From Continual Learning to Causal Discovery in Robotics

Luca Castri[1], Sariah Mghames[1], Nicola Bellotto[2]
[1]University of Lincoln, UK
[2]University of Padua, Italy

Introduction

Reconstructing accurate causal models of dynamic systems from time-series of sensor data is a key problem in many real-world scenarios, especially robotics.

The possibility to leverage the **Continual Learning** (CL) paradigm in the **Causal Discovery** process may represent a way to make the latter feasible for robotics applications where the computational resources are limited

The robot may be added as an active agent in the learning process in order to help Increasing the quality of the reconstructed causal models.

Motivation

- Causal discovery approaches require static or time-series data collected and processed in advance;
- In many real-world robotics applications, due to the limited computing and memory hardware resources, this approach could appear inefficient or even unfeasible;
- The link between **Continual Learning** (CL) and **Causality** might represent a stepping stone towards the exploitation of causal discovery algorithms that currently suffer many limitations in autonomous robots.

Research Objectives

The proposed approach, **Causal Robot Discovery (CRD)**, would allow the robot to overcome its hardware limitations and to improve the quality of the causal models by continually feeding new data for causal analysis, discarding the old collected one.

CRD overcomes current limitations in causal analysis for real-world robotics applications, addressing in particular:

- the computing and memory hardware resources of the robot, which may hinder its capability to perform meaningful causal analysis;
- the update of previous causal models with new observational and interventional data from the robot to generate more accurate ones.

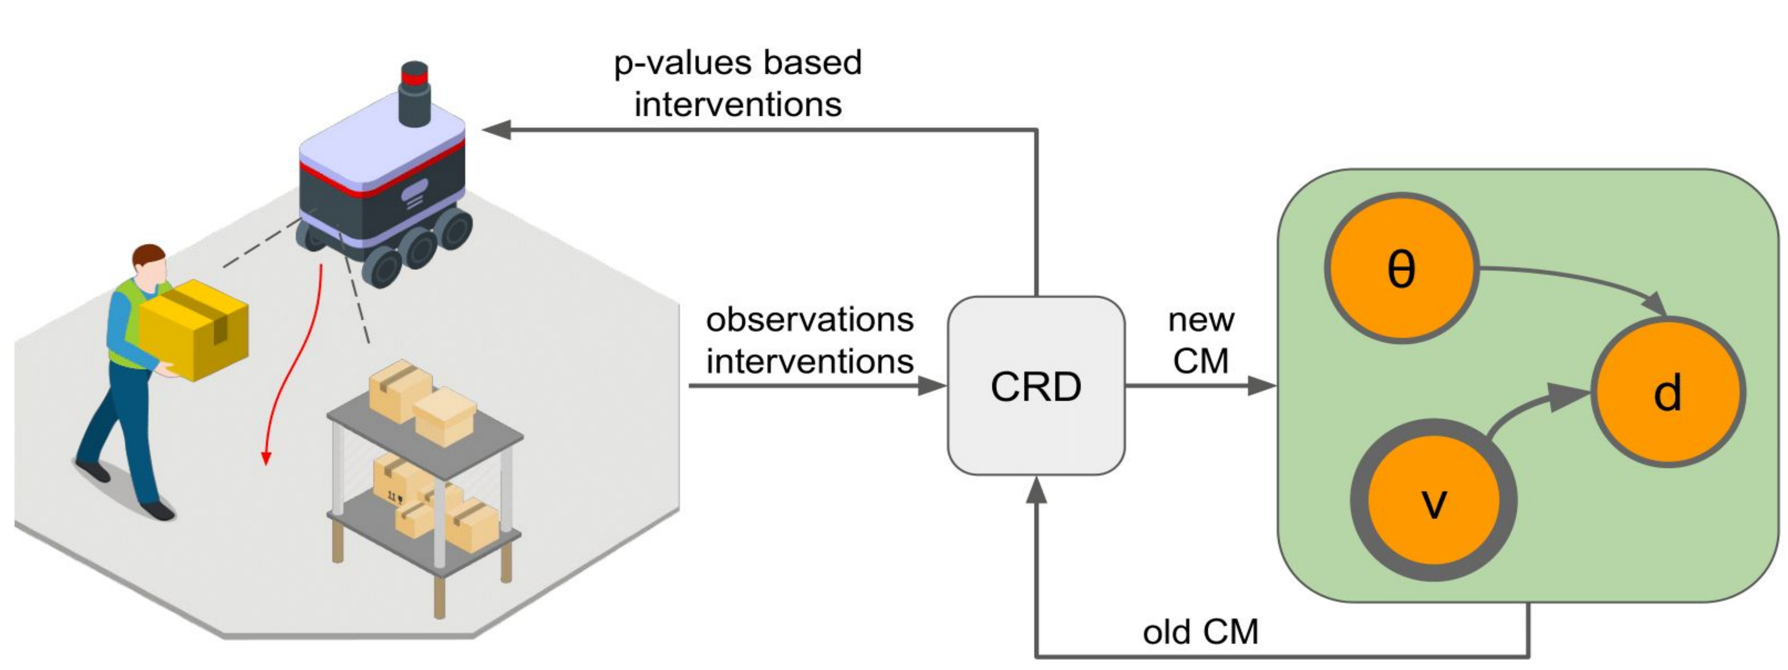

Causal Robot Discovery

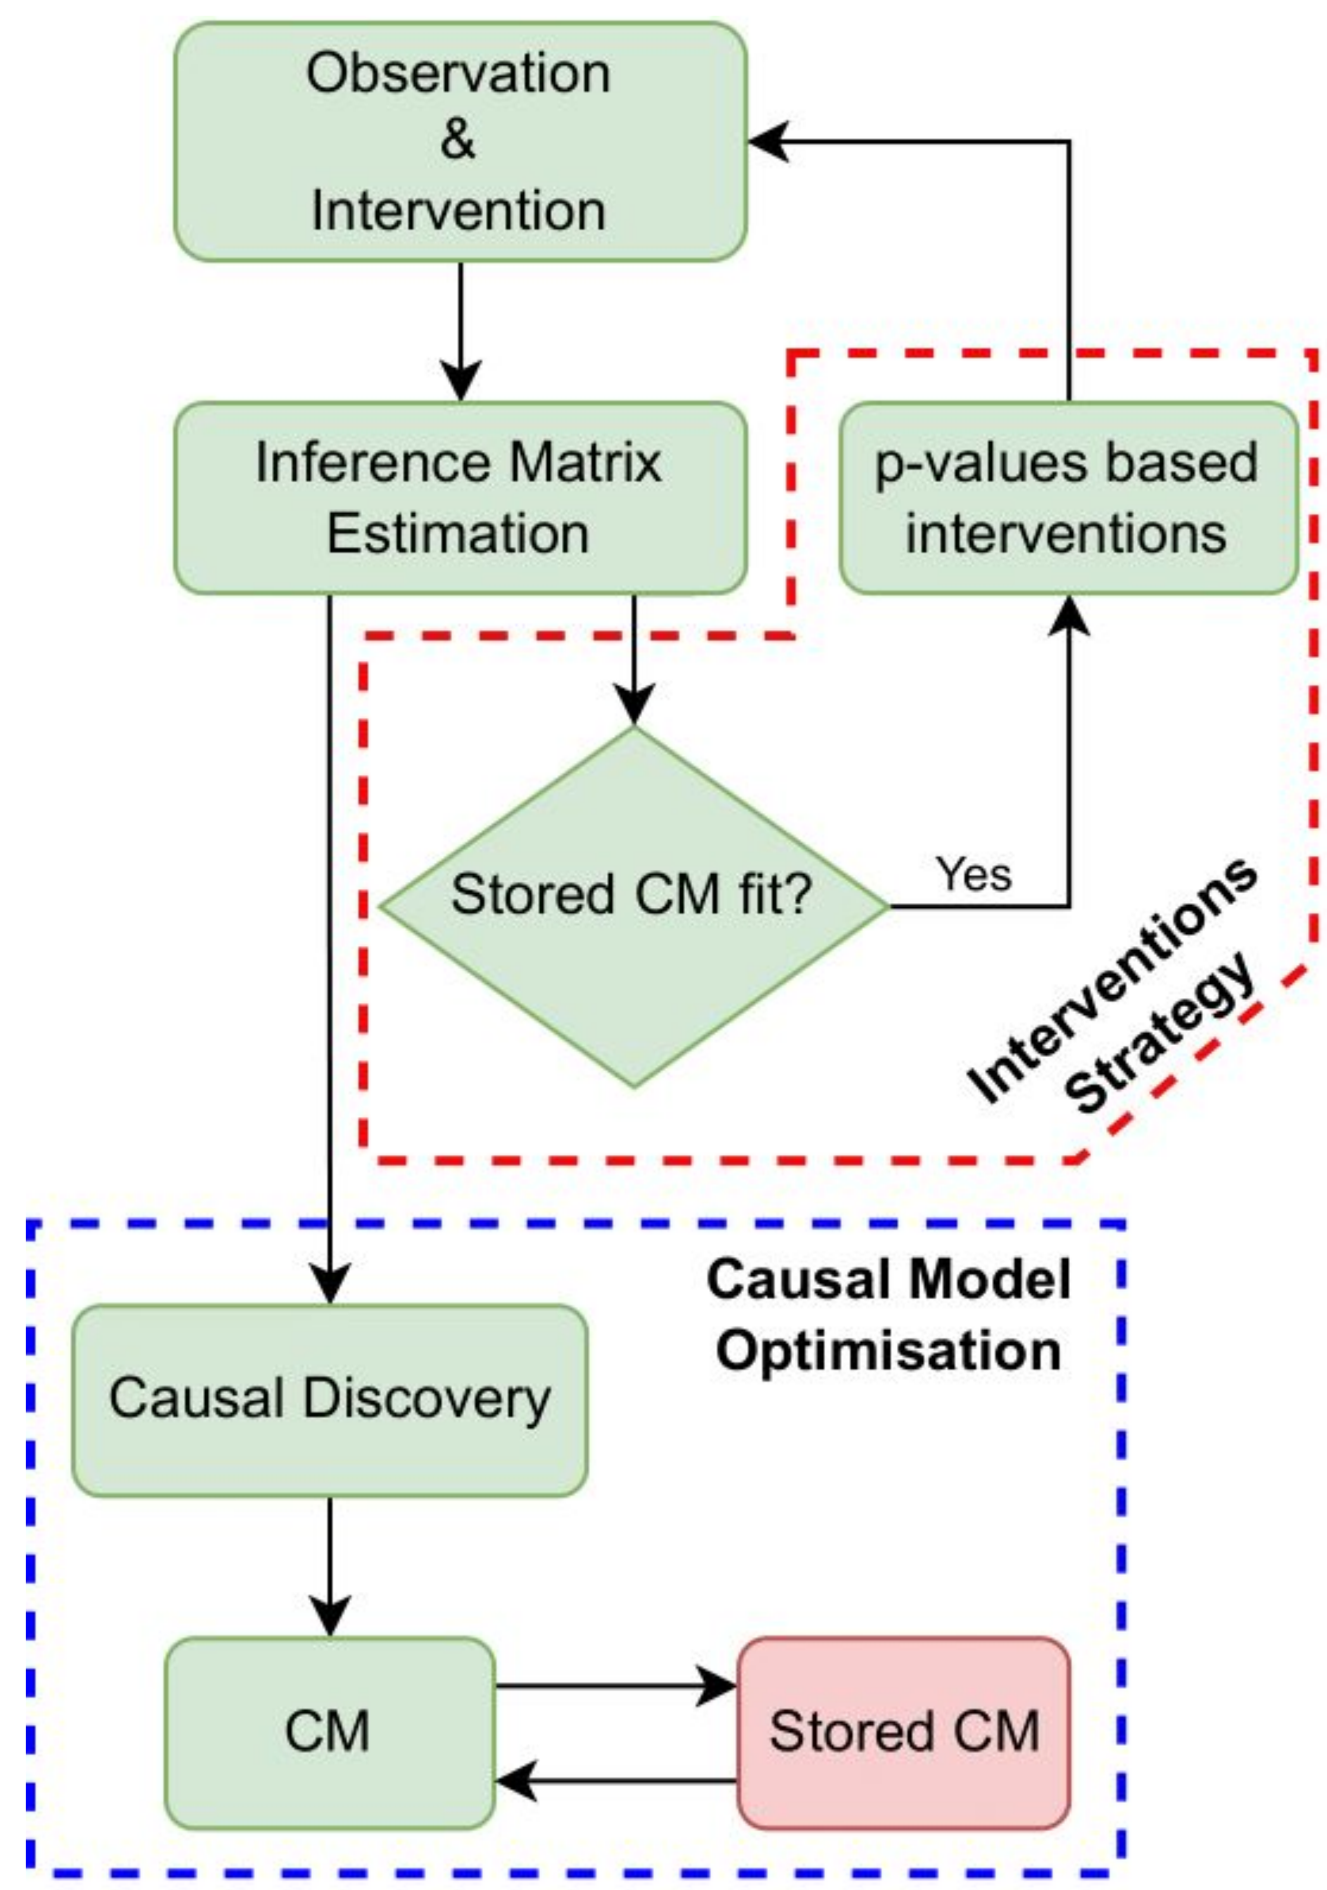

The steps (S) performed by our CRD approach are:

S1. An inference matrix is estimated on a first set of observational data collected by the robot;

S2. At the first attempt, no causal model is stored
⇒ **Interventions Strategy** is neglected;
⇒ **Causal Model Optimisation** estimates and saves the causal model (CM) by [2];

S3. The robot can improve the CM's quality by providing new data to the CRD ⇒ new data and S1 repeated. At this stage, two parallel processes are executed:
→ **Interventions Strategy** compares the stored CM with the previously estimated inference matrix. If they match ⇒ **p-values based interventions**; else ⇒ **new CM exploiting the stored one** [1];
→ **Causal Model Optimisation** performs a causal discovery on the new data and compare the obtained CM with the stored one
⇒ new CM that inherits only the links minimising the p-values of the causal graph.

References

1. Kocacoban, D.; and Cussens, J. 2019. Online Causal Structure Learning in the Presence of Latent Variables. In 2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA), 392–395.
2. Runge, J. 2018. Causal network reconstruction from time series: From theoretical assumptions to practical estimation. Chaos: An Interdisciplinary Journal of Nonlinear Science, 28: 075310.

Contact: **Luca Castri lcastri@*lincoln.ac.uk***

Lincoln Centre for Autonomous Systems (L-CAS)
School of Computer Science, University of Lincoln, UK

Website: ***https://darko-project.eu***

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101017274