Causal Inference for Intelligent Mobile Robots in Dynamic Interaction Settings

Luca Castri
University of Lincoln,
Lincoln, 15 September 2025

UNIVERSITY OF LINCOLN
UNITED KINGDOM

Ph.D. Supervisors
Prof. Nicola Bellotto
Prof. Marc Hanheide

Outline

- Causality Overview
- Motivation
- Causality & Robotics: Challenges
- Contributions
- Summary

Causality Overview

What is it?

“Science that studies the cause-and-effect relationship between events”
[Pearl, J., & Mackenzie, D. (2019). The book of why]

How can robots benefit from causality?

Causal Structure Learning

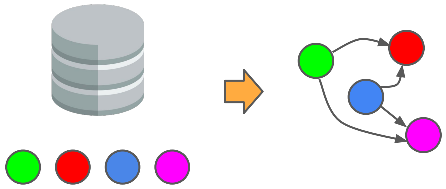

Causal Reasoning

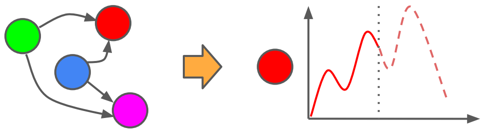

Causal Representation Learning

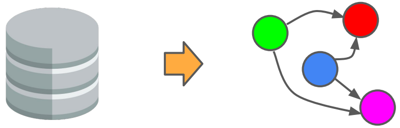

Motivation

Causality applications so far:

- Climate [Runge et al. 2014, 2018, 2019, 2020, Kretschmer et al. 2016, 2017, 2018, ...]
- Healthcare [Runge et al. 2015, Saetie et al. 2021, ...]
- Machine learning [Schölkopf et al. 2021, Seitzer et al. 2021, ...]
- Robotics
  - Imitation learning [Kats et al. 2018, Angelov et al. 2019, 2020]
  - Manipulation [Brawer et al. 2021, Lee et al. 2022, 2023, Cannizzaro et al. 2023a]
  - Autonomous Driving [Howard et al. 2023a,b, 2025]
  - Social HRI [Love et al. 2024a,b]
  - Others [Cao et al. 2021, Cannizzaro et al. 2023b]

  - Causality for modelling human spatial behaviour and robot interactions?

    [Mahata et al. 2017, Vasconez et al. 2019, Jahanmahin et al. 2022, Mukherjee et al. 2022, Dahiya et al. 2023] } Causality not employed

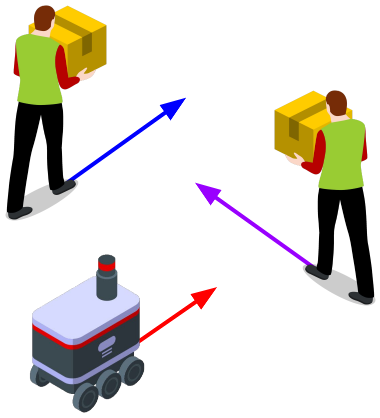

Motivation

Why do we need causal models?

Traditional modelling approaches for human spatial behaviours often ignore the factors that influence them

Having a causal model of human spatial behaviours could enable robots to reason as follows:

- “what happens if I go this way?”

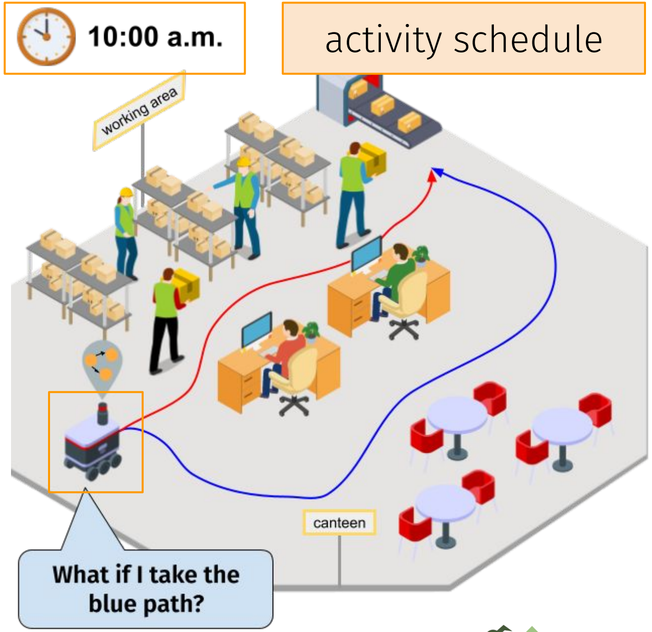

➔ deeper understanding of the scenario
➔ decision-making and forecasting

Causality & Robotics: Challenges

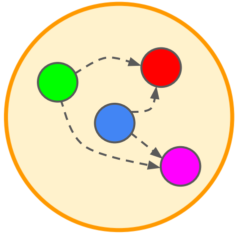

- Time-series causal discovery is too slow and costly for resource-constrained mobile robots

RQ_1 How can state-of-the-art causal discovery methods be improved to enable their use in robotic applications?

- Robots cannot use their embodiment to support causal discovery through interventions
- Causal discovery from time-series uses only observations

RQ_2 Can the robot's actions be exploited to perform interventions and enhance causal discovery?

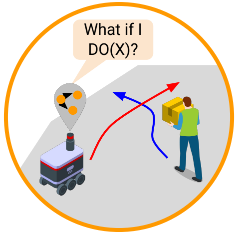

- Causal discovery not integrated into ROS
- No causal reasoning in robot decision-making

RQ_3 Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

Contributions

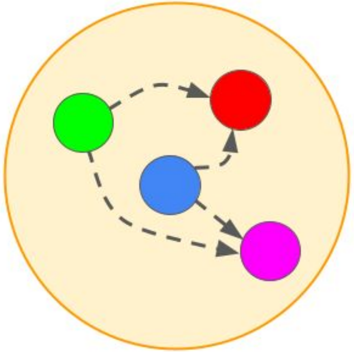

Fast and accurate causal discovery algorithm for human-robot spatial interactions

Observation and intervention-based causal discovery algorithm for time-series

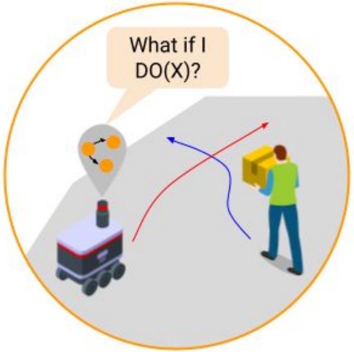

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Fast and accurate causal discovery algorithm for human-robot spatial interactions

How can state-of-the-art causal discovery methods be improved to enable their use in robotic applications?

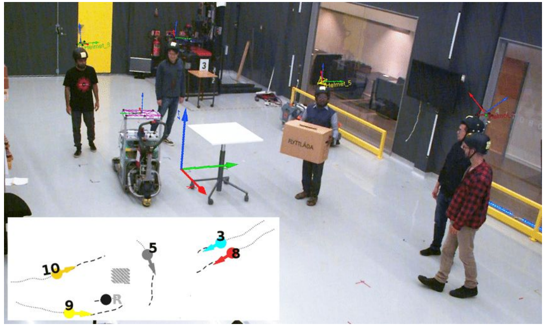

THÖR
[Rudenko et al. 2020]

ATC
[Brscic et al. 2013]

single-agent multi-agent

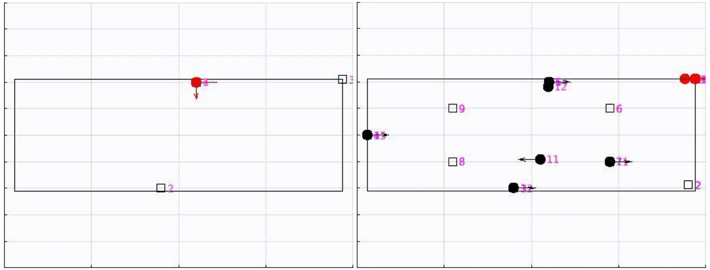

- PCMCI [Runge et al. 2019]

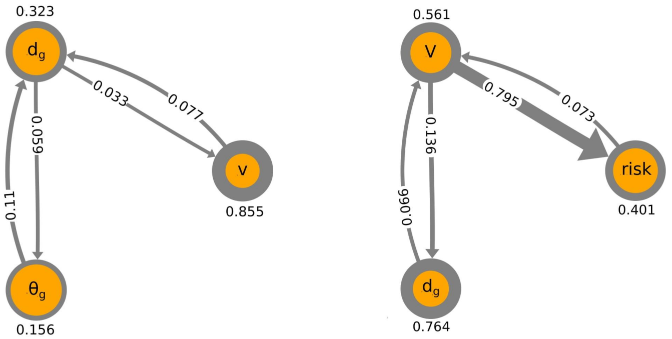

- > 2h
- 6500 samples

Limitation: PCMCI execution time

- Autonomous mobile robots
  - limited hardware resources
  - real-time requirements

→ We need a fast causal discovery method

Fast and accurate causal discovery algorithm for human-robot spatial interactions

How can state-of-the-art causal discovery methods be improved to enable their use in robotic applications?

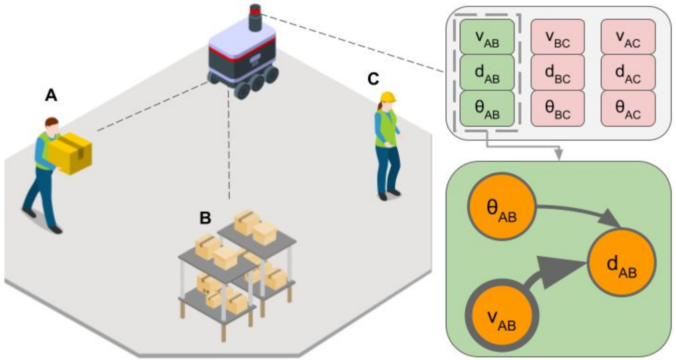

- PCMCI computational complexity: $\mathcal{O}(N^3\tau_{\max}^2 + N^2\tau_{\max})$
- Are all observed variables useful?

GOAL
Build an all-in-one solution to select key variables and reconstruct a causal model

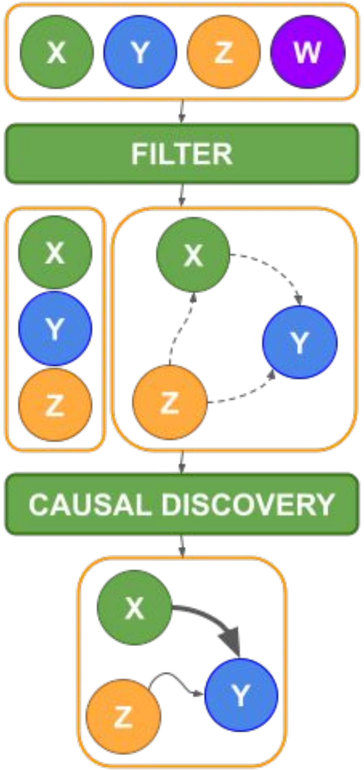

Filtered-PCMCI (F-PCMCI)

1. predefined set of variables
2. remove irrelevant variables using *Transfer Entropy*
3. build hypothetical causal structure from reduced set
4. run PCMCI on hypothetical model

➔ **Faster** and **more accurate** causal discovery

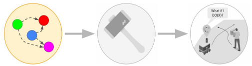

Fast and accurate causal discovery algorithm for human-robot spatial interactions

How can state-of-the-art causal discovery methods be improved to enable their use in robotic applications?

$$S_2 = \begin{cases} x_{0_t} = c_{01} x_{1_{t-2}} c_{02} x_{2_{t-1}} + \eta_{0_t} \\ x_{1_t} = \eta_{1_t} \\ x_{2_t} = c_{21} x_{1_{t-2}}^2 + \eta_{2_t} \\ x_{3_t} = c_{33} + x_{3_{t-1}} + \eta_{3_t} \\ x_{4_t} = c_{42} x_{2_{t-2}} - c_{43} * x_{3_{t-1}} + \eta_4 \\ x_{5_t} = \frac{c_{50} x_{0_{t-1}}}{1 + c_{55} x_{5_{t-2}}} + \eta_{5_t} \\ x_{6_t} = \eta_{6_t} \end{cases}$$

- 10 run tests
- 1500 samples

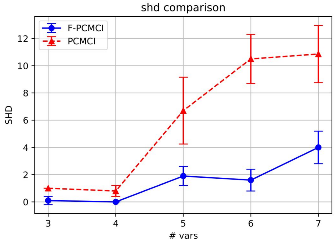

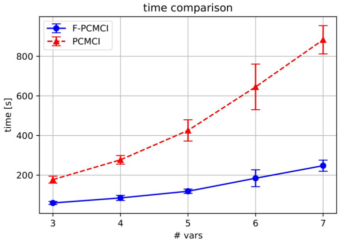

→ **F-PCMCI always faster** and **more accurate** than PCMCI

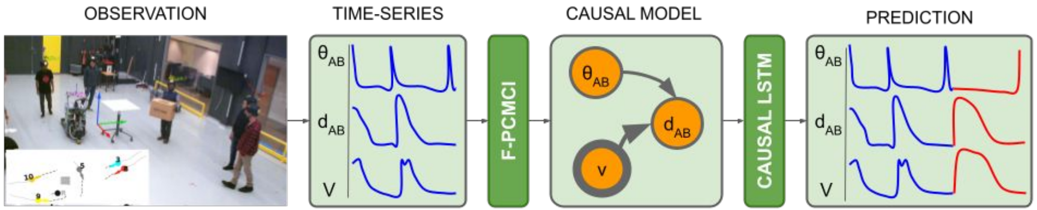

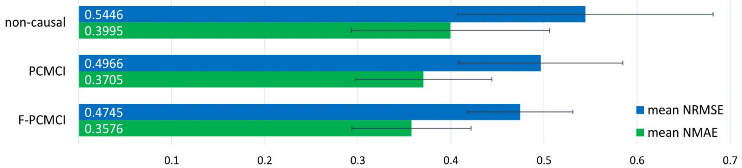

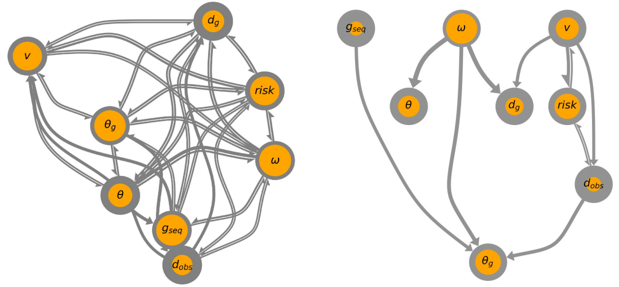

PCMCI ~80mins F-PCMCI ~18mins

- No ground-truth causal model
- Prediction accuracy as a proxy for the accuracy of causal models

Research outcomes

★ Castri et al. "Enhancing causal discovery from robot sensor data in dynamic scenarios," in Conference on Causal Learning and Reasoning, 2023.
★ Castri et al. "Causal discovery of dynamic models for predicting human spatial interactions," in International Conference on Social Robotics, 2022.

- Castri et al., "Efficient causal discovery for robotics applications," in Italian Conference on Robotics and Intelligent Machines (I-RIM 3D), 2023.
- Ghidoni et al., "From human perception and action recognition to causal understanding of human-robot interaction in industrial environments," in ItalIA Convegno Nazionale sull'Intelligenza Artificiale, 2022.

Observation and intervention-based causal discovery algorithm for time-series

Can the robot's actions be exploited to perform interventions and enhance causal discovery?

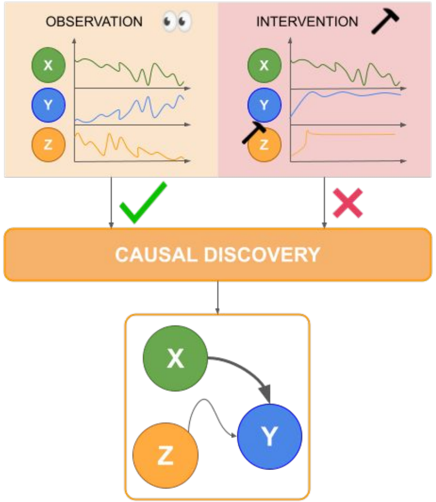

- Complex system ⇒ Observational data insufficient
- Time-series methods do not integrate interventional data

GOAL

First causal discovery method for time-series that uses both observational and interventional data

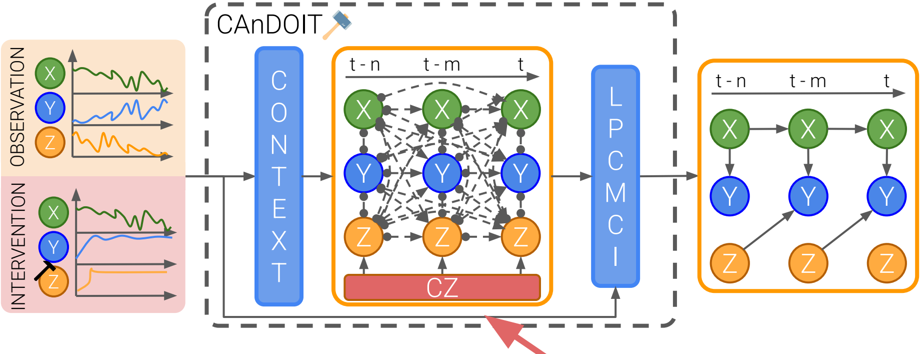

CAnDOIT

CAusal **D**iscovery with **O**bservational and **I**nterventional data from **T**ime-series

- Based on LPCMCI [Gerhardus et al. 2020]

context variables to inject interventional data [Mooij et al. 2020]

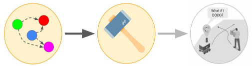

Observation and intervention-based causal discovery algorithm for time-series

Can the robot's actions be exploited to perform interventions and enhance causal discovery?

camera height — camera velocity — camera distance to the cube

$$b = K_h \frac{H}{H_{max}} + K_v \left(1 - \frac{v}{v_{max}}\right) + K_d \frac{d_c}{d_{cmax}}$$

$$\begin{cases} F_c(t) = \cancel{b(H(t-1))} \\ C_c(t) = b(H(t-1), v(t-1), d_c(t-1)) \end{cases}$$

uncertainties

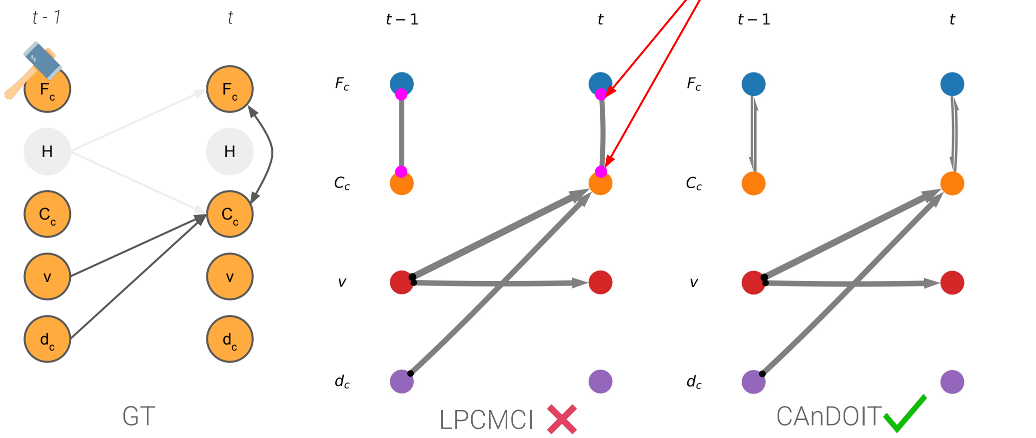

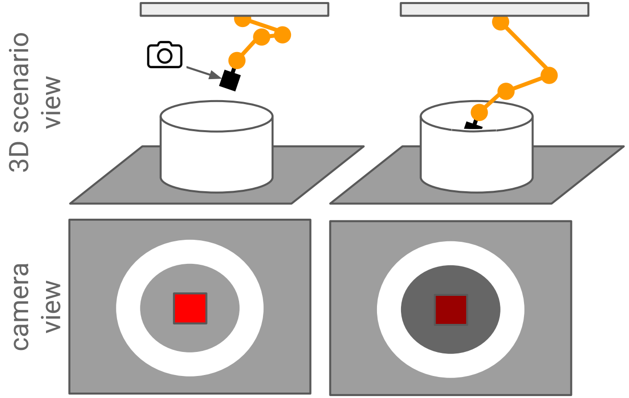

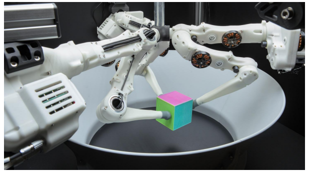

CausalWorld [Ahmed et al. 2021]

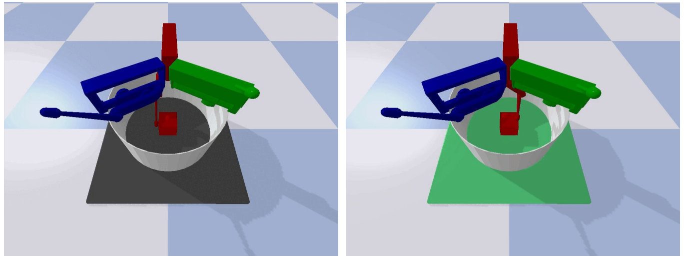

Research outcomes

★ Castri et al. "CAnDOIT: Causal Discovery with Observational and Interventional Data from Time-Series", Advanced Intelligent Systems, 2024.

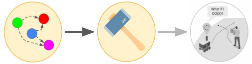

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

- Causal discovery methods are not yet deployable in robotic systems
  - cannot run directly on robots
  - requires data collection + offline analysis
  - causal models not usable in real-time

→ Need for an integration with ROS

GOAL

First ROS-based causal analysis framework

PCMCI F-PCMCI
DYNOTEARS TiMINo
VARLiNGAM TCDF tsFCI

ROS

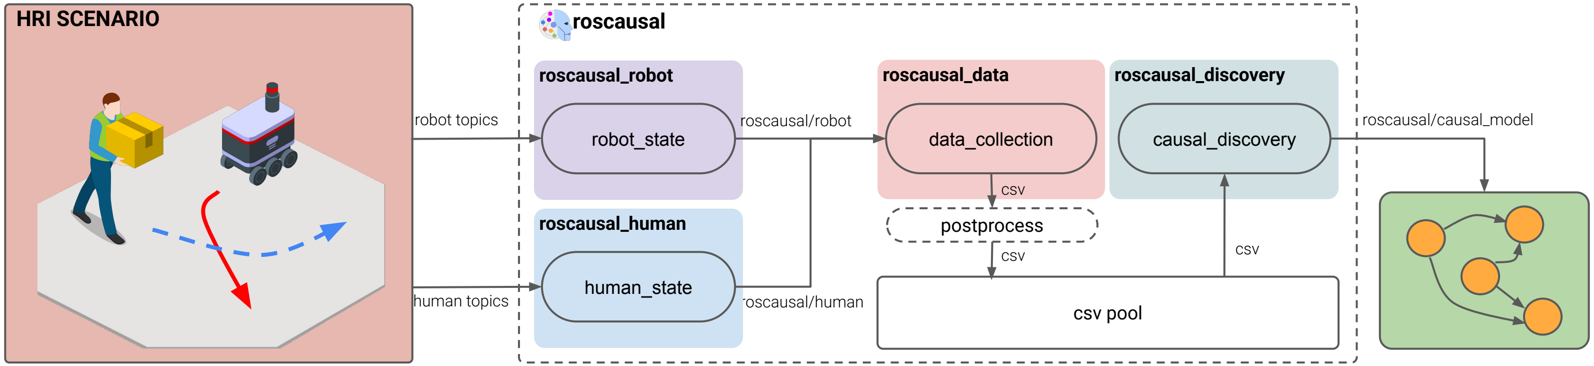

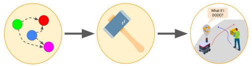

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

- **Simulation** (ROS-Causal_HRISim)
  - TIAGo robot
  - teleoperated and autonomous pedestrians

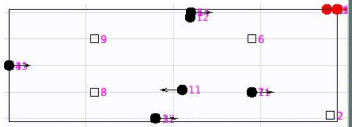

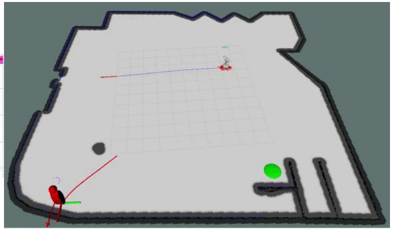

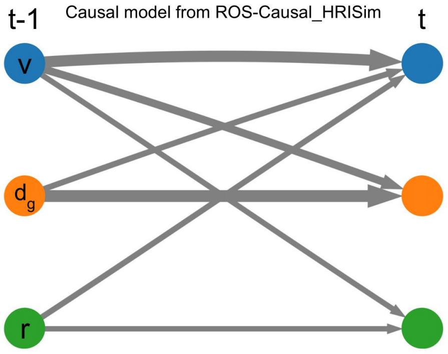

- **Real-world**
  - TIAGo task:
    - rectangular path
  - Participant task:
    - four goal positions
    - avoid the robot

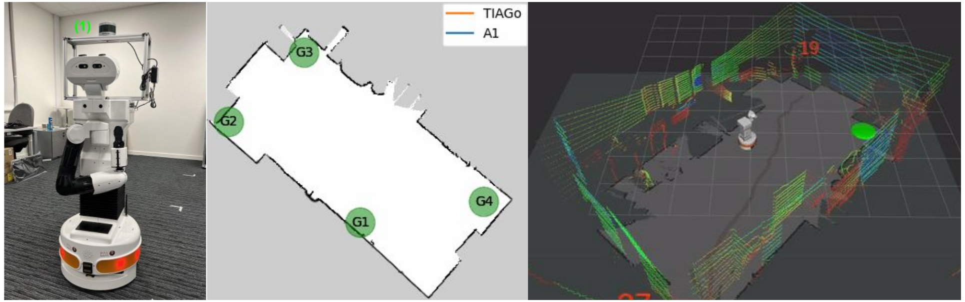

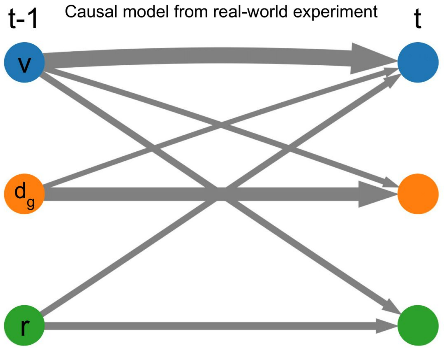

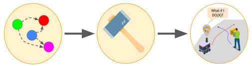

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

- Human-aware navigation
  - relies on predictive models of human motion
  - ignores contextual factors
- Potential safety and efficiency issues

GOAL

Causality-enhanced robot decision-making framework for human-aware navigation

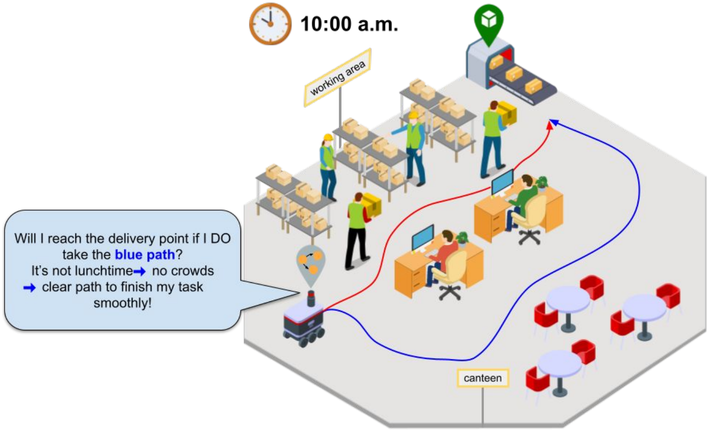

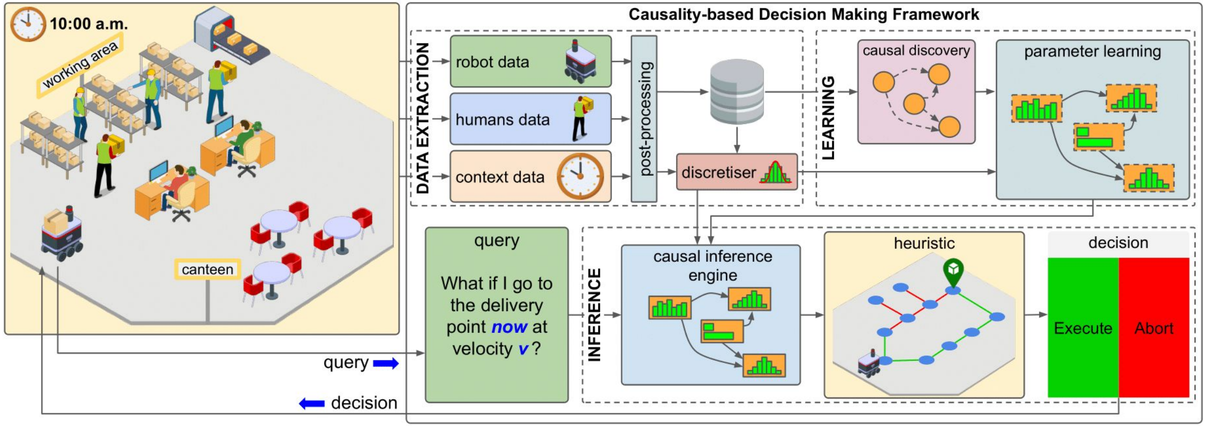

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

Robot Task

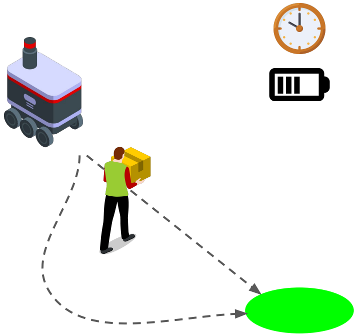

1

"What if I go to now at velocity v?"

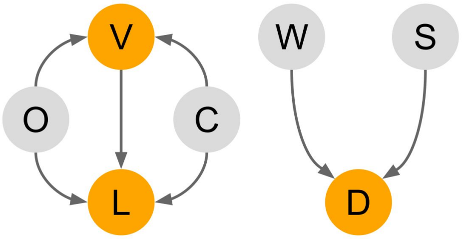

system variables	context factors
V: robot velocity	O: static obstacle
L: battery level change	C: charging status
D: people density	W: waypoint
	S: time slot

2

$$\hat{L} = \mathbb{E}[L \mid do(V = v), C = c]$$

$$\hat{D} = \mathbb{E}[D \mid do(S = s), W = w] \quad \forall w \in \Omega$$

3

$$h(w_i) = \sum_{i=1}^{n-1} \left(\lambda_\delta \cdot \delta(w_i, w_{i+1}) + \lambda_D \cdot \hat{D}(w_i) + \lambda_L \cdot |\hat{L}(w_i, w_{i+1})| \right)$$

4

Safety and efficiency take priority over distance

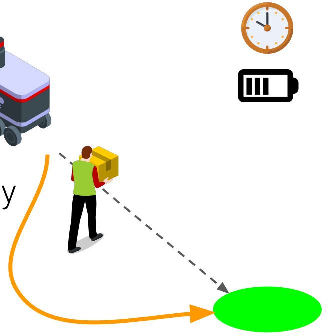

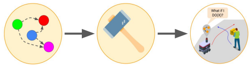

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

PeopleFlow (based on ROSCausal_HRISim)

- context-sensitive humans and robot behaviours in a warehouse setting
- TIAGo robot and autonomous pedestrians

working time

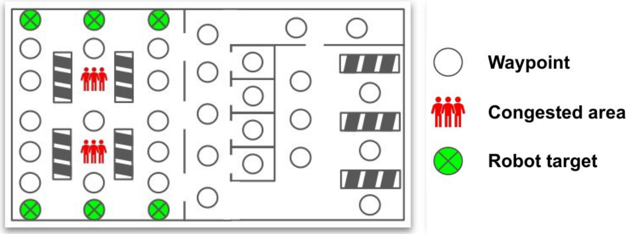

non-causal

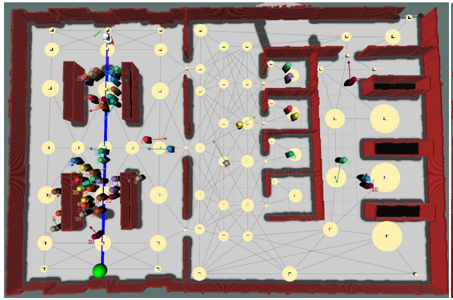

causal

efficiency

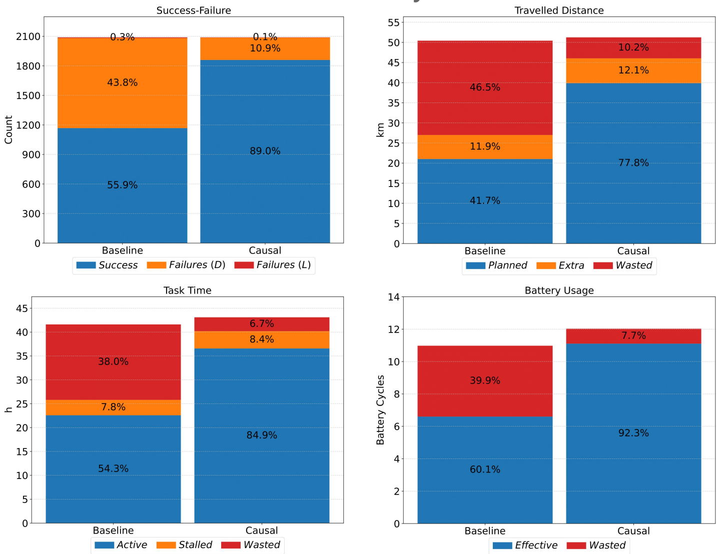

safety

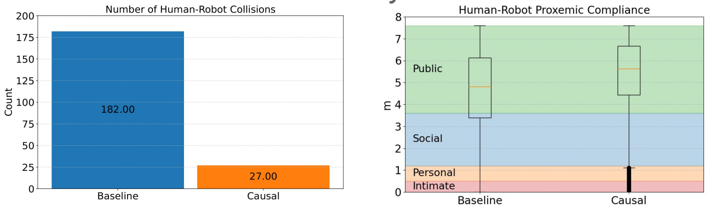

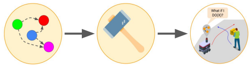

Integrating Causal Inference for Autonomous Robots in Dynamic Environments

Can the robot autonomously reconstruct and use causal models to improve decision-making and interactions in human-shared environments?

Real-world: Poster session

Data collection for parameters learning

- 2 mins rosbag
- 10 Hz

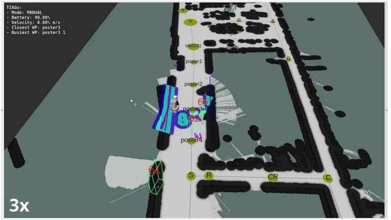

Research outcomes

★ Castri et al. "Causality-enhanced Decision-Making for Autonomous Mobile Robots in Dynamic Environments," under review.

★ Castri et al. "Experimental Evaluation of ROS-Causal in Real-World Human-Robot Spatial Interaction Scenarios," in IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2024.

★ Castri et al. "ROS-Causal: A ROS-based Causal Analysis Framework for Human-Robot Interaction Applications," Causal-HRI @ ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2024.

- Stracca et al., "Darko-nav: Hierarchical risk-and context-aware robot navigation in complex intralogistic environments," in ERF 2025.
- Rudenko et al., "Hierarchical system to predict human motion and intentions for efficient and safe human-robot interaction in industrial environments," in 1st German Robotics Conference, 2025.

Baseline approach

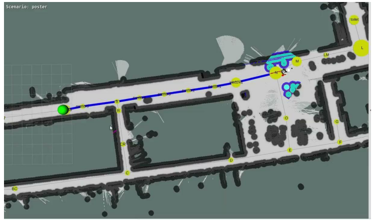

Causal approach

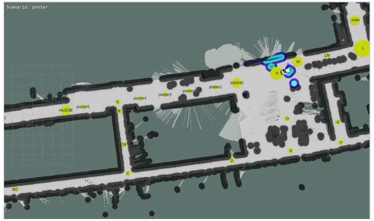

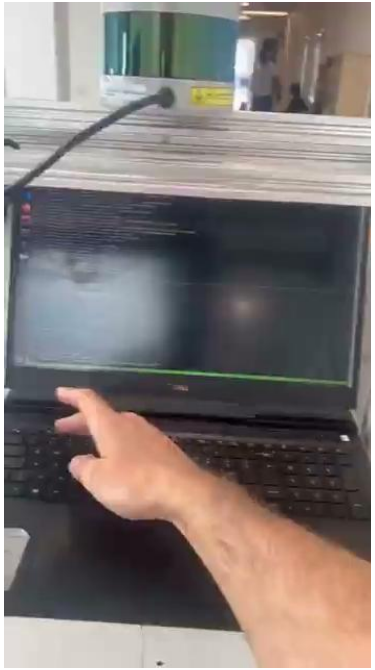

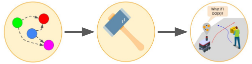

Summary

Main achievements

- **F-PCMCI**: fast and accurate causal discovery method tailored for the computational constraints of robotics
- **CAnDOIT**: observation and intervention-based causal discovery method from time-series
- **ROS-Causal**: ROS-based framework for performing causal discovery directly onboard a robot using its own sensors
- **Causal Decision-Making**: end-to-end framework that learns and uses causal models to make robot navigation safer and more efficient in dynamic and human-shared environments

Future directions

- **Causal representations learning ⇒** learn also the variables not only the causal structure
- **Causal abstraction ⇒** reason at multiple levels, from high-level intentions to low-level physics
- **Continual learning ⇒** deal with non-stationary environments
- **Counterfactual reasoning ⇒** reason about alternative pasts "*What if I had done X instead of Y?*"

List of Publications

Presented in this dissertation

- Rudenko et al., “Hierarchical system to predict human motion and intentions for efficient and safe human-robot interaction in industrial environments,” in 1st German Robotics Conference, 2025.
- Stracca et al., “Darko-nav: Hierarchical risk-and context-aware robot navigation in complex intralogistic environments,” in ERF 2025.
- Castri et al. “Causality-enhanced Decision-Making for Autonomous Mobile Robots in Dynamic Environments,” under review.
- Castri et al. “Experimental Evaluation of ROS-Causal in Real-World Human-Robot Spatial Interaction Scenarios,” in IEEE Int. Conference on Robot and Human Interactive Communication (RO-MAN), 2024.
- Castri et al. “ROS-Causal: A ROS-based Causal Analysis Framework for Human-Robot Interaction Applications,” Causal-HRI @ ACM/IEEE Int. Conference on Human-Robot Interaction (HRI), 2024.
- Castri et al. “CAnDOIT: Causal Discovery with Observational and Interventional Data from Time-Series”, Advanced Intelligent Systems, 2024.
- Castri et al., ‘Efficient causal discovery for robotics applications,’ in Italian Conference on Robotics and Intelligent Machines (I-RIM 3D), 2023.
- Castri et al. “Enhancing causal discovery from robot sensor data in dynamic scenarios,” in Conference on Causal Learning and Reasoning, 2023.
- Castri et al. “Causal discovery of dynamic models for predicting human spatial interactions,” in Int. Conference on Social Robotics, 2022.
- Ghidoni et al., ‘From human perception and action recognition to causal understanding of human-robot interaction in industrial environments,’ in ItalIA Convegno Nazionale sull’Intelligenza Artificiale, 2022.

Developed in the context of the DARKO project, but not part of this dissertation

- Mghames et al., ‘Neurosym: Deployment and evaluation of a ros-based neuro-symbolic model for human motion prediction,’ in IEEE Int. Conference on Cybernetics and Intelligent Systems (CIS) and IEEE Conference on Robotics, Automation and Mechatronics (RAM), 2024.
- Mghames et al., ‘Qualitative prediction of multi-agent spatial interactions,’ in IEEE Int. Conference on Robot and Human Interactive Communication (RO-MAN), 2023.
- Mghames et al., “A neuro-symbolic approach for enhanced human motion prediction,” in Int. Joint Conference on Neural Networks (IJCNN), 2023.
- Castri et al., ‘From continual learning to causal discovery in robotics,’ in Continual Causality Bridge Program @ AAAI, 2023.

Thank you!