From Continual Learning to Causal Discovery in Robotics

Anonymous submission

Abstract

Reconstructing accurate causal models of dynamic systems from time-series of sensor data is a key problem in many real-world scenarios. In this paper, we present an overview based on our experience about practical challenges that the causal analysis encounters when applied to autonomous robots and how Continual Learning (CL) could help to overcome them. We propose a possible way to leverage the CL paradigm to make causal discovery feasible for robotics applications where the computational resources are limited, while at the same time exploiting the robot as an active agent that helps to increase the quality of the reconstructed causal models.

1 Introduction

Causal discovery approaches generally build the causal model of the observed scenario from static or time-series data collected and processed in advance. However, in many real-world robotics applications, this approach could result inefficient or even unfeasible. The link between *Continual Learning* (CL) (Lesort et al. 2020) and *Causality* might represent a stepping stone towards the exploitation of causal discovery algorithms (Glymour, Zhang, and Spirtes 2019) that currently suffer many limitations in autonomous robots.

Causal inference is an active research area in different fields, including robotics and autonomous systems (Hellström 2021; Brawer, Qin, and Scassellati 2020; Cao et al. 2021; Katz et al. 2018; Angelov, Hristov, and Ramamoorthy 2019). However, most of these works overlooked some key features that are important for real-world application, i.e. the computational cost and the memory needed by causal analysis when long time-series are processed to reconstruct a causal model of the observed scenario. To this end, the CL's ability to enable the acquisition of more knowledge by trained models without forgetting previous information, and without using previous data recordings, might help to address these problems and to achieve better result in terms of quality of the causal analysis. For instance, a robot in an automated warehouse with humans and various objects (e.g. see Fig. 1) could observe and intervene in the interactions among them (e.g. worker and shelf) in order to build a causal model and therefore a deep understanding of the situation. Since the limited hardware resources though, the robot's causal analysis might be slow and based on a limited amount of data, leading to a low quality causal model. The solutions suggested in this paper would allow the robot to overcome its hardware limitations and, moreover, to improve the quality of the causal models by continually feeding new data for causal analysis, discarding the old collected one. This would enable a more efficient use of the robot's memory and computing's resources compared to existing causal discovery's approaches. To summarise, this paper proposes a *Causal Robot Discovery* (CRD) approach to overcome 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.

2 Related Work

Causal discovery: Several methods have been developed over the last few decades to derive causal relationships from observational data, which can be categorized into two main classes (Glymour, Zhang, and Spirtes 2019). The first one includes *constraint-based methods*, such as Peter and Clark (PC) and Fast Causal Inference (FCI), which rely on conditional independence tests as constraint-satisfaction to recover the causal graph. The second one includes *score-based methods*, such as Greedy Equivalence Search (GES), which assign a score to each Directed Acyclic Graph (DAG) and perform a search in this score space. However, many of these algorithms work only with static data (i.e. no temporal information) and are not applicable to time-series of sensor data in many robotics applications, for which time-dependent causal discovery methods are instead necessary. To this end, a variation of the PC algorithm, called PCMCI, was adapted and applied to time-series data (Runge 2018; Runge et al. 2019; Saetia, Yoshimura, and Koike 2021).

Causal robotics: Causal inference has been recently considered in robotics, for example to build and learn a Structural Causal Model (SCM) from a mix of observation and self-supervised trials for tool affordance with a humanoid robot (Brawer, Qin, and Scassellati 2020). Other applications include the use of PCMCI to derive the causal model of an underwater robot trying to reach a target position (Cao

et al. 2021) or to predict human spatial interactions in a social robotics context (Castri et al. 2022). Further causality-based approaches can be found in the robot imitation learning and manipulation area (Katz et al. 2018; Angelov, Hristov, and Ramamoorthy 2019; Lee et al. 2021). However, all these solutions rely on a fixed set of time-series for causal analysis and do not consider the computational cost and complexity for online update of the robot's causal models.

Continual learning: The concept of learning continually from experience has been present in artificial intelligence since early days (Weng et al. 2001). Recently this has been explored more systematically in machine learning (Hadsell et al. 2020; Parisi et al. 2019) and robotics (Lungarella et al. 2003; Lesort et al. 2020; Churamani, Kalkan, and Gunes 2020). To our knowledge though, few applications of the continual learning paradigm can be found in the causality field. Javed, White, and Bengio (2020) incorporate causality and continual learning with an online algorithm that continually detects and removes spurious features from a causal model. In (Kummerfeld and Danks 2012, 2013; Kocacoban and Cussens 2019, 2020), instead, algorithms for online causal structure learning are presented to deal with non-stationary data. This is a key feature of data from real-world environments, which is still under-investigated in robotics and therefore motivates our approach proposed next.

3 Causal Robot Discovery

A review of the literature revealed that the possible limitations of autonomous robots doing causal discovery with their own on-board sensors have not been taken into account. Indeed, the computational and memory requirements for long time-series of sensor data are often very demanding, making the use of previous algorithms for causal inference unfeasible on such platforms.

Our approach is partially inspired by the works of Kocacoban and Cussens (2019) for handling non-stationary data, but differs from it in two ways. First of all, we adopt the current state-of-the-art PCMCI method for causal discovery from time-series data; second, we propose to re-learn the causal model not only when the observed scenario changes, but also at each new robot's set of observations/interventions (periodically, e.g. every few minutes). In particular, the introduction of the CL paradigm could help the robot to overcome the challenge of limited hardware resources and to improve the quality of the causal analysis even with non-stationary data. In addition, a CRD approach could benefit from the fact that robots are physically embodied in the environment and can actively influence its dynamic processes (i.e. by performing interventions). That is, CRD could improve the accuracy of the causal model by enriching "passive" observational data from the sensors with "active" interventional data from robot's actions aimed at collecting specific time-series for causal discovery.

Therefore, our goal is to decrease the need of hardware resources – often scarce in autonomous systems – and to increase the causal analysis quality by using the robot as an active agent in the learning process. The proposed CRD system is thought to limit the demand for hardware resources

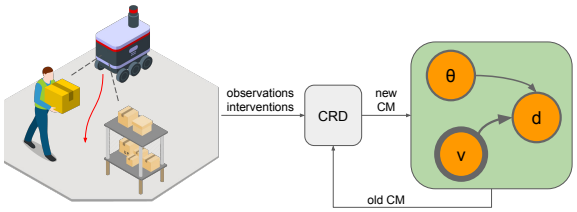

Figure 1: CRD approach: the robot provides observational and interventional data about the human-object interaction to the CRD block. The latter generates a causal model which is stored and used to compare the next causal model built on the next robot's observation and intervention.

and allow the robot to perform high quality causal discovery in a reasonable time by using its own on-board sensor data.

The proposed approach is depicted in Fig. 1: (*i*) starting from a prefixed set of variables, the robot collects meaningful data by observing and intervening in the target scenario; (*ii*) based on this data, a causal model is estimated using PCMCI (Runge 2018), which computes test statistics and p-values as causal strengths of the DAG's links. At this stage, differently from (Kocacoban and Cussens 2019), to increase the accuracy of the causal discovery, the robot keeps on collecting data by observing and intervening in the scenario to create new causal models. Periodically then, the robot compares the new causal models with the old ones, inheriting from the latter only the links that minimise the p-values of the DAG's causal relations. By repeating this process until the observed scenario changes, a stable version of the causal model with minimum uncertainty levels would be reached. The latter is useful not only for modeling the current scenario, but also when it changes. Indeed, by partially using the stored causal model that is consistent the target scenario, even the latter changes, the estimation of the new model can be significantly sped up (Kocacoban and Cussens 2019).

Note that by iteratively discarding the old time-series and storing only the built causal model helps to avoid the combinatorial explosion otherwise affecting the PCMCI discovery, therefore allowing the robot to operate and compute new models within reasonable time. The catastrophic-forgetting problem is also mitigated by the fact that, during continuous operations, the robot observes similar processes with only small incremental changes, which leads to sequences of similar causal models, and that it reconstructs the latter by exploiting relatively small variations of the previous models.

4 Conclusion

In this paper we considered the hardware resource limitations of autonomous robots, which are crucial to perform causal inference, and proposed a new approach for causal robot discovery to overcome some of the main challenges. This includes improving the quality of the causal models by using the robot as an active agent in the learning process. Future work will be devoted to the implementation and application of this approach to real-wold robotics problems, with a special interest in industrial scenarios involving human-robot interaction and collaboration.

References

Angelov, D.; Hristov, Y.; and Ramamoorthy, S. 2019. Using causal analysis to learn specifications from task demonstrations. In *Proc. of the Int. Joint Conf. on Autonomous Agents and Multiagent Systems, AAMAS*.

Brawer, J.; Qin, M.; and Scassellati, B. 2020. A causal approach to tool affordance learning. In *IEEE/RSJ Int. Conf. on Intell. Robots & Systems (IROS)*, 8394–8399.

Cao, Y.; Li, B.; Li, Q.; Stokes, A.; Ingram, D.; and Kiprakis, A. 2021. Reasoning Operational Decisions for Robots via Time Series Causal Inference. In *2021 IEEE Int. Conf. on Robotics and Automation (ICRA)*, 6124–6131.

Castri, L.; Mghames, S.; Hanheide, M.; and Bellotto, N. 2022. Causal Discovery of Dynamic Models for Predicting Human Spatial Interactions. In *International Conference on Social Robotics (ICSR)*.

Churamani, N.; Kalkan, S.; and Gunes, H. 2020. Continual learning for affective robotics: Why, what and how? In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 425–431. IEEE.

Glymour, C.; Zhang, K.; and Spirtes, P. 2019. Review of Causal Discovery Methods Based on Graphical Models. *Frontiers in Genetics*.

Hadsell, R.; Rao, D.; Rusu, A. A.; and Pascanu, R. 2020. Embracing change: Continual learning in deep neural networks. *Trends in cognitive sciences*, 24(12): 1028–1040.

Hellström, T. 2021. The relevance of causation in robotics: A review, categorization, and analysis. *Paladyn, Journal of Behavioral Robotics*, 238–255.

Javed, K.; White, M.; and Bengio, Y. 2020. Learning Causal Models Online. *CoRR*, abs/2006.07461.

Katz, G.; Huang, D. W.; Hauge, T.; Gentili, R.; and Reggia, J. 2018. A novel parsimonious cause-effect reasoning algorithm for robot imitation and plan recognition. *IEEE Trans. on Cognitive and Developmental Systems*.

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.

Kocacoban, D.; and Cussens, J. 2020. Fast Online Learning in the Presence of Latent Variables. *Digitale Welt*, 4(1): 37–42.

Kummerfeld, E.; and Danks, D. 2012. Online learning of time-varying causal structures. In *UAI workshop on causal structure learning*.

Kummerfeld, E.; and Danks, D. 2013. Tracking time-varying graphical structure. *Advances in neural information processing systems*, 26.

Lee, T. E.; Zhao, J. A.; Sawhney, A. S.; Girdhar, S.; and Kroemer, O. 2021. Causal Reasoning in Simulation for Structure and Transfer Learning of Robot Manipulation Policies. In *2021 IEEE Int. Conf. on Robotics and Automation (ICRA)*, 4776–4782.

Lesort, T.; Lomonaco, V.; Stoian, A.; Maltoni, D.; Filliat, D.; and Díaz-Rodríguez, N. 2020. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information Fusion*, 58: 52–68.

Lungarella, M.; Metta, G.; Pfeifer, R.; and Sandini, G. 2003. Developmental robotics: a survey. *Connection science*, 15(4): 151–190.

Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113: 54–71.

Runge, J. 2018. Causal network reconstruction from time series: From theoretical assumptions to practical estimation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28: 075310.

Runge, J.; Nowack, P.; Kretschmer, M.; Flaxman, S.; and Sejdinovic, D. 2019. Detecting and quantifying causal associations in large nonlinear time series datasets. *Science Advances*, 5.

Saetia, S.; Yoshimura, N.; and Koike, Y. 2021. Constructing Brain Connectivity Model Using Causal Network Reconstruction Approach. *Frontiers in Neuroinformatics*, 15: 5.

Weng, J.; McClelland, J.; Pentland, A.; Sporns, O.; Stockman, I.; Sur, M.; and Thelen, E. 2001. Autonomous mental development by robots and animals. *Science*, 291(5504): 599–600.