

Efficient Causal Discovery for Robotics Applications

Luca Castri¹, Sariah Mghames¹ and Nicola Bellotto^{1,2}

Abstract—Using robots for automating tasks in environments shared with humans, such as warehouses, shopping centres, or hospitals, requires these robots to comprehend the fundamental physical interactions among nearby agents and objects. Specifically, creating models to represent cause-and-effect relationships among these elements can aid in predicting unforeseen human behaviours and anticipate the outcome of particular robot actions. To be suitable for robots, causal analysis must be both fast and accurate, meeting real-time demands and the limited computational resources typical in most robotics applications. In this paper, we present a practical demonstration of our approach for fast and accurate causal analysis, known as Filtered PCMCI (F-PCMCI), along with a real-world robotics application. The provided application illustrates how our F-PCMCI can accurately and promptly reconstruct the causal model of a human-robot interaction scenario, which can then be leveraged to enhance the quality of the interaction.

Index Terms—causal robotics, causal discovery, human-robot spatial interactions

I. INTRODUCTION

The increased use of robots in numerous sectors, such as industrial, agriculture and healthcare, represents a turning point for their progress and growth. However, it requires also new approaches to study and design effective human-robot interactions. A robot, sharing the working area with humans, must accomplish its task taking into account that its actions may lead to unpredicted responses by the individuals around it, while at the same time taking into account the execution time and the computational cost for completing the task. Rapid and accurate understanding of the cause-and-effect relationships in the environment will allow the robot to reason on its own actions. The latter represents a crucial step towards effective human-robot interactions and collaborations.

Causal inference [1] is an active research area in various fields, including robotics [2], human-human and human-robot spatial interactions (HRSI) [3], [4]. Over the past decades, numerous causal discovery methods have been developed for static and time-series data [5], [6]. However, most of these works overlooked a feature that is crucial for real-world applications, i.e. the computational cost of the causal analysis when applied to scenario with limited hardware resources and real-time requirements, such as autonomous robotics. Indeed, causal analysis of complex and dynamical systems is extremely demanding in terms of time and hardware resources [6], making it a challenge for autonomous robotics [7].




Fig. 1. A mobile robot in a warehouse-like settings observes the interaction between agents A and B. With our method, the robot can ignore interactions AC and BC, since agent B is a static object and agent C is a stationary human not participating in the interaction.

This paper presents an overview of our approach Filtered PCMCI (F-PCMCI) along with a real-world robotics application. F-PCMCI offers an all-in-one solution that identifies the causal features characterising the system and builds a causal model directly from time-series data. As a result, the causal discovery process turns out faster and more accurate, rendering it well-suited for applications in autonomous robotics. In an automated warehouse scenario (see Fig. 1), where a robot observes interactions among objects and humans (e.g., workers and shelves), identifying relevant features from the robot’s sensors is vital. For instance, features like human-shelf distance/angle and human velocity are significant for describing observed interactions, while others (e.g., unrelated humans) can be ignored. F-PCMCI allows the robot to exclude unnecessary features, constructing a causal model solely from the crucial ones. This causal model can then be used to better understand and predict the interaction of the involved agents.

To summarise, our approach F-PCMCI is an effective algorithmic solution to select the most meaningful features from a set of variables and build a causal model from such selection. To this end, we significantly enhance speed and accuracy of the causal discovery making it suitable for robotics applications.

II. FILTERED-BASED CAUSAL DISCOVERY

Our Filtered PCMCI (F-PCMCI) [4] extends one of the state-of-the-art causal discovery methods, i.e. PCMCI [6], augmenting it with a feature-selection algorithm that is able to identify the correct subset of variables for causal analysis, starting from a predefined set of them. It comprises two main components: the feature selection step and the causal discovery process. In the feature selection step, we use a Transfer Entropy (TE)-based method to “filter” important features and

¹School of Computer Science, University of Lincoln, UK.

²Dept. of Information Engineering, University of Padua, Italy.

This work has received funding from the EU H2020 research & innovation programme – grant agreement 101017274 (DARKO).




Fig. 2. Pipeline showing the use of F-PCMCI for modeling and predicting spatial interactions in a robot's intralogistics scenario [4].

identify potential associations from a larger set of variables. This filter calculates TE between each variable and the others. If an association is found, the variables are considered in subsequent causal discovery, performed by the PCMCI, to assess their causal relationship [4]. The TE filter therefore provides a set of variables and a hypothetical causal model which contribute to speeding up and improving the accuracy of the causal discovery.

To aid the evaluation and application of our method on real-world problems, we have implemented a Python package of F-PCMCI, which is publicly available on GitHub¹ and PyPI². Additionally, we have provided a comprehensive web page and various tutorials that explain how to use F-PCMCI, starting from input time-series data to the output causal model.

III. ROBOT APPLICATION

We applied our approach to model and predict spatial interactions in a human-robot collaboration scenario for intralogistics [4], as part of a large EU project (DARKO³). As shown in Fig. 2, this involves three main steps: (i) extracting time-series of sensor data (i.e. trajectories) from human spatial interactions; (ii) reconstructing the causal model using F-PCMCI; (iii) embedding the causal model in a neural network, LSTM-based prediction system. We extracted human and robot trajectories from the THÖR dataset [8], which uses helmets and infrared cameras to track agents' motion. The dataset features a shared working area comprising people, target positions, static objects, and a moving robot in a warehouse-like environment. People carry boxes and navigate the environment alone or in groups, simulating intralogistics activities. The robot, a small autonomous forklift, moves in a safe but socially unaware manner, following a predefined path with a maximum speed of 0.34 m/s. It projects its current motion intent onto the floor in front of it using a mounted beamer. We chose this dataset for its diverse range of interactions, including humans, robot, and static objects (Fig. 2, left image). Our LSTM-based prediction system includes an encoder-decoder network with an input-attention layer, which selects the most relevant time steps from the observation window, and a self-attention one to identify the key prediction cues. The latter integrates also the causal discovery output as a non-trainable parameter.

¹<https://github.com/lcastri/fpcmci>

²<https://pypi.org/project/fpcmci>

³<https://cordis.europa.eu/project/id/101017274>

A detailed explanation of the system and experimental results can be found in our recent paper [4]. It compares our F-PCMCI to the standard PCMCI in terms of accuracy and execution time, outperforming the latter in both metrics. It also shows how F-PCMCI's causal model improves the accuracy of the spatial interaction prediction, for a possible robot's intralogistics scenario, compared to a non-causally informed system. Further details about the Python implementation of the prediction system are provided in the companion web page⁴.

IV. CONCLUSION

In this paper, we motivated the need and provided an overview of an efficient causal discovery algorithm for robotics applications. By integrating a feature-selection module based on transfer entropy, our F-PCMCI enhances the accuracy and computational efficiency of causal discovery, making it well-suited for autonomous robot applications. In particular, the paper shows a robot-assisted intralogistics scenario in which the F-PCMCI causal model improves the quality of the prediction of spatial interactions in a warehouse-like environment.

Inspired by recent research on continual learning and causality [7], our future work aims to enhance F-PCMCI with a real-time strategy for adaptive causal modeling. This would empower the robot to continually improve the causal model and adapt to changing scenarios as new data arrives.

REFERENCES

- [1] J. Pearl, *Causality*. Cambridge University Press, 2009.
- [2] T. Hellström, "The relevance of causation in robotics: A review, categorization, and analysis," *Paladyn, Journal of Behavioral Robotics*, vol. 12, no. 1, pp. 238–255, 2021.
- [3] L. Castri, S. Mghames, M. Hanheide, and N. Bellotto, "Causal discovery of dynamic models for predicting human spatial interactions," in *International Conference on Social Robotics (ICSR)*, 2022.
- [4] —, "Enhancing causal discovery from robot sensor data in dynamic scenarios," in *Conference on Causal Learning and Reasoning*, 2023.
- [5] C. Glymour, K. Zhang, and P. Spirtes, "Review of Causal Discovery Methods Based on Graphical Models," *Frontiers in Genetics*, 2019.
- [6] J. Runge, P. Nowack, M. Kretschmer, S. Flaxman, and D. Sejdinovic, "Detecting and quantifying causal associations in large nonlinear time series datasets," *Science Advances*, vol. 5, 2019.
- [7] L. Castri, S. Mghames, and N. Bellotto, "From continual learning to causal discovery in robotics," in *Continual Causality Bridge Program at AAAI23*, 2023.
- [8] A. Rudenko, T. P. Kucner, C. S. Swaminathan, R. T. Chadalavada, K. O. Arras, and A. J. Lilienthal, "Thör: Human-robot navigation data collection and accurate motion trajectories dataset," *IEEE Robotics and Automation Letters*, vol. 5, pp. 676–682, 2020.

⁴https://github.com/lcastri/cmm_ts