Ph.D. in Information Technology
Thesis Defenses

April 29th, 2022
at 15:00
Room Conferenze "Emilio Gatti" and online by Webex

Davide AZZALINI – XXXIV Cycle

Detecting Anomalies in the Behavior of Autonomous Robots

Supervisor: Prof. **Francesco Amigoni**

Abstract:

Detection of anomalies and faults is a key element for long-term robot autonomy, because, together with subsequent diagnosis and recovery, it allows to reach the required levels of robustness and persistency. A fault which is not promptly detected and addressed, in fact, may result in the robot damaging itself or, even worse, in harming surrounding people. In this thesis, multiple approaches for detecting anomalous behaviors in autonomous robots starting from data collected during their routine operations are proposed. The main idea of this thesis is to model the nominal (expected) behavior of a robot and to evaluate how far the observed behavior is from the nominal one. Two approaches are proposed, based on Hidden Markov Models (HMMs) and Variational Auto-Encoders (VAEs), respectively. Both online and offline anomaly detection is tackled. Particular attention is devoted to ensuring that the proposed methods can be easily applicable in different practical settings. Accordingly, all the approaches proposed in this thesis are designed not to make any limiting assumption on how anomalies look like and to work with small amounts of (labeled) training examples. To this end, a data augmentation and retraining technique based on adversarial learning is also presented. We show how the methods proposed in this thesis positively compare against state-of-the-art anomaly detectors commonly used in robotics in a variety of application domains involving different robotic platforms required to operate for long periods of time without interruption.

Marco CANNICI – XXXIV Cycle

Learning Efficient and Effective Representations for Event-based Cameras

Supervisor: Prof. **Matteo Matteucci**

Abstract:

Event-based cameras are bio-inspired sensors that emulate the functioning of biological retinas. Unlike traditional cameras, which generate dense frames at a constant and predefined rate, these sensors, similarly to the photoreceptors in the retina, output data only when a change in brightness is detected. The result is a sensor able to sparsely and incrementally encode visual changes with microseconds resolution, high dynamic range, and minimum requirements for power consumption and bandwidth. Nevertheless, due to their fundamentally novel way of recording appearance, these sensors cannot be directly used with typical computer vision systems, which must be redesigned to work with events. This thesis addresses this challenge by focusing on three aspects of designing deep neural networks for event-based vision. First, we look at how to efficiently compute hidden neural representations by preserving event-based cameras' properties during computation. We accomplish this by designing a framework for converting deep neural networks into systems with identical expressiveness but capable of performing asynchronous and incremental processing, thus retaining the event camera's asynchronous and data-driven nature. Then, we focus on performance and study how to learn effective input representations for a given task. We propose a recurrent mechanism that automatically learns to interface with any convolutional network by sparsely and incrementally building a frame-like representation from asynchronous events. Finally, we focus on the challenging task of training neural networks to operate effectively on a real-world event-based camera when the only source of training supervision comes from simulation. We tackle the problem from a domain adaptation perspective by learning to extract domain-invariant intermediate representations. This learning strategy enables the network to attain performance comparable to that potentially achieved by directly learning from real annotated samples, yet without performing any finetuning on a real device. Throughout the thesis, we explore the importance of representations in event-based networks, at both the input and hidden layers, and show that by focusing on these aspects, considerable gains can be achieved toward more effective and efficient processing.

PhD Committee

Prof. **Andrea Bonarini,** Politecnico di Milano
Prof. **Guillermo Gallego,** Technische Universität Berlin
Dr. **Bruno Lacerda,** University of Oxford