Ph.D. in Information Technology
Thesis Defense

July 14th, 2026
at 4:00 pm
Room Schiavoni – building 20A

Isabella POLES – XXXVIII Cycle

Advancing Deep Learning Across the Medical Imaging Pipeline: From Disease Phenotyping to Treatment Optimization

Supervisor: Prof. Marco Domenico Santambrogio

Abstract:

Clinical decision-making encompasses a complex continuum spanning diagnosis, prognosis, treatment planning, and therapy delivery, where heterogeneous sources of information, including radiological imaging, histopathology, molecular measurements, and electronic health records, must be integrated to enable patient-specific decisions. *Routine* clinical workflows largely depend on manual assessment by radiologists, pathologists, and multidisciplinary teams, which remain affected by inter- and intra-observer variability, time constraints, and fragmented coordination across specialties.

Recent advances in Deep Learning (DL), fueled by the development of increasingly powerful neural architectures and large-scale training paradigms, have demonstrated remarkable potential across medical applications, including disease diagnosis and phenotyping, image segmentation, multimodal learning, and radiotherapy dose prediction and optimization. However, translating these advances into robust, system-level improvements in clinical practice remains challenging. Many DL approaches rely on large-scale annotated datasets due to their dependence on over-parameterized supervised learning, while their generalizability is often limited by disease heterogeneity and distribution shifts across patient populations, institutions, and acquisition protocols. Moreover, conventional multimodal architectures often assume the availability of complete, paired data, limiting their applicability in real-world scenarios with missing, unpaired, or heterogeneous modalities. Finally, the computational burden associated with gradient-based optimization over complex, non-convex learning landscapes remains a significant barrier, particularly for time-sensitive applications such as radiotherapy treatment planning. Together, these challenges continue to limit the reliability, scalability, and clinical integration of DL-based solutions.

In this context, this dissertation expands DL techniques for medical imaging for clinical *routine* and *research* practice along three complementary axes: (1) mitigating labelled data scarcity and capturing fine-grained and latent phenotypic variability through generative modeling and representation-learning strategies during diagnosis; (2) enabling robust multimodal phenotyping and prognostic modeling under incomplete or unpaired observations through knowledge distillation, context-aware fusion, and gradient-steering approaches; and (3) accelerating

radiotherapy workflows through integrated rigid and deformable image registration, multiscale tumor delineation, and learning-to-optimize treatment planning.
By addressing these challenges, the proposed methodologies would help move DL models one step closer to clinical translation, improving the accuracy, efficiency, and objectivity of diagnosis, prognosis, and therapy. Ultimately, these contributions would support emerging applications such as the discovery of novel imaging biomarkers, enhanced treatment personalization, and reduced time-to-treatment in precision medicine.

PhD Committee

Prof. Giacomo Boracchi, Politecnico di Milano

Prof. Michael Riegler, Simula Metropolitan

Prof. Bogdan Georgescu, Siemens Healthineers