

Explainable and Responsible Process Prescriptive Analytics (Extended Abstract)

Alessandro Padella¹

¹University of Padova

Abstract

Within the realm of Process Mining, Process-Aware Recommender systems (PAR systems) are information systems that aim to monitor process executions, predict their future behaviour, and find optimal corrective actions to reduce the risk of failure or to maximize a given reference Key Performance Indicator (KPI). The PAR system comprises Predictive Analytics and Prescriptive Analytics. The second part regards providing recommendations for fixing the execution of processes that are predicted to have undesired KPI values. While the research has focused on generating recommendations that aim to have better and better KPIs, my Ph.D. aims to combine the high scores of recommendations with their feasibility and fairness.

Keywords

Process Mining, Prescriptive Analytics, Recommender Systems, Explainable AI, Process Improvement, Bias removal

1. Introduction

In the context of Process Mining, Process-Aware Recommender systems (hereafter shortened as PAR systems) are a class of Information Systems which aims to *predict* how the process instances are going to end and eventually *recommend* the corrective actions for improving their execution.

This has been translated into developing frameworks that predict the outcome for each running instance of a process and suggest the next-best activity to perform, aiming to improve it when this is not satisfactory. In a general sense, an outcome is measured through the so-called Key Performance Indicator (KPI) [1, 2, 3].


While several proposals have recently been put forward to provide effective recommendations (cf. Section 3), less attention has been paid to ensure their fairness, practical feasibility wrt. the process constraints, and their comprehensibility for better engagement of process participants.


This project aims to address these additional aspects.

Bolzano 2022, International Conference on Process Mining, October 23-28, Bolzano, IT

 alessandro.padella@phd.unipd.it (A. Padella)

 <https://github.com/Pado123> (A. Padella)

 © 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

2. Research Questions

As mentioned, a PAR system's starting point is defining the process outcomes. The outcome of a process is often measured through a customizable KPI function that, given an execution recorded in a log trace (which describes the life-cycle of a particular *process instance*, i.e., a case), looks at the activities executed and attribute values, returning a KPI value (e.g. the total execution time or the total cost of the procedure).

Several PAR systems are limited to providing recommendations that only consist in suggesting what activity to do as next, irrespectively whether the activity is possible in that moment in time. Indeed, human process actors and other types of resources may not be available. Furthermore, the impact of the recommended activity on the process outcome may depend on the actual resource that performs the activity. From this, the following research question arises.

Research question 1: *How can we build a prescriptive business process analytics block that can provide more feasible recommendations?*

We indeed want to design a recommendation framework for Prescriptive Analytics to provide recommendations on combinations of activities and resources. It should analyze the past executions of the process and exploit some machine (or deep) learning techniques for generating recommendations.

However, from the perspective of business reality, there is a wrong assumption: past executions can be repeated at the moment in which we provide recommendations.

Each company has its own set of rules and protocols and a pool of human resources. Over time and as the company evolves, these may change (e.g. some employees may resign, or some machines may be changed).

The challenge then becomes to make the framework capable of understanding and incorporating changes in organisational structures, thus avoiding making recommendations that are not actually executable. This leads to the second research question.

Research question 2: *How can we ensure that the context in which recommendations are generated and the context in which they are provided are compliant?*

Works [4] illustrate how some datasets, and so the models built on them, may be biased, especially regarding race and gender. The same problem can occur for PAR systems: it may, e.g., recommend corrective actions on cases with given characteristics. Also, PAR systems might decide to allocate activities to only certain process actors, who have better performances, thus ultimately causing them to be overloaded while others are seldom employed. This generates our third research question.

Research question 3: *Are the recommendations we are providing fair?*

Solving the research questions above, we focused on the quality of recommendations, answering the question "*Who should perform what, for improving our KPI?*", trying to have an efficient and responsible framework.

In addition, it is also crucial to accompany recommendations with an explanation of the rationale that brought the system to suggest this way. This increases the engagement and trust of the actors in the system, and thus the willingness to follow what is suggested. This brings to the fourth research question:

Research question 4: *How can we increase trust in our recommendations?*

We want to use some Explainable AI techniques to provide some explanations about not only

what we are suggesting but also *why* the algorithm took that decision. This can answer the question “*Why should a resource perform the suggested activity?*”. Finally, we want to understand whether the explanations provided help process actors trust our framework more. In doing so, we plan to perform some human evaluations with actual process actors, as well as we aim to use objective metrics to assess the explanation quality, e.g. discussed in [5].

Research question 5: *Does the real users feel helped by our work?*

3. Project Roadmap

We have initially started addressing the first and the fourth research questions, integrating explanations in the prescriptive-analytics framework proposed in [1] and leveraging on Cat-Boost [6], a high-performance open source framework that has shown to provide more accurate predictions and with limited computation time if compared with the literature (see [7]).

Explanations are given using the theory of Shapley values [8]. However, the explanations of the recommendations are different from those of predictions as discussed in the literature (see e.g. [3]). Indeed, the explanation of a prediction is translated into explaining how much each variable influences the final KPI value. For example, “*The fact that the variable `customer_type` assumes the value `Gold` contributes to decreasing the expected total time of the process instance by 120 hours*”. On the other hand, the explanation of a recommendation is related to the recommended activity. Indeed, given the suggestion of a certain activity, the explanation represents how much performing that activity changes the contribution of each variable on the final KPI value. For example, “*Performing the activity `Send Letter`, the fact that the variable `customer_type` assumes the value `Gold` goes from contributing to decreasing the expected total time by 120 hours to contributing to decreasing it by 230 hours*.”

Our work will be organized as follows:

- Regarding research question 1, the literature in [2, 9] proposes other interesting Prescriptive-Analytics frameworks. However, just focus on recommending activities. We plan to test the goodness of recommendations using cases from the past for which compensatory actions are known and check whether the system recommends them.
- For research question 2, we plan to analyze the work of De Smedt et al. in [10]. It presents a framework which aims to forecast the entire process model from historical event data, representing event data as multiple time series. We will start with this work and try to exploit similarity techniques for time-series data like Grid Representation and Matrix Distances from [11] for detecting the variation of the companies’ structure and resources.
- Regarding research question 3, we aim to follow the directions indicated by Mannhardt in [12] and by van der Aalst in [13]. Specifically, we plan to specialize and adjust the de-biasing technique of *Adversarial debiasing*, discussed in [14]. His goal is to make the model-independent with respect to a certain variable. In it, two predictive models are trained: the first predicts your desired value, and the second takes as input the output of the first and infers what is the value of the variable whose influence we want to remove. The goal of the model-building phase will be to have the first classifier have the highest possible accuracy while that of the second has to be equivalent to that of a random, making the model capable of providing responsible recommendations.

- For Research question 5, we plan to work on a graphical user interface capable of representing the recommendations and their relative explanations. The literature in [5] proposes the approach of *appropriate trust*: an objective method which may overcome the problem of subjective evaluations of the explanations regarding the black-box models. Furthermore, we plan to combine this type of objective evaluation with a full evaluation through its associated graphical interface. From this, both real users and domain expert will test the system to which will be associated a satisfaction survey.

Acknowledgments

Acknowledgement. My PhD. scholarship is partly funded by IBM Italy, and by the BMCS Doctoral Program, University of Padua. This research is also supported by the Department of Mathematics, University of Padua, through the BIRD project “Data-driven Business Process Improvement” (code BIRD215924/21).

References

- [1] M. de Leoni, M. Dees, L. Reulink, Design and evaluation of a process-aware recommender system based on prescriptive analytics, in: 2020 2nd International Conference on Process Mining (ICPM), 2020.
- [2] S. Weinzierl, S. Dunzer, S. Zilker, M. Matzner, Prescriptive business process monitoring for recommending next best actions, in: Business Process Management Forum, 2020.
- [3] R. Galanti, B. Coma-Puig, M. de Leoni, J. Carmona, N. Navarin, Explainable predictive process monitoring, in: Proceedings of the 2nd International Conference on Process Mining (ICPM 2020), 2020.
- [4] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: Proceedings of the 1st Conference on Fairness, Accountability and Transparency, 2018.
- [5] H. Löfström, K. Hammar, U. Johansson, A meta survey of quality evaluation criteria in explanation methods, in: Proceedings of CAiSE Forum 2022, 2022.
- [6] A. V. Dorogush, V. Ershov, A. Gulin, Catboost: gradient boosting with categorical features support, in: Proceedings of the Workshop on ML Systems at NIPS 2017, 2017.
- [7] R. Galanti, M. de Leoni, M. Monaro, N. Navarin, A. Marazzi, B. Stasi, S. Maldera, An explainable decision support system for predictive process analytics, 2022. doi:10.48550/arXiv.2207.12782.
- [8] C. Molnar, Interpretable Machine Learning, 2022. Available online at <https://christophm.github.io/interpretable-ml-book/>.
- [9] Z. D. Bozorgi, I. Teinemaa, M. Dumas, M. L. Rosa, A. Polyvyanyy, Prescriptive Process Monitoring for Cost-Aware Cycle Time Reduction, in: 3rd International Conference on Process Mining (ICPM), 2021.
- [10] J. De Smedt, A. Yeshchenko, A. Polyvyanyy, J. Weerdt, J. Mendling, Process Model Forecasting Using Time Series Analysis of Event Sequence Data, 2021, pp. 47–61. doi:10.1007/978-3-030-89022-3_5.

- [11] Y. Ye, J. Jiang, B. Ge, Y. Dou, K. Yang, Similarity measures for time series data classification using grid representation and matrix distance, Knowledge and Information Systems (2019).
- [12] F. Mannhardt, Responsible Process Mining, Springer International Publishing, 2022.
- [13] W. Aalst, Responsible data science: Using event data in a “people friendly” manner, 2017. doi:10.1007/978-3-319-62386-3_1.
- [14] D. Kenna, Using adversarial debiasing to remove bias from word embeddings, CoRR (2021).