Ph.D. in Information Technology: Thesis Defenses

February 18th, 2019

Room Conferenze "Emilio Gatti"– 11.30 am

Mirza RAMICIC – XXX Cycle

"Perception As Behaviour Inducing Mechanism: A Reinforcement Learning Perspective"

Advisor: Prof. **Andrea Bonarini**

Abstract:

Rapid advancement of machine learning make it possible to consider large amounts of data to learn from. In most of the implementations of reinforcement learning facing this type of data, approximation is obtained by neural networks and the process of drawing information from data is mediated by a short-term memory that stores the previous experiences for additional re-learning, to speed-up the learning process, mimicking what is done by people. In this work, we are proposing a range of novel computational approaches able to selectively filter the informational or cognitive load for the agent's short-term memory, thus simulating the attention mechanism characteristic also of human perception. Using the proposed attention filter block architecture, we were able to devise a variety of frameworks of agent's perception that are able to adapt to its environment by selecting the most suitable experiences. The adaptation also resulted in an emergence of different behavioural characteristics or traits among artificial learning agents.

Davide TATEO – XXX Cycle

"Building Structured Hierarchical Agents"

Advisors: Proff. **Andrea Bonarini and Marcello Restelli**

Abstract:

In modern real-world applications, autonomous agents are required to solve very complex tasks, using information taken from low-level sensors, in uncontrolled, dangerous, and unknown scenarios.

Among them, Robotics systems share some common characteristics: most of these systems use continuous state and action variables that may need a fine grain precision.

They may exhibit different dynamics between different parts of the system, leading to a natural division based on different abstraction levels. Finally, some tasks are even difficult to formalize in the framework of Reinforcement Learning, making difficult to define a reward function, while some human (or non-human) experts may be able to provide behavioral demonstrations.

Based on these assumptions, we propose two approaches to improve the applicability of Reinforcement Learning techniques in these scenarios: a new Hierarchical Reinforcement Learning framework based on the Control Theory framework, which is particularly well suited for robotics systems, and a family of Inverse Reinforcement Learning algorithms that are able to learn a suitable reward function for tasks (or subtasks) difficult to formalize as a reward function, particularly when the demonstrations come from a set of different suboptimal experts. Our proposals make it possible to easily design a complex hierarchical control structure and learn the policy both by interacting directly with the environment or providing demonstrations for some subtasks or for the whole system.

Carlo D'ERAMO – XXXI Cycle

"On the Exploitation of Uncertainty to Improve Bellman Updates and Exploration in Reinforcement Learning"

Advisor: Prof. **Marcello Restelli**

Abstract:

The issue of sample efficiency always constituted a matter of concern in Reinforcement Learning (RL) research. It is historically well-known that this issue arises from the need of the agent to explore the environment it is moving in to improve its knowledge about it, and to exploit simultaneously the actions it considers to be the best to maximize its return, creating a trade-off known in RL as exploration-exploitation dilemma. The addressing of this trade-off is central and constitutes a measure of effectiveness of any algorithm available in literature.

The purpose of this thesis is to study the previously described problems proposing novel methodologies that explicitly consider the concept of uncertainty to speed up learning and improve its stability. This solution is not new in RL research, but there is still a lot of work that can be done in this direction and this thesis takes inspiration from the available literature on the subject extending it with novel significant improvements on the state of the art. In particular, the works included in this thesis can be grouped into two parts: one where uncertainty is used to improve the behavior of the Bellman equation and the other where it is used to improve exploration.

All the works presented in this thesis are described, theoretically studied, and eventually, empirically evaluated on several RL problems. The obtained results highlight the benefits that the explicit exploitation of uncertainty in RL algorithms can provide; indeed, we show how in a large set of problems our methods prove to be more stable and faster to learn w.r.t. others available in the literature.

Alessandro RIVA – XXXI Cycle

"Development and Analysis of Algorithms for Information-Gathering in Autonomous Mobile Robotics"

Advisor: Prof. **Francesco Amigoni**

Abstract:

The gathering of information is a task that implicitly appears in a large number of autonomous mobile robotics applications, such as surveillance, search, and mapping. Despite the increasing attention received in recent years from both public institutions and private organizations, the development of proper algorithmic solutions for information-gathering applications is still particularly challenging, being most of the problems involved computationally hard. In this thesis, new efficient algorithmic solutions are provided. The design of each algorithm is supported by a thorough theoretical analysis and is validated by means of experimental simulations in realistic problem instances. The achievements constitute a solid base for the implementation and the deployment of real autonomous mobile robotics systems.

Diego CARRERA – XXXI Cycle

"Learning and adaptation to detect changes and anomalies in high-dimensional data"

Advisor: Prof. **Giacomo Boracchi**

Abstract:

Monitoring a datastream to detect whether the incoming data departs from normal conditions is a problem encountered in many important applications, ranging from quality control in industrial process to health monitoring. Many solutions in the literature adopt a model that describes normal data and check whether this model is not able to describe new data, detecting anomalous instances or permanent changes affecting the data-generating process. Pursuing this approach is challenging when data have high dimensions or feature complex structures (as in case of images or signals), that is the typical situation in real world scenarios. We address this problem from two different perspectives, depending on the assumptions on the data-generating distribution. At first, we assume that data can be described by a smooth probability density function and we focus on the change-detection problem, where the goal is to detect permanent changes affecting the data-generating process. We show the detectability loss phenomenon, namely that performance of a popular change-detection algorithm that monitors the likelihood decreases when the data dimension increases. Moreover, we propose QuantTree, a novel algorithm to define histograms as density models for high dimensional data that enables the non-parametric monitoring of datastreams.

In the second part we focus on data that feature complex structures, that cannot be described by a smooth probability density function. We adopt dictionaries yielding sparse representations to characterize normal data and propose a novel anomaly-detection algorithm that detects as anomalous any data that do not conform to the learned dictionary. To make our anomaly-detection algorithm effective in practical application, we propose two domain adaptation algorithms that adapt the anomaly detector when the process generating normal data changes. The proposed algorithms have been successfully tested in two real world applications: a quality inspection system monitoring the production of nanofibrous materials through the analysis of Scanning Electron Microscope (SEM) images, and ECG monitoring using wearable devices.

PhD Committee:
Prof. **Francesco Amigoni**, DEIB
Dr. **Alessandro Lazaric**, INRIA, Parigi
Dr. **Alessandro Giusti**, IDSIA, Lugano