamic real-world scenarios.

PhD Committee

Prof. Giacomo Boracchi, Politecnico di Milano

Prof. Elisabetta Fersini, Università degli Studi di Milano - Bicocca

Prof. Roberto Vezzani, Università di Modena e Reggio Emilia