A Causal Approach to Prescriptive Process Monitoring

Zahra Dasht Bozorgi[1]

[1]University of Melbourne, Level 4, 700 Swanston St, Carlton, 3053, VIC, Australia

1. Introduction

A business process is a collection of events, activities, and decisions that collectively lead to an outcome that can be of value to a customer [1]. Process mining is a family of techniques that extract information about processes using historical process execution data, generally known as event logs [2]. Recent advances in process mining techniques have allowed companies to manage and improve their processes more efficiently. Particularly, process mining has benefited from advances in machine learning techniques to provide accurate predictions of the future state of business processes in an area known as *predictive process monitoring*. Recent predictive monitoring methods can produce highly accurate predictions of upcoming events, future event suffixes, remaining time, and the outcome of the process [3, 4, 5, 6, 7, 8]. However, having reliable predictions does not always lead to improvement of the process. A study by Dees *et al.* [9], shows that if good predictions are followed by bad recommendations, the desired improvement is not achieved. This study illustrates the need for new techniques that find the best interventions based on the context of each case.

Causal inference is a field in statistics that is concerned with estimating causal effect of a variable when that variable is changed. This field has been widely used in other domains such as medicine, social sciences, and marketing. In medicine, causal inference techniques are used to determine the effectiveness of new drugs. In social sciences and marketing, new policies and campaigns are evaluated using state-of-the-art causal inference methods. The success of causal inference in these fields serves as a motivation to use these techniques in the area of process mining.

In this Ph.D. project, we aim to address the need for effective recommendations by devising new prescriptive process monitoring methods which are based on *causal relationships*. These methods will make recommendations, at tactical and operational levels, about what actions should be taken to achieve a given process objective.

BPM 2021: Proceedings of the Demonstration & Resources Track, Best BPM Dissertation Award, and Doctoral Consortium at BPM 2021 co-located with the 19th International Conference on Business Process Management, September 6-10, 2021, Rome, Italy

zdashtbozorg@student.unimelb.edu.au (Z. D. Bozorgi)
https://www.linkedin.com/in/zahradbozorgi/ (Z. D. Bozorgi)

2. Research Problem

Starting from a set of pre-defined treatments (a.k.a. process interventions) and associated costs, this thesis explores the use of causal inference in prescriptive monitoring, to recommend what treatments are to be used in what context such that a given benefit function is maximised for the organisation. Specifically, the overall aim of this project is to answer the following question:

How to provide cost-aware recommendations based on causal relations between a proposed intervention and a target of interest?

Since the area of prescriptive process monitoring remains largely under-explored, we select three mains dimensions of exploration for scoping purposes. These are:

- The objective of the recommendations. In this thesis, we focus on two main goals, namely outcome improvement and remaining time reduction.
- Explainability of the recommendations. We explore the use of both interpretable and black-box (hard to explain) methods for producing recommendations and assess relative merits and shortcomings.
- Prescriptive monitoring with multiple treatment options. We explore the use of reinforcement learning solutions to choose the best course of action during the execution of the process.

Accordingly, we define the following research questions to explore these dimensions:

1. RQ1: How can the use of causal inference in prescriptive process monitoring result in recommendations that optimise process outcome and cycle time while providing measurable benefits to the organisation?
2. RQ2: How do black-box approaches for causal inference compare with interpretable techniques when applied to prescriptive process monitoring?
3. RQ3: How can we use reinforcement learning in combination with causal inference to devise a prescriptive process monitoring method that can select the best actions that lead to the achievement of an objective?

3. Approach

To address the identified research problems, we propose three studies each addressing one research question. To the best of our knowledge, this is the first time causal inference has been used in a prescriptive process monitoring context.

3.1. RQ1: Recommendations for varying objectives

The main purpose of this study is to investigate the use of causal inference for prescriptive monitoring. We hypothesise that if we can establish a causal relationship between a possible intervention and the outcome of interest, the recommendations that result from such causal relationships will result in more benefit for the company than correlation-based recommendations. We divided this study into two sub-studies, one addressing outcome improvement and the other cycle time reduction. We chose this two objectives because one is an example of a binary outcome, while the other is described as a real-valued attribute.

Our approach proposed in [10] is a rule-based prescriptive system. First, a set of interventions which are highly correlated with a positive outcome of interest are identified. Then using a causal effect estimation method called Uplift Tree [11], we identify contexts in which those interventions causally influence the outcome. We also propose a cost model that identifies the Return-on-Investment (ROI) of an intervention. The decision to intervene in a process is then made based on the ROI. We show that using this method, we can automatically identify improvement actions that are traditionally done by humans.

In another study [12], we propose a cost-aware prescriptive monitoring method that is designed to reduce cycle time. The core of this recommendation system is the Orthogonal Random Forest (ORF) [13], which is a causal estimation method that works with continuous outcome variables. The approach consists of two phases: an offline phase in which given an event log, a causal effect estimation model is trained, and the best policy for applying the intervention is selected to maximise the gain in applying it. The second phase is the online phase where the causal effect estimator and the selected policy are used to determine which on-going cases should receive an intervention. Our results show that selecting an intervention policy based on causal models leads to a higher net-gain than policies based on traditional machine learning methods for prediction.

3.2. RQ2: Explainability of the recommendations

In the second study, we plan to investigate the use of black-box and white-box models. Many methods for causal estimation have been proposed in the causal inference literature. Most of these models are black-box, meaning that we cannot understand how different variables are combined to estimate a causal effect. Conversely, some models such as uplift trees are interpretable or white-box. These models are constrained to give a better understanding of how the causal effect estimation is made. Many models, however, are not designed to be interpretable, but to be accurate estimators of causal effects. Due to an issue known as *the fundamental problem of causal inference*, it is very difficult to establish whether black-box models have higher accuracy in real-world settings and whether the more accurate estimates yield higher benefit. Since prescriptive monitoring methods are designed to aid humans in decision-making, transparency of the models they are based on is important. Therefore, in this study, we plan to investigate whether the use of black-box models is preferable to white-box ones. To do this, we first need to address the problem of not observing ground truth in the data. So we will first create simulated event logs by using the data generation approach proposed in [14] based on a flow-based deep generative model called Sigmoidal Flow. We will train a Sigmoidal Flow on real-life event logs, and then create simulated data that contain the ground truth causal effect while preserving the statistical properties of our real-world data. Using this simulated data we will investigate different black-box and white-box causal inference approaches and develop a cost-benefit model to determine in which circumstances using a black-box model is warranted.

3.3. RQ3: Multiple treatment options

This study will be conducted in the final year of this Ph.D. In the previous studies, we assume that there is a binary treatment that can potentially influence a process performance metric. However, in practice, many actions can be taken during the execution of the process that influence performance metrics. With the methods proposed in the first study, many causal models need to be trained separately, and this is computationally inefficient. Furthermore, using many standalone causal models will not consider how the different actions influence each other (i.e., the action interplay). In such circumstances, reinforcement learning frameworks can be employed to design a prescriptive monitoring method that deals with a multitude of actions/treatments. In particular, *contextual bandits algorithms* have benefited from the causal inference literature to make them less prone to problems in estimation bias [15]. Contextual bandits are an extension of multi-armed bandits. They output an action conditional on the state of the environment. The aim of this study is to devise a new prescriptive monitoring method based on contextual bandits that can recommend the best action at each step of the process execution. In a prescriptive monitoring method based on contextual bandits, the context will be defined by case prefixes extracted from event logs and ongoing process executions, and the reward will be a function based on one or multiple process performance metrics.

3.4. Methodology

This project will follow a Design Science research method [16]. The rigour of the approaches will be ensured by conducting an extensive literature review and constructing a comprehensive evaluation benchmark, using well defined selection and assessment criteria. The relevance of the solutions will be ensured via an extensive evaluation of the developed techniques with real-life and simulated data sets, and where possible, through case studies with relevant organisations.

4. Limitations of the Study

The main limitation of this study is that we assume that all the variables that influence intervention and outcome are observed in the event log. This might lead to biased estimations that cause the prescriptive system to make a sub-optimal recommendation. We try to alleviate this threat to validity by performing sensitivity analysis on our models to measure their robustness to unmeasured confounding. Another limitation of this study is the evaluation of the proposed methods. While we use cost models and other evaluation methods in the causal inference literature such as Qini coefficients to evaluate our prescriptive methods, the best and most rigorous way of evaluating these methods is conducting an A/B test.

References

[1] M. Dumas, M. L. Rosa, J. Mendling, H. A. Reijers, Fundamentals of Business Process Management, Second Edition, Springer, 2018.

[2] W. M. P. v. d. Aalst, Process Mining: Data Science in Action, second ed., Springer, 2016.

[3] F. Taymouri, M. La Rosa, S. Erfani, Z. D. Bozorgi, I. Verenich, Predictive business process monitoring via generative adversarial nets: The case of next event prediction, in: International Conference on Business Process Management, Springer, 2020.

[4] F. Taymouri, M. L. Rosa, S. M. Erfani, A deep adversarial model for suffix and remaining time prediction of event sequences, CoRR (2021).

[5] V. Pasquadibisceglie, A. Appice, G. Castellano, D. Malerba, Using convolutional neural networks for predictive process analytics, in: International Conference on Process Mining, ICPM 2019, Aachen, Germany, June 24-26, 2019, IEEE, 2019.

[6] N. Tax, I. Verenich, M. L. Rosa, M. Dumas, Predictive business process monitoring with LSTM neural networks, in: Advanced Information Systems Engineering - 29th International Conference, CAiSE 2017, Proceedings, Springer, 2017.

[7] M. Camargo, M. Dumas, O. G. Rojas, Learning accurate LSTM models of business processes, in: T. T. Hildebrandt, B. F. van Dongen, M. Röglinger, J. Mendling (Eds.), Business Process Management - 17th International Conference, BPM 2019, Vienna, Austria, September 1-6, 2019, Proceedings, Springer, 2019.

[8] L. Lin, L. Wen, J. Wang, Mm-pred: A deep predictive model for multi-attribute event sequence, in: T. Y. Berger-Wolf, N. V. Chawla (Eds.), Proceedings of the 2019 SIAM International Conference on Data Mining, SDM 2019, Calgary, Alberta, Canada, May 2-4, 2019, SIAM, 2019.

[9] M. Dees, M. de Leoni, W. M. P. van der Aalst, H. A. Reijers, What if process predictions are not followed by good recommendations?, in: Proceedings of the Industry Forum at BPM 2019 co-located with 17th International Conference on Business Process Management (BPM 2019), Vienna, Austria, September 1-6, 2019, CEUR-WS.org, 2019.

[10] Z. D. Bozorgi, I. Teinemaa, M. Dumas, M. L. Rosa, A. Polyvyanyy, Process mining meets causal machine learning: Discovering causal rules from event logs, in: 2nd International Conference on Process Mining, ICPM 2020, Padua, Italy, October 4-9, 2020, IEEE, 2020.

[11] P. Rzepakowski, S. Jaroszewicz, Decision trees for uplift modeling with single and multiple treatments, Knowl. Inf. Syst. (2012).

[12] Z. D. Bozorgi, I. Teinemaa, M. Dumas, M. La Rosa, Prescriptive process monitoring for cost-aware cycle time reduction, arXiv preprint arXiv:2105.07111 (2021).

[13] M. Oprescu, V. Syrgkanis, Z. S. Wu, Orthogonal random forest for causal inference, in: Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, PMLR, 2019.

[14] B. Neal, C.-W. Huang, S. Raghupathi, Realcause: Realistic causal inference benchmarking, arXiv preprint arXiv:2011.15007 (2020).

[15] M. Dimakopoulou, Z. Zhou, S. Athey, G. Imbens, Balanced linear contextual bandits, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2019.

[16] R. H. von Alan, S. T. March, J. Park, S. Ram, Design science in information systems research, MIS quarterly (2004).