Multimodal Process Prediction (Extended Abstract)

Johannes Lahann[1,2]

[1]Saarland University, Saarbrücken, Germany
[2]German Research Center for Artificial Intelligence (DFKI), Saarbrücken, Germany

Abstract
Process prediction enables the forecast of a process instance based on event log data from process-oriented information systems. Recently, deep learning methods have improved the state-of-the-art in various process prediction tasks. Currently, most research addresses the control flow or time perspective. However, critical information is often hidden in context variables and represented in various data types such as tabular data, image data, and sensor data. This work explores the influence of context variables on process prediction. The main objective is to develop new deep learning-based methods to increase the efficiency and applicability of multimodal process prediction.

1. Introduction

Process prediction, also referred to as predictive business process monitoring or predictive process analytics is a subset of process mining that tries to gain predictive insights into the future of ongoing process instances. Organizations can leverage these predictions to adjust ongoing process execution in real time to prevent undesirable outcomes [1, 2].

In recent years, deep learning-based approaches to process prediction have attained popularity and led to breakthrough results in various tasks such as next step prediction, remaining time, and outcome prediction [3, 4, 5]. Accordingly, numerous deep learning architectures that differ in data processing, modeling, and evaluation have been developed.

Existing work mainly focuses on common event log attributes, especially activities, resources, or timestamps. In practice, however, a few event log attributes rarely determine the future of a process instance alone. Other context variables represented in various data types might also be relevant. These may be tabular data directly associated with the event log, including event attributes such as resource or timestamp or case attributes such as customer or product. In addition, context variables can contain unstructured data such as text, images, and sensor data. These contextual variables are often critical, as they can directly influence the continuation of the process and thus significantly improve the prediction quality.

2. Research Goal and Research Questions

While deep learning techniques have been successfully applied to structured and unstructured data, research on process prediction has mainly involved structured information such as ac-

ICPM 2022 Doctoral Consortium and Tool Demonstration Track
johannes.lahann@dfki.de (J. Lahann)

tivities, resources, or timestamps. To address this research gap, this thesis investigates how context variables of various data types can be processed and combined for process prediction. Designing and implementing deep learning-based methods for multimodal process prediction is the primary goal of this thesis. In particular, the following three research questions shall be addressed.

RQ1. What are the strengths and weaknesses of existing deep learning-based process prediction architectures regarding data processing, modeling, and evaluation?

RQ2. How can different contextual variables be effectively incorporated into multimodal process prediction?

RQ3. How can process prediction enable other predictive process analytics such as trace clustering or anomaly detection?

RQ1 is an empirical knowledge question that involves identifying existing approaches for process prediction and determining their capabilities, limitations, and tradeoffs. In particular, it examines whether there is a superior method or whether specific methods are preferable depending on situational circumstances.

In contrast, RQ2 aims to provide a deeper understanding of the different use cases in which process prediction is used in practice. For this purpose, several use cases from different application domains with different context variables should be examined. Furthermore, the context variables that have the most significant effect on the continuation of the process should be identified. Based on these findings, effective multimodal prediction methods should be developed. These methods should support relevant contextual variables and build on sound design decisions that have been benchmarked against alternative approaches in RQ1.

Finally, RQ3 investigates the potential of neural networks for other predictive process monitoring tasks. In natural language processing, language models have been exploited for solving complex downstream tasks like machine translation, summarization, or question answering. Similarly, with this research question, we want to explore whether neural networks originally trained for process prediction can be advantageously leveraged for anomaly detection, trace clustering, etc.

3. Research Method

The research project follows the design science research (DSR) paradigm for Information Systems (IS) research [6]. Hevner et al. propose seven guidelines (G1-G7) to address a research problem in the IS domain.

G1. *Design as an Artifact.* The result of the research project should be one or more novel artifacts.

G2. *Problem Relevance.* The addressed problem should be relevant to the IS research area.

G3. *Design Evaluation.* The artifacts should be evaluated through an appropriate method to prove their efficiency and utility.

G4. *Research Contribution.* The artifacts should either solve an unsolved problem or improve on the existing solutions for the problem.

G5. *Research Rigor.* The artifacts should be defined through formal definitions, pseudo code, or source code in order to ensure consistency and reproducibility.

G6. *Design as a Search Process.* The artifacts should be designed through an iterative search process. The process consists of defining problem space and solution criteria, designing the solution, and verifying the solution based on the criteria.

G7. *Communication of Research.* The results and outcomes of the research project should be communicated to academics and practitioners.

We plan to implement the above guidelines as follows. To focus on relevant research problems, we performed a structured literature review (SLR) in order to screen existing deep learning-based process prediction methods and reveal trade-offs and limitations (G2). Through the SLR, we identified two major research problems: (i) a lack of a documented and reproducible benchmark for process prediction methods; (ii) a focus on a selection of event log variables leading to a neglect of relevant context information. Accordingly, we set to develop the following novel artifacts (G1, G4): (i) a modular framework for benchmarking novel prediction methods against existing methods. (ii) three new deep learning-based process prediction methods that can effectively incorporate tabular data, image data, and sensor data for multimodal process prediction. We show how these developed predictive methods can be utilized to solve a variety of predictive process monitoring tasks, including next-step prediction, outcome prediction, remaining-time prediction, trace clustering, and anomaly detection. Each artifact is designed formally by justifying all design decisions and reporting pseudo-code or source code (G5). Additionally, we describe the experimental setup of each artifact in detail and release the conducted experiments whenever possible. All artifacts are assessed by empirical evaluations on real-live event logs or synthetic datasets (G3). Finally, all artifacts that we introduce in this thesis are developed through an iterative design flow (G6), while major milestones are presented at conferences and published in journals (G7).

4. Completed Research

To address RQ1, we performed a structured literature review concerning the state-of-the-art of deep learning-based process prediction [7]. The review compares 32 process prediction approaches that are classified along different carefully selected dimensions such as neural network architecture, prediction target, input features, and data processing methods. In particular, the review showed the benefits and drawbacks of the existing approaches. It revealed research gaps that played a significant role in determining the research direction of this dissertation project.

Most of the other author's contributions have addressed RQ2 (e.g. [8, 9, 10]). In [11], we covered a use case from the steel industry where sensor data are processed through an LSTM autoencoder to detect the quality of semi-finished products. Furthermore, we used the learned representations of the autoencoder to predict the next step in the production process. In contrast to most existing prediction approaches, we consider sensor data as an additional input source next to typical event log data.

In [12], we proposed a novel prediction method that can process a flexible number of inputs while supporting categorical and continuous variables by using gramian angular fields and convolutional neural networks. The approach can be adapted through configuration to handle different prediction tasks, including next step prediction, next resource prediction, remaining time prediction, and outcome prediction. Its effectiveness is measured by comparing the prediction quality with existing approaches on publicly available datasets. The results show that the proposed method is an effective alternative to the more frequently used LSTM-based approaches.

In [13], we investigated how deep learning can be applied to computer-aided designs (CAD) to support the planning process in a "one-of-a-kind production" scenario of a manufacturing firm. To this end, we utilized convolutional neural networks to cluster process instances based on their CAD images and connected them with event log data to gather predictive insights.

Only initial research has been conducted concerning RQ3. In [12], we showed that the hidden representations of process prediction models could also be utilized for process trace clustering. Additionally, in [9], we examined the potential of anomaly detection through process prediction for a purchase order handling process. To this end, we investigated the effectiveness of an unsupervised language model pretraining with a target task classifier finetuning applied to process prediction and anomaly detection.

Our other studies examined the strengths and weaknesses of machine learning methods in various use cases [14, 15].

5. Planned Research

Future work comprises activities to analyze the integration of other context attributes of additional data types, such as text data. In addition, thorough benchmarking of design decisions in process prediction research will be conducted, with a focus on comparability and transparency. Finally, further research is needed to deepen the understanding of how an underlying predictive model affects the quality of downstream tasks. To this end, we plan to further explore the extent of a predictive model's impact on anomaly detection.

References

[1] W. Van der Aalst, M. Pesic, M. Song, Beyond process mining: from the past to present and future, in: International Conference on Advanced Information Systems Engineering, Springer, 2010, pp. 38–52.

[2] F. M. Maggi, C. Di Francescomarino, M. Dumas, C. Ghidini, Predictive monitoring of business processes, in: International conference on advanced information systems engineering, Springer, 2014, pp. 457–472.

[3] J. Evermann, J. R. Rehse, P. Fettke, Predicting process behaviour using deep learning, Decis. Support Syst. 100 (2017) 129–140. doi:10.1016/j.dss.2017.04.003.

[4] N. Tax, I. Verenich, M. La Rosa, M. Dumas, Predictive business process monitoring with lstm neural networks, in: Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif.

Intell. Lect. Notes Bioinformatics), volume 10253 LNCS, Springer Verlag, 2017, pp. 477–492. doi:10.1007/978-3-319-59536-8_30.

[5] I. Verenich, M. Dumas, M. L. Rosa, F. M. Maggi, I. Teinemaa, Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring, ACM Transactions on Intelligent Systems and Technology (TIST) 10 (2019) 1–34.

[6] A. R. Hevner, S. T. March, J. Park, S. Ram, Design science in information systems research, MIS quarterly (2004) 75–105.

[7] D. A. Neu, J. Lahann, P. Fettke, A systematic literature review on state-of-the-art deep learning methods for process prediction, Artificial Intelligence Review (2021) 1–27.

[8] S. Dadashnia, P. Fettke, P. Hake, J. Lahann, P. Loos, S. Klein, N. Mehdiyev, T. Niesen, J.-R. Rehse, M. Zapp, Exploring the potentials of artificial intelligence techniques for business process analysis (2017).

[9] O. Gutermuth, J. Lahann, J.-R. Rehse, M. Scheid, S. Schuhmann, S. Stephan, P. Fettke, Efficient and compliant purchase order handling - a contribution to bpi challenge 2019 (2019).

[10] S. Klein, J. Lahann, L. Mayer, D. Neu, P. Pfeiffer, A. Rebmann, M. Scheid, B. Willems, P. Fettke, Business process intelligence challenge 2020: Analysis and evaluation of a travel process, in: 10th Business Process Intelligence Challenge at the Int. Conf. on Process Mining (ICPM), 2020.

[11] N. Mehdiyev, J. Lahann, A. Emrich, D. Enke, P. Fettke, P. Loos, Time series classification using deep learning for process planning: A case from the process industry, Procedia Computer Science 114 (2017) 242–249.

[12] P. Pfeiffer, J. Lahann, P. Fettke, Multivariate business process representation learning utilizing gramian angular fields and convolutional neural networks, in: International Conference on Business Process Management, Springer, Cham, 2021, pp. 327–344.

[13] N. Mehdiyev, L. Mayer, J. Lahann, P. Fettke, Deep learning-based clustering of processes and their visual exploration: An industry 4.0 use case for small, medium-sized enterprises, Expert Systems n/a (????) e13139. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/exsy.13139. doi:https://doi.org/10.1111/exsy.13139. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/exsy.13139.

[14] J. Lahann, M. Scheid, P. Fettke, Utilizing machine learning techniques to reveal vat compliance violations in accounting data, in: 2019 IEEE 21st conference on business informatics (CBI), volume 1, IEEE, 2019, pp. 1–10.

[15] J. Lahann, M. Scheid, P. Fettke, Towards optimal free trade agreement utilization through deep learning techniques (2020).