

Explainable Predictive Monitoring of Temporal Measures of Business Processes

Ilya Verenich^{1,2}

¹ Queensland University of Technology, Brisbane, Australia

² University of Melbourne, Melbourne, Australia

ilya.verenich@unimelb.edu.au

Keywords: Business Process Mining · Explainable Artificial Intelligence · Predictive Monitoring · Transparent Models · Machine Learning · Deep Learning · Flow Analysis.

1 Introduction

Business process monitoring is concerned with the analysis of events produced during the execution of a business process in order to assess the fulfillment of compliance requirements and performance objectives [2]. Process monitoring can take place *offline*, on a tactical level via periodically generated reports, or *online*, on an operational level via dashboards displaying the performance of ongoing process instances (*cases*) in terms of performance indicators such as cycle time, resource utilization, and defect rate. These traditional process monitoring techniques allow analysts and managers to identify deviating patterns in the performance of a process. However, these techniques are in fact reactive, in the sense that they detect process issues only once they have happened, and hence they do not allow users to act in a preventive manner.

Advances in machine learning combined with the availability of process execution data have led to the emergence of *predictive process monitoring* techniques. Predictive business process monitoring refers to a family of online process monitoring techniques that seek to predict the future state or properties of ongoing (incomplete) cases, stored in an event stream. These predictions are based on models extracted from historical (complete) cases recorded in event logs.

In the past years, numerous approaches have been proposed to address typical prediction targets, such as compliance violations [5,8], the next activity or the remaining sequence of activities of a case [3,11], or quantitative process performance indicators such as the remaining cycle time of a case [1,9]. These predictions can be used to alert process workers to problematic cases or to support resource allocation decisions. For example, if the cycle time of a case is expected to exceed a given threshold stipulated by a service-level agreement (e.g., 10 working days), this prediction can be used to inform the customer of the expected delay in advance or to allocate a dedicated resource to this case in order to put it back on track towards timely completion.

This thesis focuses on predicting *temporal* measures, i.e. performance measures that deal with time. Temporal performance measures include cycle time (also known as throughput time, lead time, or case duration), processing time and waiting time.

The problem of predictive process monitoring can be unfolded as shown in Figure 1. A *prediction point* is the point in time where the prediction takes place. A *predicted point* is a point in the future where the performance measure has the predicted value. A prediction is thus based on the predictor’s knowledge of the history of the process until the prediction point as well as knowledge of the future until the predicted point. The former is warranted by the predictor’s *memory* while the latter is based on the predictor’s *forecast*, i.e., predicting the future based on trend and seasonal pattern analysis. Finally, the prediction is performed based on a *prediction algorithm*.

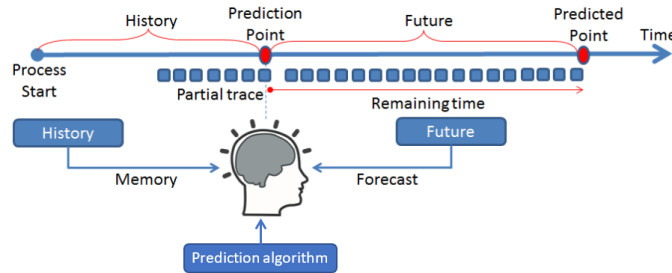


Fig. 1. Overview of predictive process monitoring.

2 Summary of Contributions

As mentioned earlier, a wide range of predictive process monitoring approaches have been proposed in related work. However, there is no unified approach to evaluate existing business process monitoring methods, as different authors have used different datasets, experimental settings, evaluation measures, and baselines to assess their proposed methods. As such, there is a wide range of disparate methods that have not been compared to each other in a uniform setting.

To address this gap, this thesis starts with a systematic literature review of methods for predictive monitoring of temporal measures of business processes. The review was conducted using a range of search keywords and exclusion criteria, and retrieved 25 relevant studies containing unique contribution [13]. These studies were surveyed along several classification dimensions, such as the type of input data, awareness of the underlying business process, family of the algorithm employed, evaluation scenario and availability of the implementation. As a result, a taxonomy of predictive monitoring methods was proposed based on three main aspects: (i) the type of input data required; (ii) the process-awareness of the methods; (iii) and the algorithms employed to derive a predictive model

from the data. It was found that most studies employ discriminative machine learning algorithms to train predictive models, apply them at runtime.

A key gap in the existing body of methods, identified through the literature review, is the lack of support for *explainability* in existing approaches. Indeed, a major drawback of black-box prediction methods in practical deployments is that users have difficulties trusting the prediction when the reasoning behind the prediction is unclear [10]. In real-life applications, users do not only need to get predictions, but they also need to be able to act upon these predictions. In turn, to do so, they need to understand the rationale for these predictions. The decision on how to respond to a prediction is largely driven by the user’s interpretation of the predictive model’s decision logic. If this interpretation matches the intuition of a human decision maker, it might enhance trust in the model [10].

A recent survey by Guidotti et al. [4] provides a comprehensive review of methods for explaining black box models, as well as presents a taxonomy of explainability methods according to the type of problem faced.

The first group of research proposals focus on *explaining* how the model makes certain predictions. To this end, reverse engineering is typically exploited to understand the black box model. These approaches are also known as *post-hoc explainability*, and they provide interpretations based on natural language explanations, visualizations or by examples. Thus, we can separate two processes – decision making (prediction) and decision explaining. Post-hoc interpretations are also most suitable to provide intuition for complex non-transparent models, such as deep neural networks [10]. For example, Evermann et al. [3] made an investigation into understanding the behavior of deep neural networks, using network hallucinations and steady-state analysis.

Rather than providing post-hoc explanations for a black-box classifier, the second group of proposals focus on directly designing a *transparent* classifier that solves the same classification problem. This principle, known as “Transparent Box Design”, is used in this thesis. Specifically, our approach is grounded in the field of transparent box design for *local* interpretability, i.e. we provide explanations for individual predictions, rather than globally for all predictions. To this end, we apply the idea of *decompositional explanation* [6,7] to enhance understanding of a prediction by decomposing it into its elementary components.

In one embodiment of this general idea, we used deep learning models, specifically RNNs and LSTMs, that had been shown to achieve higher levels of accuracy in predictive monitoring, owing to their ability to learn relevant features from trace prefixes automatically. In order to make these models explainable, we used an instance of multi-task learning where several related predictive monitoring tasks were integrated into a single system which is trained jointly, as a way to decompose the predictions. Namely, our technique provides a prediction of the temporal measure in terms of the predicted temporal measures of each of the activities in the most likely remaining path of the current trace.

In a second embodiment, we propose an explainable predictive process monitoring method by automatically discovering a BPMN process model from the event log [14]. An ongoing case is then replayed on this process model and for each

activity and each decision gateway reachable from the current execution state, we estimate their performance measure and branching probabilities respectively, using pre-trained regressors and classifiers. These estimations are then aggregated at the level of a process instance via flow analysis techniques, to produce a final prediction.

We conducted a comparative evaluation of two proposed explainable approaches with each other and with various state-of-the-art baselines, using a unified experimental set-up and 17 real-life event logs originating from different business domains. The results of the evaluation highlighted the trade-off between the prediction accuracy and explainability of predictions in the context of predictive monitoring of temporal measures. The most accurate predictions were typically obtained via black-box models. Specifically, in 13 out of 17 datasets, LSTM-based models trained in a single-task learning setting achieve the best accuracy. At the same time, explainable models are generally less accurate. This trade-off should be considered by business process analysts when choosing a suitable prediction method. Another observation yielded by our evaluation is the correlation between structural characteristics of event logs and the performance of various prefix bucketing and encoding combinations. As such, the former is also an important factor when selecting a prediction method.

The research contributions of the thesis have been consolidated into an open-source toolset for predictive business process monitoring, namely Nirdizati¹. It is a Web-based predictive business process monitoring engine that can be used to train predictive models using the methods described in this thesis, as well as third-party methods, and then to make predictions at runtime, for ongoing process instances. Currently, the tool is able to predict remaining time of process instances, their binary outcome, the next activity to be executed as well as various static case attributes. Predictions can be presented visually in a dashboard or exported for periodic reporting.

Nirdizati has been validated in collaboration with industry partners, providing important feedback as to how this technology is likely to be used in an enterprise setting. One important observation noted in our validation is that the designed predictive monitoring platform shall be used in conjunction with other business intelligence (BI) or monitoring systems. To this end, the platform should offer an API to allow integration with other tools. This task is beyond the scope of the thesis; however, we provide functionality to export the predictions in a CSV file, for importing into third-party BI tools. In this way, analysts will be able to build customized dashboards to highlight the required measures.

Another important lesson learned by testing Nirdizati is that there exist several use cases for predictive process monitoring in practice. One of them is the *real-time dashboard* use case where users can see the current predictions updated continuously. This is especially relevant for very dynamic processes with high throughput. For slower processes, the *regular reports* use case may be sufficient where a user wants to get reports on a regular basis with the current set of predictions. These predictions are possibly filtered so that the report focuses

¹ <http://nirdizati.org>

on the cases that are most likely to become deviant and therefore need attention. Finally, the *alarm* use case may be helpful, where users only want to know about a prediction when a case is likely to be deviant.

The last but not least contribution of the thesis is the application of the proposed explainable predictive monitoring methods to improve the efficiency of business processes [12]. Specifically, we showed how reordering certain activities at runtime, based on the predicted branching probabilities and cycle times of individual activities determined via predictive models, can reduce the unnecessary process work (overprocessing waste), and therefore, the overall process cost.

References

1. van der Aalst, W.M.P., Schonenberg, M.H., Song, M.: Time prediction based on process mining. *Information Systems* **36**(2), 450–475 (2011)
2. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A.: *Fundamentals of Business Process Management*, Second Edition. Springer (2018)
3. Evermann, J., Rehse, J., Fettke, P.: Predicting process behaviour using deep learning. *Decision Support Systems* **100**, 129–140 (2017)
4. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. *ACM Comput. Surv.* **51**(5), 93:1–93:42 (2019)
5. Leontjeva, A., Conforti, R., Di Francescomarino, C., Dumas, M., Maggi, F.: Complex symbolic sequence encodings for predictive monitoring of business processes. In: *Business Process Management Conference*, Proceedings. pp. 297–313 (2015)
6. Levy, A.: *Machine-likeness and explanation by decomposition*. Ann Arbor, MI: Michigan Publishing, University of Michigan Library (2014)
7. Lipton, Z.C.: The mythos of model interpretability. *Commun. ACM* **61**(10), 36–43 (2018)
8. Maggi, F., Di Francescomarino, C., Dumas, M., Ghidini, C.: Predictive monitoring of business processes. In: *CAiSE 2014*. Proceedings. pp. 457–472 (2014)
9. Rogge-Solti, A., Weske, M.: Prediction of business process durations using non-markovian stochastic petri nets. *Information Systems* **54**, 1–14 (2015)
10. Schneider, J., Handali, J.P., vom Brocke, J.: Increasing trust in (big) data analytics. In: *CAiSE 2018 International Workshops*. pp. 70–84 (2018)
11. Tax, N., Verenich, I., La Rosa, M., Dumas, M.: Predictive business process monitoring with LSTM neural networks. In: *CAiSE 2017*, Proceedings. pp. 477–492 (2017)
12. Verenich, I., Dumas, M., La Rosa, M., Maggi, F., Di Francescomarino, C.: Minimizing overprocessing waste in business processes via predictive activity ordering. In: *CAiSE 2016*, Proceedings. pp. 186–202 (2016)
13. Verenich, I., Dumas, M., La Rosa, M., Maggi, F., Teinemaa, I.: Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring. *ACM Transactions on Intelligent Systems and Technology (ACM TIST)* (2019), forthcoming
14. Verenich, I., Dumas, M., La Rosa, M., Nguyen, H.: Predicting process performance: A white-box approach based on process models. *Journal of Software: Evolution and Process* **31**(6), e2170 (2019)