Towards Process Representation Models for Business Process Management

Peter Pfeiffer[1,2]

[1] German Research Center for Artificial Intelligence (DFKI), Campus D3.2, Saarbrücken, Germany
[2] Saarland Informatics Campus, Saarbrücken, Germany
peter.pfeiffer@dfki.de

Abstract

Language representation models have reached impressive performance in various language understanding tasks, especially due to their ability to learn rich representations of the concepts found in the data. Transferring and adapting these abilities to event logs may have huge potential, as it would enable a more effective and efficient analysis for process mining and business process management (BPM) tasks. However, the development of such models for event logs is challenging due to the specifics of the data and real-world processes. This paper presents the challenges to be faced when developing neural process representation models, recent advances gained in this PhD project as well as fruitful interactions between AI and BPM on this topic.

Introduction and Goal

Analyzing data describing business process executions using process mining methods has gained a lot of attention over the past years. The research field of process mining aims at developing methods that allow to extract and gain knowledge of operational processes from data stored in event logs. In this paper, we introduce the problem of accurately learning the behavior of a process captured in event log data with neural networks and how to use such models to solve tasks in the field of process mining and business process management. In contrast to process discovery techniques, which create a symbolic representation of the sequence of activities found in the event log, we aim at training neural networks to produce a vector representation of process executions, including contextual information, which serve as a basis for solving various tasks. Instead of focusing on methods to solve one specific task, e.g., making predictions about how executions will unfold in the future, we try to develop methods that enable a wider range of applications on event log data. Hence, this research project is about developing neural-network-based approaches for event log data, with the final goal to realize a process representation model with capabilities similar to language models (Devlin et al. 2019) but for the BPM field. While the realization of a universal process representation model is probably out of scope for this PhD, the contributions are supposed to support their development.

Problems and Challenges

In an event log, process executions are captured as sequences of events, so-called cases, where each event has several attributes with different attribute values. The attribute *case id* groups events together, indicating which events belong to the same process execution. We often expect events to have at least the attributes *activity*, *resource*, and a *timestamp*. However, different event logs can have different attributes making them very heterogeneous in terms of attributes and their values - these are varying in number, type, values, and dependencies. For instance, *activity* and *resource* are usually categorical attributes while *timestamp* is continuous. Certain activities are often executed by certain resources, at specific points in time, or in specific contexts. Most event logs feature additional attributes like *costs*, *workload*, textual descriptions, or even CAD data (Mehdiyev et al. 2022) adding important contextual information.

Describing cases as single sequences of homogeneous tokens like words in natural language sentences misses the point that there are multiple attributes of different types. In comparison to multivariate time-series data, the attribute values are typically not determined by the point in time, but by the behavior of the underlying process. For instance, the next activity to execute is determined by what is allowed by the process, contextualized by the past attribute values. While attribute values carry a lot of information, attribute names give additional semantic information. Furthermore, adding the temporal perspective of the event data into predictive models often distorts their predictive quality (Rivera Lazo and Ñanculef 2022) although an information gain should be expected. Thus, processing event log data with existing neural network architectures is not straightforward.

Current neural approaches for event log data use existing network architectures like RNNs and CNNs. Special architectures have been developed for event logs with textual information (Teinemaa et al. 2016) or changing processes (Venkateswaran et al. 2021). However, each existing method has its shortcomings when it comes to including all relevant information that can be found in the data. Furthermore, existing approaches are trained per event log, learning one specific process instead of transferable concepts describing business processes that can be shared across different settings. Learning such features is a key aspect of representation learning which, in the end, enables transfer learning

(Goodfellow, Bengio, and Courville 2016).

The upcoming object-centric event logs, describing process executions as interactions between events and objects (both having different attributes), are even more challenging to process with neural networks as novel features are required to accurately describe the data structure (Adams et al. 2022). Initial studies indicate that applying existing process prediction techniques build for standard event logs suffer from not being able to capture important interactions between events and objects (Galanti et al. 2023).

Given the current state of research, we noted several open challenges to be tackled when developing process representation models:

1. Developing network architectures that allow to process the data modality of an event log appropriately.
2. Combine them with training objectives that allow learning the behavior of the underlying system as well as transferable features to enable transfer learning.
3. Design evaluation methods, standardized tasks as well as datasets and benchmarks to demonstrate and assess the usefulness of process representation models.

Work done

Existing work in this project includes different architectures based on LSTMs, CNNs, or transformers networks - especially the *MPPN* (Pfeiffer, Lahann, and Fettke 2021) which is a CNN-based flexible architecture that allows processing any number of categorical, temporal, and numerical attributes and has shown to be effective for different tasks on case- or event-level such as multivariate process prediction (next step and outcome), case retrieval and multi-perspective event abstraction (Rebmann et al. 2022) with little or no fine-tuning effort after pre-training. Furthermore, we investigated the ambiguity problem (Pfeiffer, Lahann, and Fettke 2022) in next activity prediction, i.e., that running process instances do often not only have one but various valid next steps which negatively affects the evaluation and comparison of process prediction approaches. In (Lahann, Pfeiffer, and Fettke 2022) we investigated whether LSTM can be used to classify cases into fitting and unfitting as well as their performance for case-based anomaly detection on a large number of datasets. Finally, (Pfeiffer 2022) describes the challenges of representation learning on event log data in more detail.

Future Work

Currently, we are working on a review of existing machine-learning-based approaches on event log data to classify them by type of task, features they learn, and datasets being used. The obtained knowledge will support the development of new encoding methods, network architectures, and assessment methods.

Furthermore, we design encoding methods and neural network architectures that allow to process event data while preserving their semantics. For instance, a neural network can hardly distinguish and interpret different categorical attributes if their value encoding does not contain the attribute name. While the *MPPN* is effective in learning features describing process executions or single events, it yet cannot capture the semantics of event attributes.

To assess the capabilities of process representation models to learn the behaviour of the underlying system, we are looking for alternative evaluation methods to next-step prediction accuracy, which has long been used to benchmark neural approaches on event log data but suffers from the ambiguity problem. Instead, we want to explore if a two-step evaluation procedure consisting of intrinsic and extrinsic methods allows obtaining more insights into the capabilities of different architectures and approaches. While intrinsic measures are supposed to test how well the behavior of the process in the data has been learned, extrinsic methods test the performance on different downstream tasks.

In the future, we want to investigate whether self-supervised pre-training or few-shot methods can help to learn transferable features and how they influence the performance. Finally, more applications of process representation models in the process mining and business process management domain, which are challenging and interesting to tackle, are required to test the efficiency and adaptability of these approaches in comparison to existing techniques. For this, we plan to implement a comprehensive benchmark consisting of different tasks.

Interactions between AI and Process Representation Learning

Several open challenges, where interaction between AI and process representation learning might be fruitful, exist - for this PhD project but also beyond. We are very interested in encoding schemes, processing methods, and neural network architectures that enable to process event log data appropriately, i.e., such that as much information and semantics as possible is preserved without loosing accuracy. Similarly, training objectives and strategies to learn process behavior, i.e., concepts and other important features for rich representations as well as evaluation approaches that allow to assess how well the behavior has been learned, are required. For learning transferable features, few-shot learning might be a strategy worth investigating. Object-centric event logs are a very interesting but challenging data structure that benefits from interdisciplinary work.

Process mining sometimes lacks a clear standardization of tasks and representative datasets. This problem has already been solved in several sub-domains of AI like NLP or computer vision. For instance, problems like process prediction and anomaly detection can be framed as classification problems while event abstraction can be seen as a semantic segmentation task. However, we have yet not reached a common understanding for many tasks nor representative datasets with high-quality ground truth labels or optimal performance numbers which hinders comparability.

In the other direction, AI might benefit from new methods on how to process and learn from event log data and how to archive transfer learning on highly variable data modalities. Furthermore, some problems in process mining might be very interesting for the AI community to tackle.

Acknowledgments

This work has been supported by the project *SmartVigilance* (FKZ: 01IS20028C) with financial support by the Federal Ministry of Education and Research (BMBF).

References

Adams, J. N.; Park, G.; Sergej, L.; Schuster, D.; and van der Aalst, W. M. P. 2022. A framework for extracting and encoding features from object-centric event data. In *International Conference on Service-Oriented Computing*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*.

Galanti, R.; de Leoni, M.; Navarin, N.; and Marazzi, A. 2023. Object-centric process predictive analytics. *Expert Systems with Applications*, 213.

Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press.

Lahann, J.; Pfeiffer, P.; and Fettke, P. 2022. LSTM-based Anomaly Detection of Process Instances: Benchmark and Tweaks. In *International Conference on Process Mining Workshops*.

Mehdiyev, N.; Mayer, L.; Lahann, J.; and Fettke, P. 2022. Deep learning-based clustering of processes and their visual exploration: An industry 4.0 use case for small, medium-sized enterprises. *Expert Systems*.

Pfeiffer, P. 2022. Business Process Representation Learning. In Janiesch, C.; Di Francescomarino, C.; Grisold, T.; Kumar, A.; Mendling, J.; Pentland, B.; Reijers, H.; Weske, M.; and Winter, R., eds., *Proceedings of the Best Dissertation Award, Doctoral Consortium, and Demonstration & Resources Track at BPM 2022 co-located with 20th International Conference on Business Process Management (BPM 2022)*.

Pfeiffer, P.; Lahann, J.; and Fettke, P. 2021. Multivariate Business Process Representation Learning Utilizing Gramian Angular Fields and Convolutional Neural Networks. Business Process Management, 327–344. Springer International Publishing.

Pfeiffer, P.; Lahann, J.; and Fettke, P. 2022. The Label Ambiguity Problem in Process Prediction. In *Business Process Management Workshops*.

Rebmann, A.; Pfeiffer, P.; Fettke, P.; and Van der Aa, H. 2022. Multi-perspective Identification of Event Groups for Event Abstraction. In *International Conference on Process Mining Workshops*.

Rivera Lazo, G.; and Ñanculef, R. 2022. Multi-attribute Transformers for Sequence Prediction in Business Process Management. Discovery Science, 184–194. Springer Nature Switzerland.

Teinemaa, I.; Dumas, M.; Maggi, F. M.; and Francescomarino, C. D. 2016. Predictive business process monitoring with structured and unstructured data. In *International Conference on Business Process Management*, 401–417. Springer.

Venkateswaran, P.; Muthusamy, V.; Isahagian, V.; and Venkatasubramanian, N. 2021. Robust and Generalizable Predictive Models for Business Processes. In *International Conference on Business Process Management*, 105–122. Springer.