POLITECNICO
MILANO 1863

Thesis Proposals

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

September 2022


"It's hard to think of a single technology
that will shape our world more in the next
50 years than Artificial Intelligence"
Barack Obama

Programma strategico

Intelligenza Artificiale

2022-2024

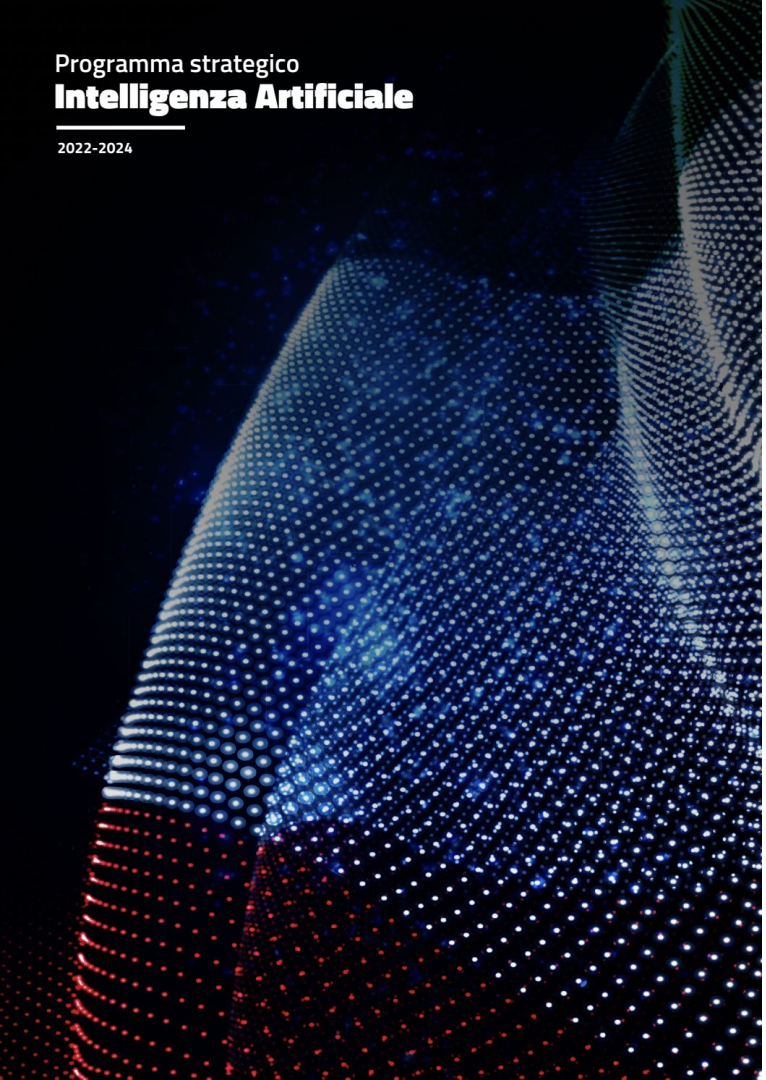

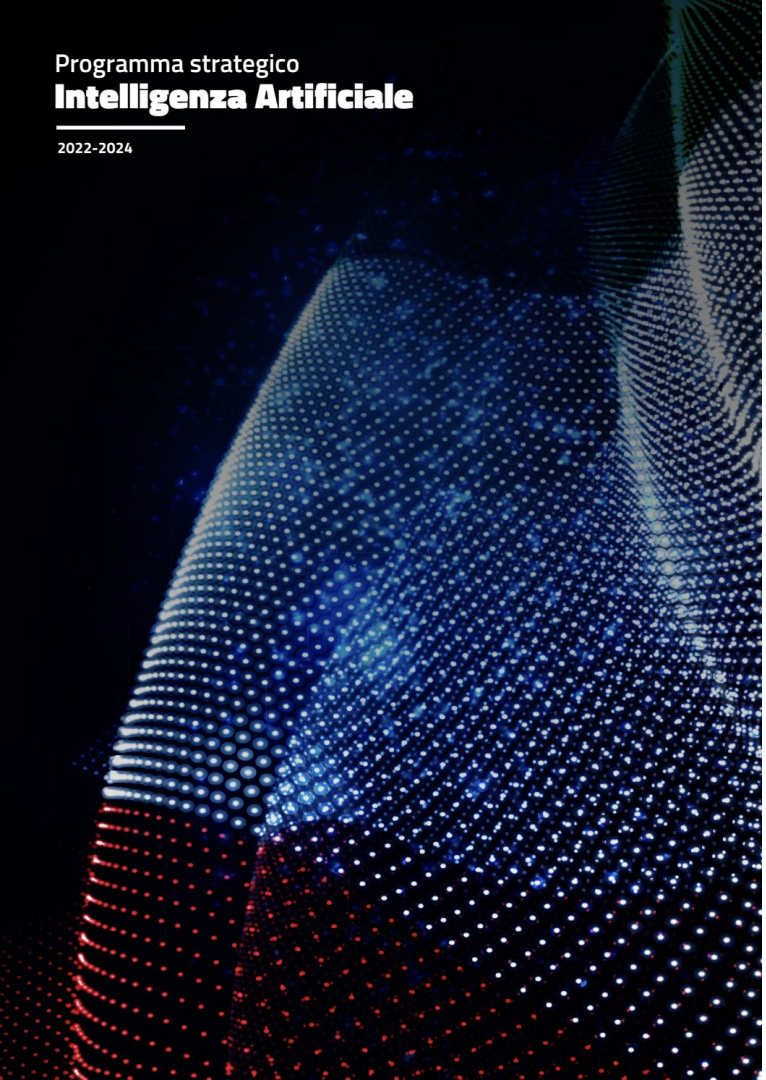

PNRR (Recovery Plan): 160 million Euros
PoliMi: Excellence Center for Machine Learning

AI Institutional Initiatives @Polimi

Foundational Research
(TRL 1-3)

ellis
European Laboratory for Learning and Intelligent Systems

AI Institutional Initiatives @Polimi

Industrial Research
(TRL 3-7)

AIRIC

Foundational Research
(TRL 1-3)

ellis
European Laboratory for Learning and Intelligent Systems

AI Institutional Initiatives @Polimi

Market Research

Industrial Research
(TRL 3-7)

Foundational Research
(TRL 1-3)

AI Institutional Initiatives @Polimi

Market Research — osservatori.net digital innovation

Industrial Applications (TRL 5-9) — ML cube

Industrial Research (TRL 3-7) — AIRIC

Foundational Research (TRL 1-3) — ellis European Laboratory for Learning and Intelligent Systems

Research Team

Nicola Gatti
Associate Professor

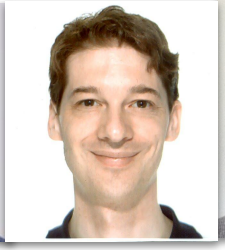

Marcello Restelli
Associate Professor

Francesco Trovò
Assistant Professor

Alberto Marchesi
Assistant Professor

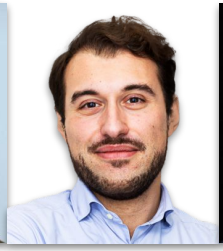

Alberto M. Metelli
PostDoc Researcher

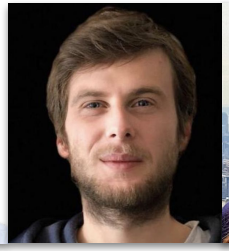

Matteo Castiglioni
PostDoc Researcher

Giulia Romano
PhD Student

Mirco Mutti
PhD Student

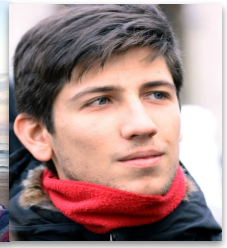

Luca Sabbioni
PhD Student

Amarildo Likmeta
PhD Student

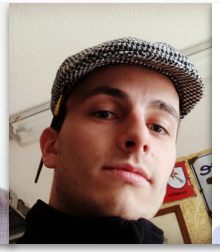

Pierre Liotet
PhD Student

Marco Mussi
PhD Student

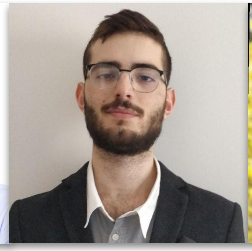

Federico Cacciamani
PhD Student

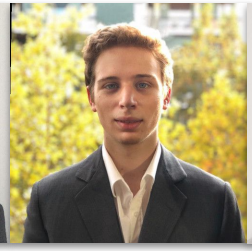

Martino Bernasconi
PhD Student

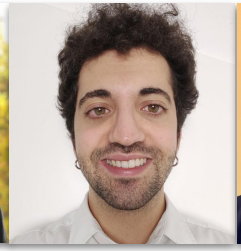

Alessio Russo
PhD Student

Gianluca Drappo
PhD Student

Riccardo Poiani
PhD Student

Paolo Bonetti
PhD Student

Research Team

Jacopo Germano
PhD Student

Riccardo Zamboni
PhD Student

Luca Carminati
PhD Student

Francesco Bacchiocchi
PhD Student

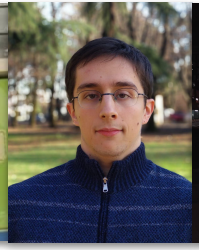

Davide Maran
PhD Student

Francesco E. Stradi
Research Intern

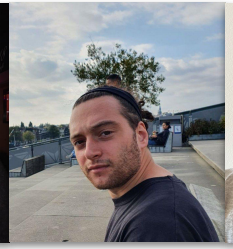

Pierriccardo Olivieri
Research Intern

Antonio Riva
Research Intern

AIRIC

Tommaso Bianchi
AI Engineer

Luca Alessandrelli
AI Engineer

Gianmarco Genalti
AI Scientist

Matteo Sacco
AI Scientist

ML cube

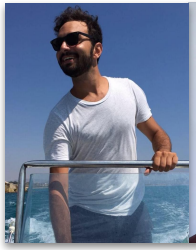

Alessandro Nuara
CTO

Alessandro Lavelli
ML Engineer

Lorenzo Bisi
AI Specialist


ellis
European Laboratory for Learning and Intelligent Systems
MILAN
HOMEPAGE
PEOPLE
SPONSORSHIP
NEWS
1863
POLITECNICO
DI MILANO
ELLIS Unit Milan
European Laboratory for Learning and Intelligent Systems

What is a thesis in foundational research?

Target: a scientific result published in a paper

Living and working in international top-ranked research team

Examples:

- Novel algorithms and experimental analysis
- Novel algorithms and theoretical properties
- Both

Research areas

Machine Learning (ML)

Reinforcement Learning (RL)

Online Learning (OL)

Game Theory (GT)

CLINT: Detection (ML)

Supervised: Paolo
Expected time for graduation 9-12 months
Start: now
Contact: **paolo.bonetti@polimi.it**

Context: 4-year H2020 research project involving 13 partners from France, Germany, Netherlands, Spain, Sweden, UK

Goal: extreme event (droughts) detection using **machine learning**

Data: temporal and spatial high-resolution climate data

Techniques: feature selection, dimensionality reduction, supervised multi-task (deep) learning

Website: https://climateintelligence.eu

Adversarial Attacks in Reinforcement Learning (RL)

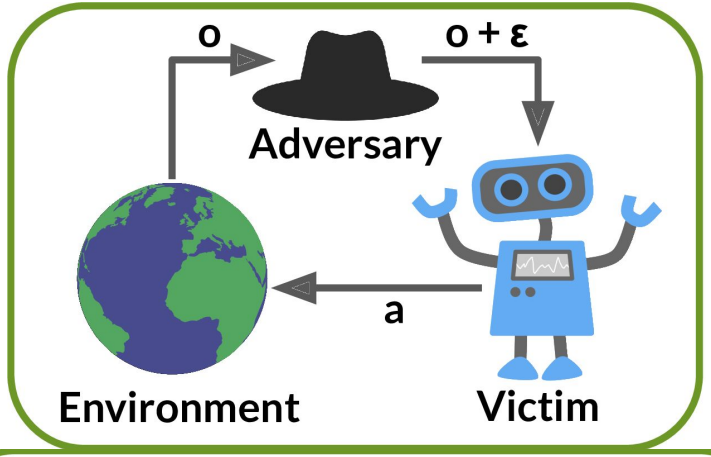

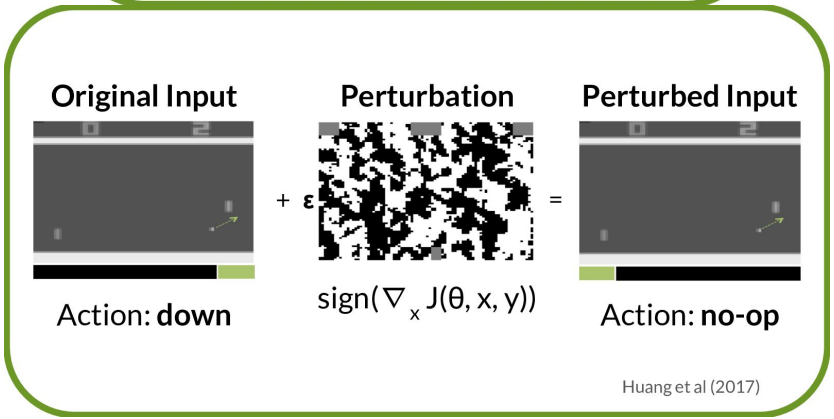

- **Adversarial RL**: Generate an adversarial noise on the state on an MDP in order to fool an agent into making the wrong actions
- **Goal**: answer the questions
  - Which environments/policies are more vulnerable?
  - Can we learn a robust policy in this sense?

Supervised: Davide e Pierriccardo
Expected time for graduation 9-12 months
Start: now
Contact: **davide.maran@polimi.it**

Real-Time Bidding (RL)

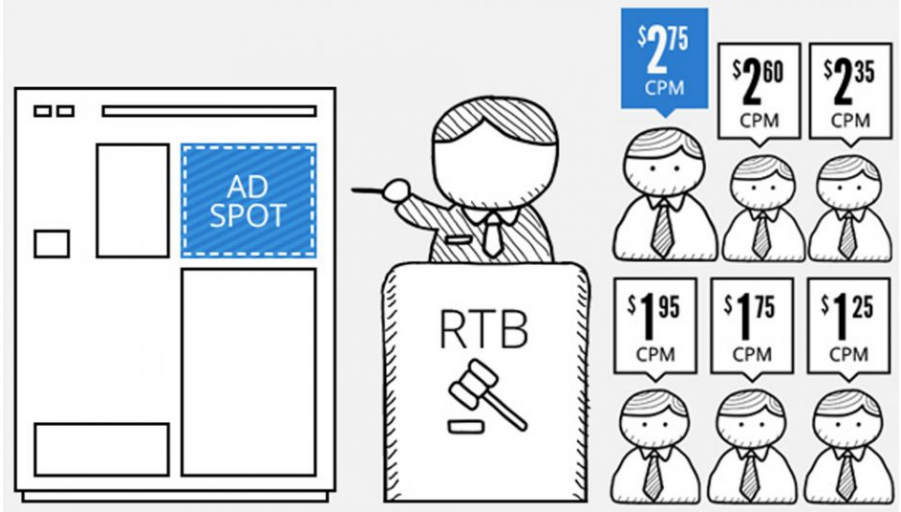

- Every time an **online advertising** slot is available, an auction is performed to assign it
- The auction is performed in real-time and the advertiser which offers more get the impression
- The objective of this thesis is to design and implement an efficient **Real-Time bidding algorithm** based on **Deep Reinforcement Learning techniques**

Supervised: Marco, Gianmarco
Expected time for graduation: 8-12 months
Start: now
Contact: **marco.mussi@polimi.it**

Online Machine Learning for Joint Pricing & Advertising (OL)

Selling a product online requires nowadays to manage multiple constraints and complex objective functions

Dynamic Pricing → Selling a product by continuously varying the price in function of market variables

Online Advertising Optimization → Determine the optimal budget allocation between ad campaigns

Goal of the thesis:

1. Identify a particular framework where the interaction between pricing and advertising is not trivial
2. Design, validation and testing of a novel online learning algorithm to tackle both problems

Supervised: Gianmarco, Marco
Expected time for graduation: 8-12 months
Start: now
Contact: **gianmarco.genalti@polimi.it**

Online Selection of Learning Algorithms (OL)

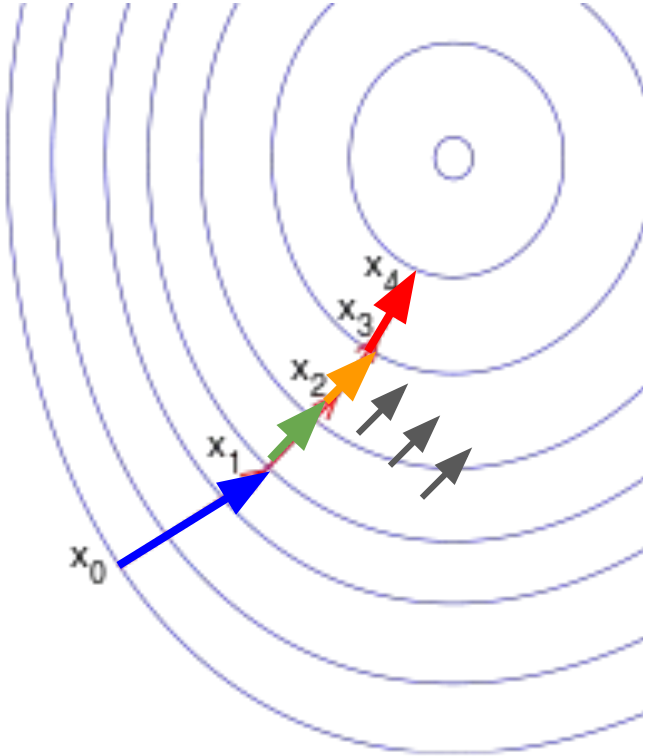

Setting: several learning algorithms optimizing the **same** ML model are available:

- Which one is the most appropriate one?
- Can we find a sequence of methods to get to the optimum point with the smallest number of iterations?

Goal: **select online** to which algorithm assign current data in order to **speed up** finding the optimum

Supervised: Alberto and Francesco
Expected time for graduation 9-12 months
Start: in a few months
Contact: **albertomaria.metelli@polimi.it, francesco1.trovo@polimi.it**

Non-Stationary MAB (OL)

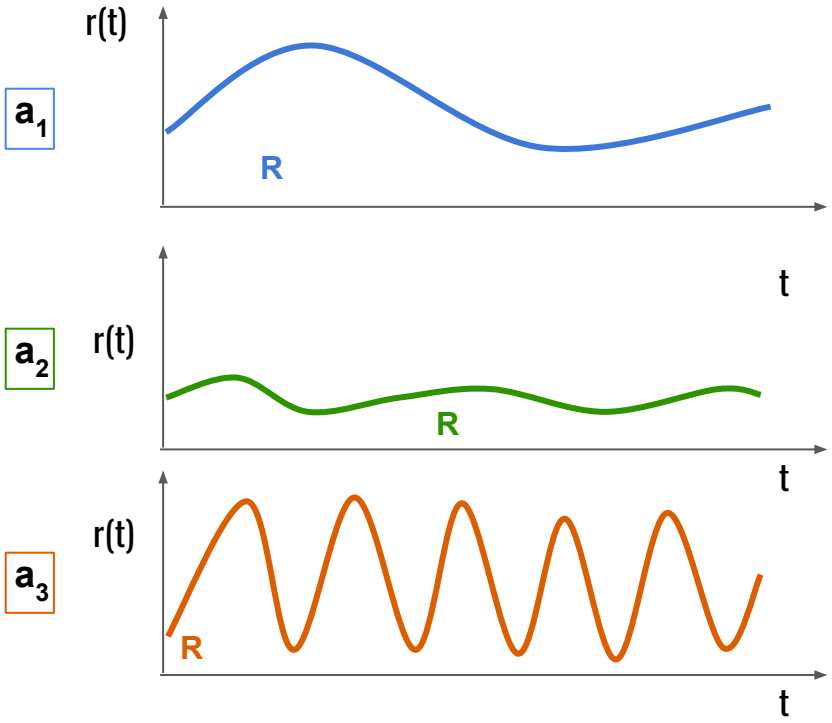

The literature provides already a wide spectrum of algorithms:

- Active (abrupt changes)
- Passive (smooth changes)

The thesis want to use change point detection techniques to improve the existing algorithms and design smart policies to avoid discarding informative data

Supervised: Francesco e Giacomo
Expected time for graduation: 9-12 months
Start: Now or in a few months
Contact: **francesco1.trovo@polimi.it**

Ranking Bandits (OL)

Ranking Bandits → Pull an arm and observe a rank between the empirical means of the rewards

Goal of the thesis → Develop algorithms and methodologies working in this framework

- *Prove the regret bounds of the developed algorithms*
- *Apply the algorithms to real world scenarios (such as Recommender Systems)*

Supervised: Francesco Emanuele, Davide, Francesco
Expected time for graduation: 9-12 months
Start: now
Contact: **francescoemanuele.stradi@polimi.it**
davide.maran@polimi.it
francesco.bacchiocchi@polimi.it

Learning to Collude in Digital Advertising

Ad Auctions → Determine how to allocate ads to advertising slots on a web page

Media agencies collectively manage many (competing) advertisers by placing bids on their behalf

- *How should they collude to maximize revenue?*
- *Which monetary incentives should the media agency give them to have them staying in?*

Goal of the thesis → Design of algorithms working for repeated auction settings by using techniques from the field of online machine learning

- *Maximize revenue of the agency over time*
- *Guarantee that colluders receive fair incentives*

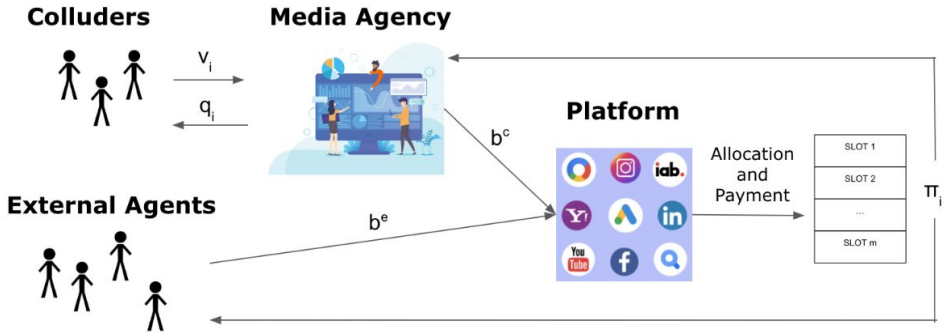

Supervised: Matteo, Giulia and Alberto Marchesi
Expected time for graduation: 7-9 months
Start: now
Contact: **matteo.castiglioni@polimi.it**
giulia.romano@polimi.it
alberto.marchesi@polimi.it

Online Learning for Bayesian Persuasion (OL)

Persuading agents to do desirable actions is a core task in many applications in which there is a central authority possessing information which is valuable for them

Example: ridesharing apps (such as *Uber*) that have to decide which clients are proposed to the drivers, in which the driver receives only partial information on the user (average tip/rating)

Goal: Design of online learning algorithms for the persuasion problem, extending current work to

- Sequential actions and private information
- Multiple agents

Supervised: Matteo, Martino and Alberto Marchesi
Expected time for graduation: 7-9 months
Start: now
Contact: **matteo.castiglioni@polimi.it**
martino.bernasconideluca@polimi.it
alberto.marchesi@polimi.it

Online Learning for Contract Design (OL)

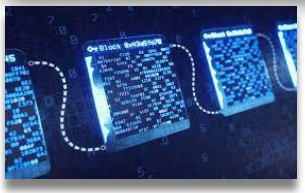

Algorithmic game theory: Studies the computational properties of micro-economic models

Contract theory studies how to design payment schemes to incentivize an agent to take a costly, desirable action

- Salesperson-company interaction
- Pervasive in digital economies: *crowdsourcing platforms, blockchains, healthcare systems*

Goal: Design of algorithms for computing adaptive contracts in dynamic settings by using online learning techniques

Supervised: Matteo, Francesco and Alberto Marchesi
Expected time for graduation: 7-9 months
Start: now
Contact: **matteo.castiglioni@polimi.it**
francesco.bacchiocchi@polimi.it
alberto.marchesi@polimi.it

Communication Equilibrium: No-regret Learning and Costly Communications (GT)

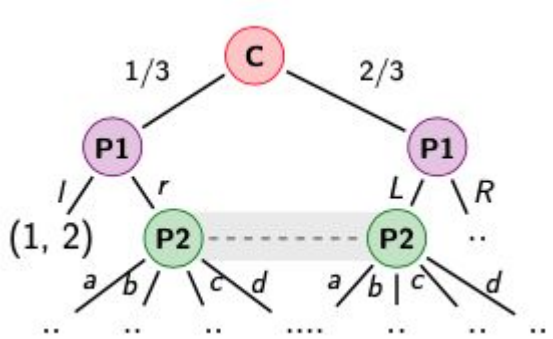

Classical notions of correlation in *multi-player games* (e.g., *correlated equilibria*) do not allow any kind of **strategic communication** between the agents.

However, if communication is introduced, the corresponding solution concept (i.e., **communication equilibrium**) brings advantages both in terms of computational complexity and utility.

Goal: Model different scenarios of strategic communication (e.g., with costly communications) and develop no-regret learning dynamics that converge to the equilibrium

Supervised: Federico
Expected time for graduation: 7-9 months
Start: now
Contact: **federico.cacciamani@polimi.it**

Multi-Fidelity Best-Arm Identification (OL)

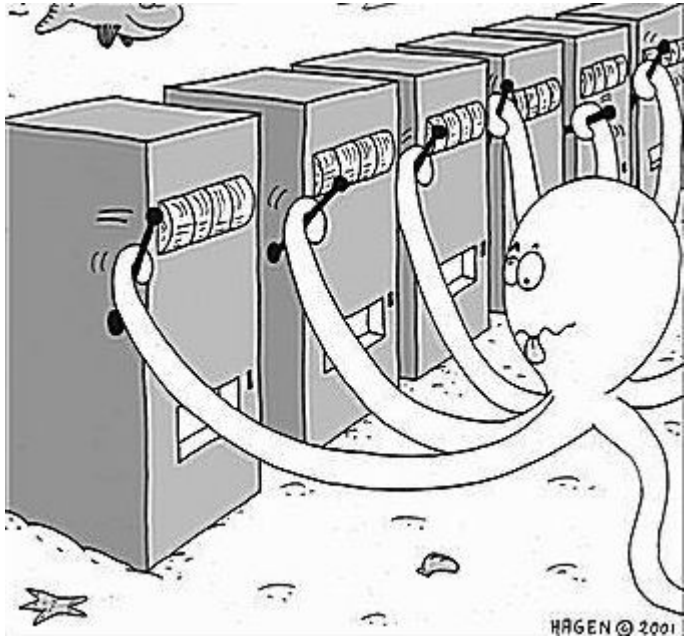

- **MF-BAI**: multiple approximations of different precision and cost of the same arm are available
- Models a huge number of ***real problems***
  - Neural Architecture Search
  - Physic simulation studies
  - Simulator selection in F1 cars
- **Goal**:
  - Pure exploration
  - Identify the best arm while minimizing the total cost
  - **Theoretical study**

Supervised: Riccardo P.
Expected time for graduation 9-12 months
Start: now
Contact: **riccardo.poiani@polimi.it**

Lipschitz Best-Arm Identification (OL)

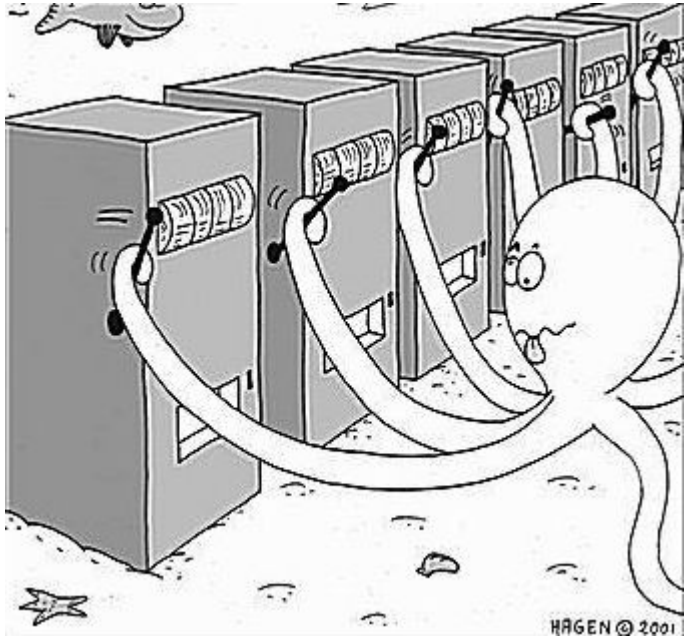

- **Lipschitz BAI**: arms values are coupled in some given metric space
- Models a huge number of ***real problems***
  - Select best F1 car setting
  - Medical trials
  - Recommender Systems
- **Goal**:
  - Pure exploration
  - Identify the best arm while minimizing the total number of samples
  - **Theoretical study**

Supervised: Riccardo P.
Expected time for graduation 9-12 months
Start: now
Contact: **riccardo.poiani@polimi.it**

Adaptive Communication in Distributed RL

Problem:
RL in distributed settings where cooperation is needed despite not knowing exactly what is happening. The **communication** between the agents is essential, even though the full potentials of it are not known.

Issues: many different communication protocol available, choosing the **right protocol** while obtaining the **right effect** is hard

Goal: Find new ways to **share useful information** to guarantee a more scalable and **effective learning process**, possibly in a **adaptive fashion** during learning. The study will focus on **theoretical aspects**.

Supervised: Riccardo Z.
Timespan: 9-12 months
Start: now
Contact: **riccardo.zamboni@polimi.it**

Integrative science, Intelligent data platform for Individualized LUNG cancer care with Immunotherapy

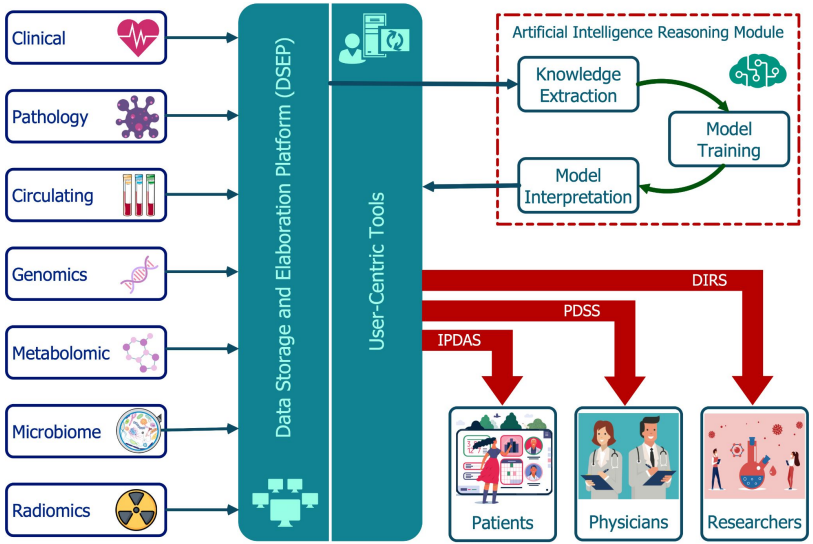

- **Goal**: Implement a platform to integrate the expertise and data provided by different institutions (e.g., Istituto Nazionale dei Tumori, Istituto Europeo di Oncologia, Lung Cancer Europe) to provide the access to a ML-based predictive tool for immunotherapy treatments

- **Requirements**:
  - Good Programming and Software develent skil

Supervised: Francesco
Expected time for graduation: 9-12 months
Start: in a few months
Contact: **francesco1.trovo@polimi.it**

Handling Non-stationarity in Linear Bandits (OL)

Problem: What if the θ^* parameter changes w.r.t. time?

Goal: Find an efficient way to handle a change in the parameter θ^* in order to suffer a lower regret

Linear Bandits are a generalization of the standard bandit problem

- Many applications in real-life problems
- **Reward** in Linear Bandits is

$$X_t = A_t^\intercal \theta^* + \eta_t$$

Supervised: Alessio
Expected time for graduation 9-12 months
Start: now
Contact: **alessio.russo@polimi.it**

Artificial Intelligence Research and Innovation Center

L'Artificial Intelligence Research and Innovation Center del Politecnico di Milano porta AI e innovazione dai laboratori di ricerca del Politecnico di Milano alla tua azienda.

POLITECNICO
MILANO 1863

What is a thesis in industrial research?

Target: developing new solutions and deploy in real-world settings

Working with the research and development team of companies

Examples:

- Novel algorithms and experimental analysis in real-world settings

SIEMENS

SIEMENS

Distributed RL for Industry 4.0

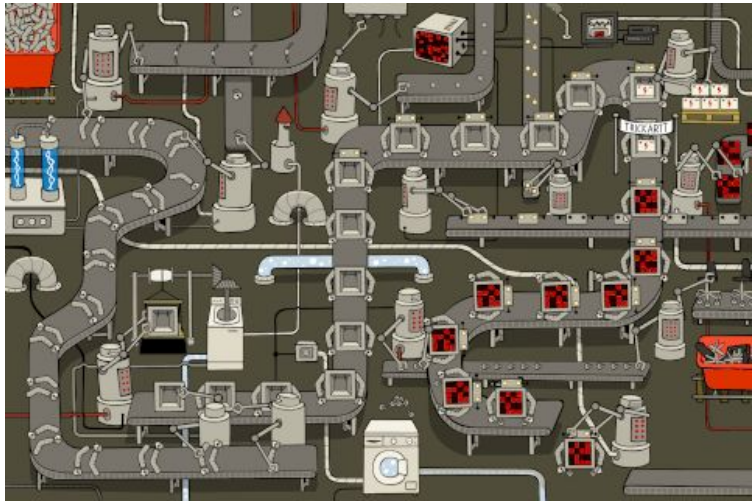

Supervised by: Riccardo Z.
Timespan: 9 Months
Start: now
Contact: **riccardo.zamboni@polimi.it**

Setting

Provided a digital twin of a factory plant with several robots, we want to **develop effective distributed RL algorithms** for solving a problem of the **cooperation** with **partial information**, **decentralized decision** and **delayed feedbacks from the environment.**

Goal

Develop and adapt state-of-the-art algorithms to the specific setting of Industry 4.0

Requirements: Programming skills in Python (and possibly some C++)

LEONARDO

Hierarchical RL for autonomous aircraft

Supervised by: Gianluca
Timespan: 9-12 months
Start: Now
Contact: **gianluca.drappo@polimi.it**

Setting

Autonomous management of a heterogeneous team of aircraft during mission execution in a contested operating environment, using **Hierarchical Reinforcement Learning** approaches.

Objectives:

- Application and adaptation of state of the art algorithms in a custom simulation environment
- Update simulator functionalities
- Eventual extension to the **Multi-Agent** framework

LEONARDO

Goal-based RL for autonomous aircraft

Setting

Mission evaluation through goal-based Reinforcement Learning for Human-autonomous team interaction.

Objectives:

- Application of state of the art algorithms in custom simulation environment.

Supervised by: Luca
Timespan: 9-12 months
Start: Now
Contact: **luca.carminati@polimi.it**

Baker Hughes

Reinforcement Learning for controlling a Compressor

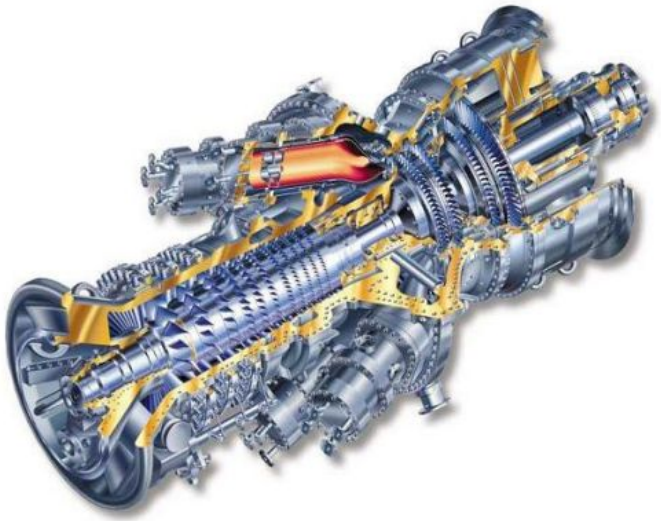

- The compressor is controlled via a **PID** (Proportional, Integral, Derivative)
- **Current Status**: manual tuning of the PID gains
- **Goal**: use RL to automatically tune the PID gains to avoid ***surge*** and regulate ***speed***
- **Available Data**:
  - Historical time series of real measurements
  - Simulator of the turbine

Supervised by: Alberto Metelli
Timespan: 9-12 months
Start: in a few months
Contact: **albertomaria.metelli@polimi.it**

IMICIB - Intesa Sanpaolo

IMI CORPORATE & INVESTMENT BANKING
INTESA SANPAOLO

Reinforcement Learning for Automatic Trading

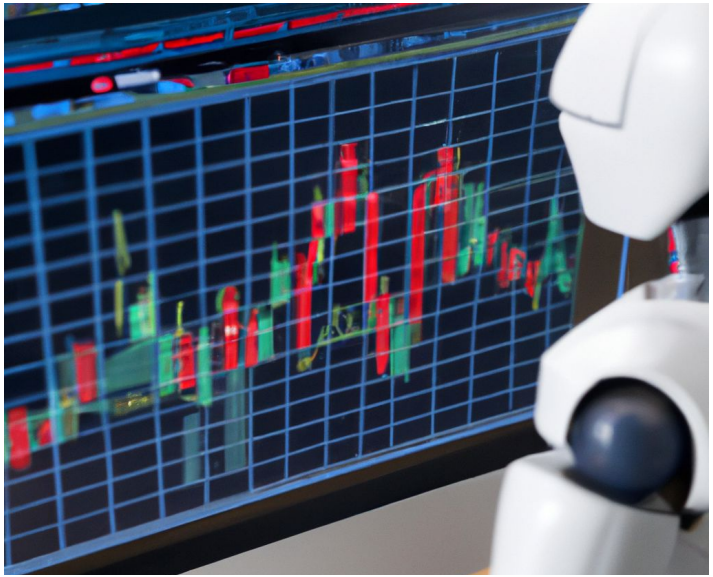

- **Setting**:
  - Realistic trading framework;
  - Historical dataset (€/$, crypto...);
- **Current Challenges**:
  - Increase performance;
  - Deal with non-stationarity;
  - Reduce training time.
- **Objectives**:
  - Many different directions are possible:
    - *Algorithmic*: Online/Deep approaches;
    - *Technologic*: Accelerate training with GPU based implementations;
    - *Financial*: Statistical analysis of financial time series for feature extraction;

Supervised by: Antonio Riva, Luca Sabbioni, Pierre Liotet.
Timespan: 9-12 months
Start: Now
Contact: **antonio5.riva@polimi.it**

IMICIB - Intesa Sanpaolo

ABIDES Arena

- **Setting**:
  - Realistic market simulator (ABIDES);
  - Realistic trading agents;
- **Objectives**:
  - Design a (realistic) software framework for algorithms to trade against one another;
  - Develop and test such software;
  - Realise an online challenge for ICAIF 2023!
- **Requirements**:
  - Recommended:
    - Software engineering skills;
    - Programming skills;
  - Plus: Knowledge of virtualization technologies and Docker.

Supervised by: Pierre Liotet, Luca Sabbioni, Antonio Riva, Lorenzo Bisi
Timespan: 9-12 months
Start: Now
Contact: **pierre.liotet@polimi.it, luca.sabbioni@polimi.it**

RSE - Ricerca Sistema Energetico

Photovoltaic Fault Diagnosis

Photovoltaic systems are affected by multiple kind of faults (physical or electrical, localised or systemic)

Design ML techniques to detect, isolate, and classify faults on photovoltaic systems

Data are collected using a real-world fault facility

Supervised by: Alessandro Lavelli
Timespan: 9 months
Start: immediately
Contact: **alessandro.lavelli@mlcube.com**


ML cube

What is a thesis in a company?

Target: developing solutions/products for specific clients' needs

These thesis involve a **paid internship** in ML cube and require further selection process (**interview**)

Examples:

- Development of a software product
- Design and implementation of software solutions

ML cube Platform

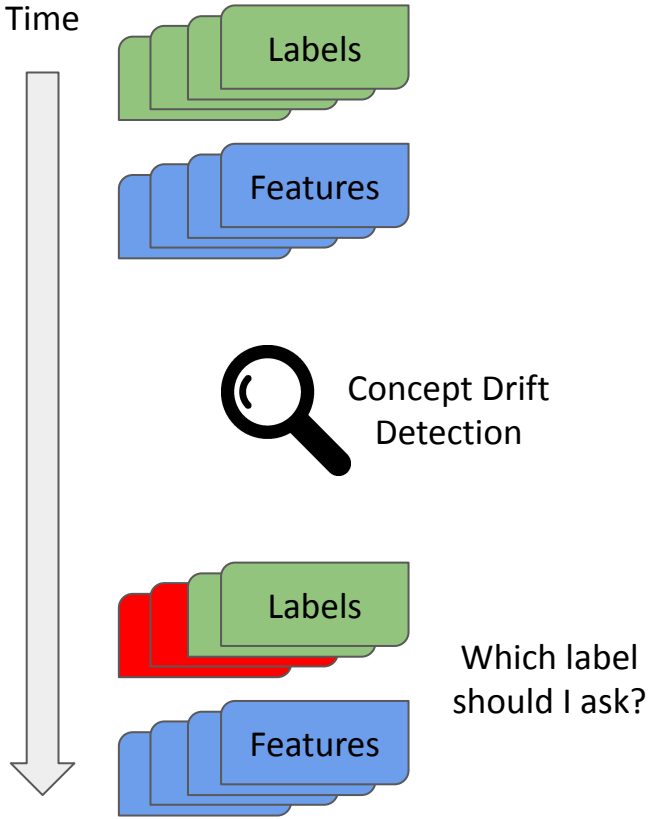

- **ML cube Platform:** *detects* concept drift and suggests how to *weight* past samples.
- Asking *new labels* can be more effective, but it has a *cost*.

Real estate market example: "*Which houses should be re-estimated after a market drift?*"

- **Goal:** Develop an *active learning* approach to select which samples to relabel after a drift, in order to optimize a *cost/performance* trade-off.

Supervised: Alessio Russo, Lorenzo Bisi
Expected time for graduation 9-12 months
Start: now
Contact: **alessio.russo@polimi.it**
lorenzo.bisi@mlcube.com

Customer Behavior Prediction

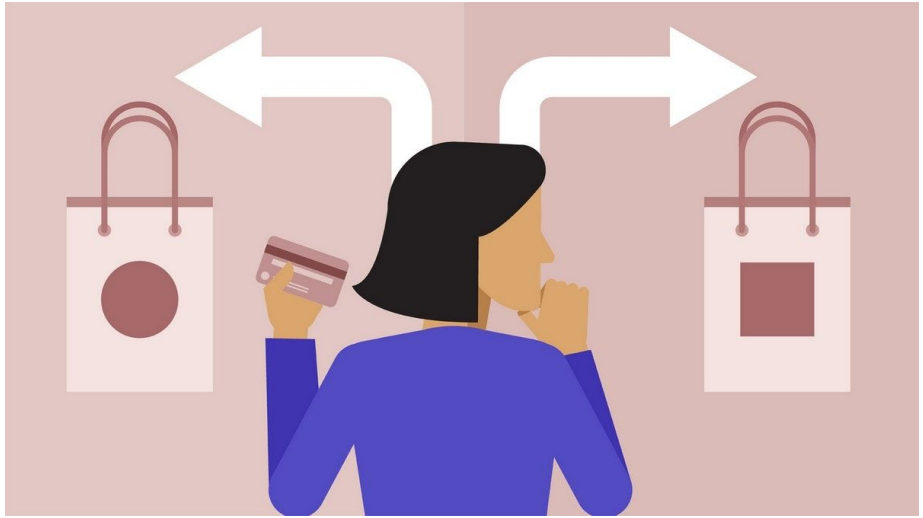

- Mass **marketing campaigns** are expensive.
- *Which customers* are more *willing to buy* a certain product?
- Companies are often allowed to gather anonymized data,

e.g., *history of previous orders, personal profile.*

- **Goal:** Realize an *explainable* ML model to predict customer behavior, in order to provide targets for marketing campaigns.

Supervised: Lorenzo Bisi
Expected time for graduation 9-12 months
Start: now
Contact: **lorenzo.bisi@mlcube.com**

Banca di Ragusa
Automatic Data extraction from Documents

- **Problem**: Bank operators have to process hundreds of documents to extract customers information and start procedures. This task can be automatized by AI systems!
- **Goal**: design and develop a system able to automatically extract specific information from documents (structured and not-structured)
- Required skills:
  - Good programming skills
- **Skills you will acquire:**
  - OCR and NLP technologies
  - Design MLOps architectures on AWS

Supervised: Lorenzo Bisi, Alessandro Nuara
Expected time for graduation 9-12 months
Start: now
Contact: **lorenzo.bisi@mlcube.com**

Gazzetta - Automatic Articles Generation

- **Problem**: Writing an article is an expensive task that requires relevant effort by journalists. However, in some cases (e.g., biographies, sport events) this task can be supported (or automatized) by AI algorithms.
- **Goal:** Realize an AI based platform able to automatically generate Gazzetta dello Sport articles.

Supervised: Lorenzo Bisi, Alessandro Nuara
Expected time for graduation 9-12 months
Start: now
Contact: **alessandro.nuara@mlcube.com**

Other information

- The usual duration of a thesis ranges from 8 to 14 months (according to the number of credits left and your effort)
- If you are interested in one of these topics, you have to fill this form:
  - https://forms.gle/kYmNm2WDLVB2Bt6L8
- You can find this address on these webpages:
  - https://rl.airlab.deib.polimi.it/
  - https://gatti.faculty.polimi.it/
  - https://trovo.faculty.polimi.it/
- Deadline: September 19