POLITECNICO
MILANO 1863

Thesis Proposals

N. Gatti, M. Restelli, F. Trovò, A. Marchesi, A.M. Metelli, M. Castiglioni, M. Papini

Fall 2023

Foundational Research Theses

Machine Learning

Detection and forecast of Tropical Cyclones with machine learning: is it possible to forecast the occurrence of a tropical cyclone through meteorological features? (CLINT – European Project)

Given a set of measurements of meteorological variables at different scales, this work aims to perform supervised learning classification to predict the probability of occurrence of tropical cyclones at different locations. This work is mainly applicative and builds upon preliminary work done in the last months on this topic.

Supervisors: Marcello Restelli, Alberto Maria Metelli

Co-supervisors: Paolo Bonetti

Expected graduation time: 9-12 months

Start: now

Contacts: paolo.bonetti@polimi.it

Radiogenomics to predict HER2 expression and tailor treatments in Breast Cancer patients (collaboration with Istituto Nazionale Tumori, INT)

Novel treatments targeting HER2 protein revolutionized the landscape of therapies and dramatically improved the prognosis of patients affected by breast cancer in the advanced stage. However, reproducible and easy-to-access biomarkers are lacking in the clinics, leaving uncertainty about the optimal use of anti-HER2 treatments. We will apply ML models of radiomic data to build a solid ML model to predict HER2 status and treatment outcomes from CT scan images of metastatic breast cancer patients.

Supervisors: Francesco Trovò

Co-supervisors: Medical Personnel from INT

Expected graduation time: 9-12 months

Start: now

Contacts: francesco1.trovo@polimi.it

Machine learning for predicting immunotherapy resistance in melanoma and NSCLC patients using digital pathology images (I3LUNG European project)

Melanoma and Non-small cell lung cancer (NSCLC) are tumors initially responsive to immunotherapy, but patients often develop resistance to the treatment. Current clinical biomarkers are not enough to predict this resistance. Machine Learning is revolutionizing the analysis of medical images, including digital pathology images, revealing clinically relevant information hidden within these images. This technological advancement is paving the way for the discovery of novel biomarkers. The objective of this thesis is to develop predictive models to predict immunotherapy resistance using digital pathology images and clinical data. The student will work with commonly used architectures in the field of digital pathology, such as attention-based models and Generative Adversarial Networks (GANs).

Supervisors: Francesco Trovò

Co-supervisors: Medical Personnel from INT

Expected graduation time: 9 months

Start: now

Contacts: francesco1.trovo@polimi.it

Reinforcement Learning

Increasing efficiency in policy search algorithms

Policy search algorithms have found extensive use across a broad spectrum of domains, often achieving astonishing performance. Nevertheless, these approaches are data-hungry, and a significant number of interactions need to be collected within the environment to obtain reasonable results. This research delves into enhancing the sample efficiency of policy search techniques. The problem will be tackled both from theoretical and practical perspectives.

Supervisors: Marcello Restelli, Alberto Maria Metelli

Co-supervisors: Riccardo Poiani

Expected graduation time: 9-12 months

Start: now

Contacts: riccardo.poiani@polimi.it

Variance Reduction in Policy Gradient Methods

Policy Gradient is one of the most versatile methods in Reinforcement Learning. In the last five years, several variance-reduction techniques from the stochastic optimization literature have been applied to Policy Gradient to improve its sample complexity (that is, learn faster). The theoretical improvement is significant, but the effectiveness of this approach in practice is nowhere near to the original (supervised) setting. One reason is that these theoretical results rely on unrealistic assumptions. The goal of this thesis is to remove or weaken these assumptions and develop more robust algorithms.

Supervisor: Matteo Papini

Expected graduation time: 9-12 months

Start: now

Contacts: matteo.papini@polimi.it

Deterministic Policy Optimization

In general, exploration is fundamental for online Reinforcement Learning. In realistic tasks where function approximation (for instance, with neural networks) is necessary, random actions are, by far, the most used exploration device. However, in some applications, random actions may be unacceptable for safety or ethical reasons. So, it would be of great value to characterize those problems for which exploration is not strictly necessary. Previous works on contextual linear bandits (a sub-problem of RL with function approximation) show that such problems do in fact exist. The goal of this thesis is to obtain similar guarantees for the full Reinforcement Learning setting and develop deterministic policy optimization algorithms.

Supervisors: Matteo Papini

Expected graduation time: 9-12 months

Start: now

Contacts: matteo.papini@polimi.it

Inverse Reinforcement Learning

Inverse Reinforcement Learning (IRL) refers to a class of algorithms designed to address the Imitation Learning problem aiming at deriving a representation of the objective function optimized by an expert agent. This problem involves using an expert's trajectories to learn the optimal strategy. However, IRL algorithms have not yet found widespread practical application. This research aims to gain new insights into IRL from a theoretical perspective and, potentially, to identify new real-world applications.

Supervisors: Alberto Maria Metelli, Marcello Restelli

Co-supervisors: Filippo Lazzati

Expected graduation time: 9-12 months

Start: now

Contacts: filippo.lazzati@polimi.it

Policy Gradient methods with adaptive policy spaces

Policy gradient methods constitute a category of reinforcement learning approaches, wherein the optimization of parameterized policies is achieved through gradient descent techniques to maximize the expected return. However, when the complexity of the parameteric policy function grows, the effectiveness of such methods often diminishes significantly. Consequently, deploying these techniques in practical real-world situations remains challenging. In this thesis, we propose a novel approach to solve this issue. Our solution is based on the idea of automatically choosing and adapting the class of policies during the learning process, splitting iteratively the policy space and thus deriving new policies at each iteration.

Supervisor: Marcello Restelli, Alberto Maria Metelli, Matteo Papini

Co-supervisors: Gianmarco Tedeschi

Expected graduation time: 9-12 months

Start: now

Contacts: gianmarco.tedeschi@polimi.it

Reinforcement Learning for Network Infrastructures (AI4REALNET – European Project on AI)

AI4REALNET aims to successfully create an overarching multidisciplinary approach, combining emerging AI algorithms, existing open-source AI-friendly digital environments – to test and benchmark AI in industry-relevant use cases – socio-technical design of AI-based decision systems and human-machine interaction (HMI), to better operate network infrastructures in real-time and predictive mode. The goal of the thesis consists in developing new approaches combining distributed and hierarchical RL techniques for the operation of large-scale network infrastructures. To address the complexity (and curse of dimensionality) of such scenarios, two approaches are considered: i) hierarchical task (temporal) decompositions (e.g., the option framework) and ii) multi-agent RL, including spatial and/or temporal decompositions and independent learning of the agents vs joint learning with communication.

Supervisor: Marcello Restelli, Alberto Maria Metelli

Co-supervisors: Marco Mussi, Gianvito Losapio

Expected graduation time: 9-12 months

Start: in a few months

Contacts: albertomaria.metelli@polimi.it

Online Learning

Online Learning for Linear MDPs with Constraints (FAIR – PNRR Excellence Project on AI)

Fascinating AI applications, such as autonomous driving and financial trading, require learning sequentially in multi-state environments in the presence of constraints; these problems are customarily modeled as Constrained Linear MDPs. The thesis aims to develop innovative techniques in order to deal with Linear Markov Decision Processes when the agent has to learn the constraints online. The techniques developed in the thesis will be then evaluated in the context of the aforementioned applications.

Supervisors: Nicola Gatti, Matteo Castiglioni, Alberto Marchesi

Co-Supervisor: Francesco Emanuele Stradi

Expected graduation time: 6-9 months

Start: now

Contacts: francescoemanuele.stradi@polimi.it

Online Learning in MDPs with Hard Constraints (FAIR – PNRR Excellence Project on AI)

Investigating how constraining the learning algorithms' exploration is one of the hottest topics in machine learning. In particular, learning in safety-critical scenarios, such as robotics, healthcare, and trading, presents unique challenges compared to standard Online Learning since exploration should obey some safety constraints. The thesis aims to model sequential decision-making scenarios as MDPs with hard constraints and to develop innovative techniques tailored for safety-critical contexts. The techniques will be evaluated in some applications of high practical interest.

Supervisors: Nicola Gatti, Matteo Castiglioni, Alberto Marchesi

Co-Supervisor: Francesco Emanuele Stradi

Expected graduation time: 6-9 months

Start: now

Contacts: francescoemanuele.stradi@polimi.it

Algorithmic Game Theory and Multi-Agent Learning

Sequential Information Selling (ELIAS – European project on AI)

Information plays a crucial role in AI for democracy. In particular, nowadays, there is a terrific amount of information being collected on digital platforms. As a result, the problem of optimally selling information is of paramount importance for information holders. The objective of this thesis is to study the problem of selling information over a sequential decision-making problem, where one has to decide which information to sell and when. The thesis will investigate how to extend the model to an online learning scenario in which the seller does not know some of the parameters of the environment.

Supervisors: Nicola Gatti, Matteo Castiglioni, Alberto Marchesi

Co-Supervisors: Federico Cacciamani

Expected graduation time: 6-9 months

Start: now

Contacts: matteo.castiglioni@polimi.it

Online Information Structure Design (FAIR – PNRR Excellence Project on AI)

The possibility of altering or manipulating the beliefs of human and/or virtual agents may have a terrific impact on AI for democracy. For instance, in several scenarios, an agent (called leader) like a company has information that influences the utility of other agents (called followers) that are unable to access this information. As a consequence, the goal of the leader is to determine how much information to disclose to maximise revenue. In this thesis, we focus on the problem of computing how much information the principal should reveal in an online manner. This is crucial in real-world scenarios where the principal starts with no prior knowledge about the follower but gradually acquires it during their interactions. The algorithms will be evaluated in synthetic settings generated from real-world applications.

Supervisors: Nicola Gatti, Matteo Castiglioni, Alberto Marchesi

Co-Supervisors: Francesco Bacchiocchi

Expected graduation time: 6-9 months

Start: now

Contacts: francesco.bacchiocchi@polimi.it

Deception in Information Structure Design (ELIAS – European project on AI)

Information structure design is a novel research area aiming at prescribing which information should be revealed, when, and to whom. Controlling such a task will be crucial in AI for democracy. Classical tools of information structure design are based on the assumption that the information signaler is trusted, but this is not the case in many applications. Understanding the best signaler's deception is a challenging problem that can be faced resembling adversarial machine learning. In this thesis, techniques will be developed to maximize the signaler's deception and they will be evaluated in real-world situations.

Supervisors: Nicola Gatti, Matteo Castiglioni, Alberto Marchesi

Co-Supervisors: Federico Cacciamani

Expected graduation time: 6-9 months

Start: now

Contacts: matteo.castiglioni@polimi.it

Online MDPs Configuration (ELIAS – European project on AI)

Many intriguing AI applications, such as financial Market, sport betting and Poker encompass strategic interactions between interested parties. Developing algorithms capable of dealing with these scenarios is a central problem in AI. In particular, Configurable MDPs model scenarios where a learner and a configurator interact in an online fashion. The thesis focuses on the development of novel learning techniques for Configurable MDPs and their application to real-world problems.

Supervisors: Nicola Gatti, Matteo Castiglioni, Alberto Marchesi

Co-Supervisors: Francesco Bacchiocchi, Francesco Emanuele Stradi

Expected graduation time: 6-9 months

Start: now

Contacts: francesco.bacchiocchi@polimi.it, francescoemanuele.stradi@polimi.it

Industrial Research Theses

Developement and Simulation of Trajectory Planning Algorithms for a Swarm of Drones (with MBDA)

This thesis is part of a joint research project between MBDA (a leading European Aerospace and Defense company) and Politecnico's AI Research & Innovation Center (AIRIC). The thesis focuses on the Trajectory Planning problem, which consists of finding an optimal trajectory from a start to an endpoint. The aim is to develop and compare state-of-the-art real-time Trajectory Planning algorithms for a swarm of drones, employing Matlab and Simulink for simulation and experimentation.

Supervisors: Nicola Gatti, Marcello Restelli

Co-Supervisors: Pierriccardo Olivieri, Tommaso Bianchi, Simone Drago, Piergiuseppe Pezzoli

Expected graduation time: 6-12 months

Start: now

Contacts: pierriccardo.olivieri@polimi.it, tommaso.bianchi@polimi.it, simone.drago@polimi.it, piergiuseppe.pezzoli@polimi.it

Developement and Simulation of Formation Control Algorithms for a Swarm of Drones (with MBDA)

This thesis is part of a joint research project between MBDA (a leading European Aerospace and Defense company) and Politecnico's AI Research & Innovation Center (AIRIC). The issue to address regards the Formation Control problem for a swarm of drones where the objective is to maintain a specific spatial configuration between the drones while following a trajectory. The objective is to develop and compare state-of-the-art real-time Formation Control algorithms for a swarm of drones, employing Matlab and Simulink for simulation and experimentation.

Supervisors: Nicola Gatti, Marcello Restelli

Co-Supervisors: Pierriccardo Olivieri, Tommaso Bianchi, Simone Drago, Piergiuseppe Pezzoli

Expected graduation time: 6-12 months

Start: now

Contacts: pierriccardo.olivieri@polimi.it, tommaso.bianchi@polimi.it, simone.drago@polimi.it, piergiuseppe.pezzoli@polimi.it

Dynamic Pricing in Multi-state Environments (with lastminute.com)

Online Learning techniques applied to dynamic pricing often assume the environment is single state. This thesis aims to extend the framework when the objective is to price complementary products (as a flight and the insurance associated), and the environment can be modeled as an MDP. This is an industrial thesis that will be developed in collaboration with LastMinute.com.

Supervisors: Nicola Gatti, Marcello Restelli

Co-Supervisor: Francesco Emanuele Stradi

Expected graduation time: 6-9 months

Start: now

Contacts: francescoemanuele.stradi@polimi.it

Automated Market Monitoring (with Vehnta)

This thesis is part of a joint research project between Vehnta (italian marketing startup) and Politecnico's AI Research & Innovation Center (AIRIC). The aim is to develop models and algorithms for the quantitative analysis of the level of interest for a specific product/service and for pinpointing the geographical areas such interest is coming from, using data from advertising campaign in an online learning scenario.

Supervisors: Nicola Gatti, Marcello Restelli

Co-Supervisors: Francesco Bacchiocchi, Tommaso Bianchi, Luca Alessandrelli

Expected graduation time: 6-9 months

Start: now

Contacts: francesco.bacchiocchi@polimi.it, tommaso.bianchi@polimi.it, luca.alessandrelli@polimi.it

Algorithmic Startup Intelligence (with Lutech Spa)

This thesis is part of a joint research project between Lutech Spa (italian leading ICT company) and Politecnico's AI Research & Innovation Center (AIRIC). The aim is to develop models and algorithms for the analysis and predictions of KPIs estimating the likelihood of success of entrepreneurial initiatives from unstructured market data.

Supervisors: Nicola Gatti, Marcello Restelli

Co-Supervisors: Tommaso Bianchi

Expected graduation time: 6-9 months

Start: ~October 2023

Contacts: tommaso.bianchi@polimi.it

Generative AI applications

The aim of this thesis is to study techniques and algorithms to make Generative AI models more effective for specific application-related tasks, such as in-domain question answering, constrained generation and personalized recommendations. The work will be part of the joint research projects between Politecnico's AI Research & Innovation Center (AIRIC) and partner companies.

Supervisors: Nicola Gatti, Marcello Restelli

Co-Supervisors: Tommaso Bianchi, Luca Alessandrelli

Expected graduation time: 6-9 months

Start: ~November 2023

Contacts: tommaso.bianchi@polimi.it, luca.alessandrelli@polimi.it

AI Algorithms for Photovoltaic System Maintenance

PV systems are subject to different kinds of faults that reduce the amount of produced power or that can be harmful. This thesis is the continuation of a project in collaboration with Ricerca Sistema Energetico (RSE) which goal is to create an AI based solution for Fault Detection and Diagnosis in PV Plants. The student will expand the developed algorithms considering new types of faults like soiling detection and hotspots that require a different approach in what is done until now.

Supervisors: Nicola Gatti, Marcello Restelli

Co-supervisors: Alessandro Lavelli

Expected graduation time: 6-8 months

Start: now

Contacts: alessandro.lavelli@mlcube.com

Boosting Renewable Energy Sector with Learning Techniques (with RSE)

Nowadays, the energy sector is under the spotlight of governments and industries for the climate changing problem and the urgent need to study different renewable energy sources. What we can do as AI experts is to elaborate solutions to optimize usage and management of produced energy in a smart way, employing suitable Machine Learning frameworks. In this thesis, you will work in collaboration with RSE to extend the functionalities of a Digital Twin modeling Energy Storage Systems, by employing and designing proper Online Learning and Reinforcement Learning solutions. With your research, we want to pave the way for a constantly growing integration of the old electrical infrastructure with modern AI techniques.

Supervisors: Nicola Gatti, Marcello Restelli

Co-supervisors: Davide Salaorni

Expected graduation time: 9 months

Start: now

Contacts: davide.salaorni@polimi.it

Distributed RL for Industry 4.0 (with Siemens)

In a production system, several robots need to act concurrently to be able to go over the production line while optimizing for some efficiency scores (energy, time, etc.), often with just partial knowledge of what the others are doing. Industry 4.0 is then an interesting testbed for distributed decision-making problems. Thus, provided a digital twin of a factory plant with several robots and products to be produced, the objective is to leverage state-of-the-art RL algorithms in a distributed fashion, for solving the problem of cooperation with decentralized decisions and information sharing.

Supervisors: Marcello Restelli

Co-supervisors: Riccardo Zamboni

Expected graduation time: 9 months

Start: now

Contact: riccardo.zamboni@polimi.it

Reinforcement Learning for Financial Trading Support (with Intesa Sanpaolo)

Trading machines based on Reinforcement Learning attract a rich profusion of interest, but suffer from the noisiness of the financial movements in the short terms. On the other hand, human traders can have more effective vision on future market behaviors, but have a slower reaction time in the detection of profitable patterns. The goal of this thesis is to investigate and develop reinforcement learning algorithms for supporting human traders to detect sudden anomalies and profitable patterns.

Supervisors: Marcello Restelli

Expected graduation time: 9-12 months

Start: flexible

Contacts: marcello.restellipolimi.it

6 months Internship Theses @ ML cube

Article Generation using Large Language Models

This thesis consists in applying Generative AI to automatically write newspaper articles for a specific context. Besides requiring prompt engineering this kind of task requires the implementation of adequate safeguards, in order to avoid LLMs producing hallucinations, and to alert the user whenever the quality of the user is below expectations.

Supervisors: Nicola Gatti, Marcello Restelli

Co-supervisors: Lorenzo Bisi

Expected graduation time: 9 months

Start: now

Contacts: lorenzo.bisi@mlcube.com

Roof surface prediction for the installation of photovoltaic panels

This thesis is about the development of AI algorithms to perform a forecast of the maximum number of solar panels given the image of a roof. To do that, the algorithm should detect the roof surface, possible obstacles, and occlusions. The goal is to support companies in designing a first draft of the solar panel plant by using computer vision to detect the free portion of the roof, and by optimizing the panels layout, once given the available surface.

Supervisors: Nicola Gatti, Marcello Restelli

Co-supervisors: Lorenzo Bisi

Expected graduation time: 9 months

Start: now

Contacts: lorenzo.bisi@mlcube.com

Drift Explainability feature for the ML cube Platform

The ML cube platform provides a tool for automatic monitoring and re-training of Machine Learning models. However, knowing that a drift has happened may be not sufficient for taking the right corrective action. The goal of this thesis is to develop a set of explainability tools to allow ML cube Platform customers to have a deeper understanding of the kind of non-stationarity their models are dealing with.

Supervisors: Nicola Gatti, Marcello Restelli

Co-supervisors: Lorenzo Bisi, Alessandro Lavelli

Expected graduation time: 9 months

Start: now

Contacts: lorenzo.bisi@mlcube.com