

# Performance-Preserving Event Log Sampling for Predictive Monitoring

Mohammadreza Fani Sani<sup>1\*</sup>, Mozhgan  
Vazifehdoostirani<sup>2</sup>, Gyunam Park<sup>1</sup>, Marco  
Pegoraro<sup>1</sup>, Sebastiaan J. van Zelst<sup>1,3</sup> and Wil M.P. van der  
Aalst<sup>1,3</sup>

<sup>1</sup>Process and Data Science Chair, RWTH Aachen University,  
Aachen, Germany.

<sup>2</sup>Industrial Engineering and Innovation Science, Eindhoven  
University of Technology, Eindhoven, the Netherlands.

<sup>3</sup>Fraunhofer FIT, Birlinghoven Castle, Sankt Augustin, Germany.

\*Corresponding author(s). E-mail(s):

[fanisani@pads.rwth-aachen.de](mailto:fanisani@pads.rwth-aachen.de);

Contributing authors: [m.vazifehdoostirani@tue.nl](mailto:m.vazifehdoostirani@tue.nl);

[Park@pads.rwth-aachen.de](mailto:Park@pads.rwth-aachen.de); [Pegoraro@pads.rwth-aachen.de](mailto:Pegoraro@pads.rwth-aachen.de);  
[s.j.v.zelst@pads.rwth-aachen.de](mailto:s.j.v.zelst@pads.rwth-aachen.de); [wvdaalst@pads.rwth-aachen.de](mailto:wvdaalst@pads.rwth-aachen.de);

## Abstract

Predictive process monitoring is a subfield of process mining that aims to estimate case or event features for running process instances. Such predictions are of significant interest to the process stakeholders. However, most of the state-of-the-art methods for predictive monitoring require the training of complex machine learning models, which is often inefficient. Moreover, most of these methods require a hyper-parameter optimization that requires several repetitions of the training process which is not feasible in many real-life applications. In this paper, we propose an instance selection procedure that allows sampling training process instances for prediction models. We show that our instance selection procedure allows for a significant increase of training speed for next activity and remaining time prediction methods while maintaining reliable levels of prediction accuracy.

**Keywords:** Process Mining, Predictive Monitoring, Sampling, Machine Learning, Deep Learning, Instance Selection

## 1 Introduction

The main goal of predictive process monitoring is to provide timely information by predicting the behavior of business processes [1] and enabling proactive actions to improve the performance of the process. It provide various predictive information such as the next performing activity of a process instance [2], e.g., patient and product, its waiting time for an activity, its remaining time to complete the process, etc [3]. For instance, by predicting the long waiting time of a patient for a registration, one can bypass the activity or add more resources to perform it.

A plethora of approaches have been proposed to support predictive process monitoring. In particular, with the recent breakthroughs in machines learning, various machine learning-based approaches have been developed [3]. The emergence of ensemble learning methods leads to improvement in accuracy in different areas [4]. eXtreme Gradient Boosting (XGBoost) [5] has shown promising results, often outperforming other ensemble methods such as Random Forest or using a single regression tree [6, 7]. Furthermore, techniques based on deep neural networks, e.g., Long-Short Term Memory (LSTM) networks, have shown high performance in different tasks [8].

However, machine learning-based techniques are computationally expensive due to their training process [9]. Moreover, they often require exhaustive hyperparameter-tuning to provide acceptable accuracy. Such limitations hinder the application of machine learning-based techniques into real-life business processes where new prediction models are required in short intervals to adapt to changing business situations. Business analysts need to test the efficiency and reliability of their conclusions via repeated training of different prediction models with different parameters [3]. Such long training time limits the application of the techniques considering the limitations in time and hardware [10].

In this regard, instance selection has been studied as a promising direction of research to reduce original datasets to a manageable volume to perform machine learning tasks, while the quality of the results (e.g., accuracy) is maintained as if the original dataset was used [11]. Instance selection techniques are categorized into two classes based on the way they select instances. First, some techniques select the instances at the boundaries of classes. For instance, Incremental Reduction Optimization Procedure (DROP) [12] selects instances using  $k$ -Nearest Neighbors by incrementally discarding an instance if its neighbors are correctly classified without the instance. The other techniques preserve the instances residing inside classes, e.g., Edited Nearest Neighbor (ENN) [13] preserves instances by repeatedly discarding an instance if it does not belong to the class of the majority of its neighbors.

However, it is limited to directly apply existing techniques for instance selection to predictive process monitoring training, since such techniques assume independence among instances [12]. Instead, in predictive process monitoring, instances are computed from event data that are recorded by the information system supporting business processes [14]. Thus, they are highly correlated [15] with the notion of *case*, e.g., a manufacturing product in a factory and a customer in a store.

In this work, we suggest an instance selection approach for predicting the next activity and the remaining time of the process that are two of the main applications of predictive business process monitoring. By considering the characteristics of the event data, the proposed approach samples event data such that the training speed is improved while the accuracy of the resulting prediction model is maintained. We have evaluated the proposed methods using three real-life datasets and state-of-the-art techniques for predictive business process monitoring, including LSTM [16] and XGBoost [5].

This paper extends our earlier work presented in [17] in the following dimensions. We have extended our method to be applied to multiple prediction tasks including next activity and remaining time prediction [1]. For the next activity prediction, we have conducted additional experiments using other real-life event logs and sampling settings. Moreover, we have evaluated the proposed approach for the remaining time prediction. To this end, we have conducted experiments using three real-life event logs. In addition, the implementation of the sampling now supports both Java and Python platforms, expanding the accessibility of the proposed approach. Furthermore, we have extended the related work section to ... Finally, we have included a discussion section to discuss the implications and limitations of this work.

The remainder is organized as follows. We discuss the related work in [section 2](#). Next, we present the preliminaries in [section 3](#) and proposed methods in [section 4](#). Afterward, [section 5](#) evaluates the proposed methods using real-life event data and [section 6](#) provides discussions. Finally, [section 7](#) concludes the paper.

## 2 Related Work

Predictive process monitoring is an exceedingly active field, both currently and historically, thanks to the compatibility of process sciences and other branches of data science that include inference techniques, such as statistics and machine learning. At its core, the fundamental component of many predictive monitoring approaches is the abstraction technique it uses to obtain a fixed-length representation of the process component subject to the prediction (often, but not always, process traces). In the earlier approaches, the need for such abstraction was overcome through model-aware techniques, employing process models and replay techniques on partial traces to abstract a flat representation of event sequences. Such process models are mostly automatically discovered from a set of available complete traces, and require perfect

fitness on training instances (and, seldomly, also on unseen test instances). For instance, van der Aalst et al. [1] proposed a time prediction framework based on replaying partial traces on a transition system, effectively clustering training instances by control-flow information. This framework has later been the basis for a prediction method by Polato et al. [18], where the transition system is annotated with an ensemble of SVR and Naïve Bayes classifiers, to perform a more accurate time estimation. Some more recent approaches split the predictive contribution of process models and machine learning models in a perspective-wise manner: for instance, Park et al. [19] obtain a representation of the performance perspective using an annotated transition system, and design an ensemble with a deep neural network to obtain the final predictive model. A related approach, albeit more linked to the simulation domain and based on a Monte Carlo method, is the one proposed by Rogge-Solti and Weske [20], which maps partial process instances in an enriched Petri net.

Recently, predictive process monitoring started to use a plethora of machine learning approaches, achieving varying degrees of success. For instance, Teinmaa et al. [21] provided a framework to combine text mining methods with Random Forest and Logistic Regression. Senderovich et al. [6] studied the effect of using intra-case and inter-case features in predictive process monitoring and showed a promising result for XGBoost compared to other ensemble and linear methods. A comprehensive benchmark on using classical machine learning approaches for outcome-oriented predictive process monitoring tasks [7] has shown that the XGBoost is the best-performing classifier among different machine learning approaches such as SVM, Decision Tree, Random Forest, and logistic regression.

More recent methods are model-unaware and perform based on a single and more complex machine learning model instead of an ensemble. In fact, such an evolution of predictive monitoring mimics the advancement in the accuracy of newer machine learning approaches, specifically the numerous and sophisticated models based on deep neural networks that have been developed in the last decade. The LSTM network model has proven to be particularly effective for predictive monitoring [8, 22], since the recurrent architecture can natively support sequences of data of arbitrary length. It allows performing trace prediction while employing a fixed-length event abstraction, which can be based on control-flow alone [8, 22], data-aware [23], time-aware [24], text-aware [25], or model-aware [19].

The latest research developments in the field of predictive monitoring, much like both this paper and the application of machine learning techniques in other domains, shifts its focus away from increasing the quality of prediction (with respect to a given error metric or a given benchmark), and specializes in enriching the results of the prediction on additional aspects or perspectives. Unlike the methods mentioned in the previous paragraph, this is usually obtained through dedicated machine learning models—or modifications thereof—rather than designing specific event- or trace-level abstractions. For instance, many scholars have attempted to make predictions more transparent through the

use of explainable machine learning techniques [26–29]. More related to our present work, Pauwels and Calders [30] propose a technique to avoid the time expenditure caused by the retrain of machine learning models; this is necessary when they are not representative anymore—for instance, when changes occur in the underlying process (caused e.g. by concept drift). While in this paper we focus on the data, and we propose a solution based on sampling, Pauwels and Calders intervene on the model side, devising an incremental training schema which accounts for new information in an efficient way.

Another concept similar to the idea proposed in this paper, and of current interest in the field of machine learning, is *dataset distillation*: utilizing a dataset to obtain a smaller set of training instances that contain the same information (with respect to training a machine learning model) [31]. While this is not considered sampling, since some instances of the distilled dataset are created ex-novo, it is an approach very similar to the one we illustrate in our paper. Moreover, recently some instance selection algorithms have been proposed to help process mining algorithms. For example, [32, 33] proposed to use instance selection techniques to improve the performance of process discovery and conformance checking procedures.

In this paper, we examine the underexplored topic of event data sampling and selection for predictive process monitoring, with the objective of assessing if and to which extent prediction quality can be retained when we utilize subsets of the training data.

### 3 Preliminaries

In this section, some process mining concepts such as event log and sampling are discussed. In process mining, we use events to provide insights into the execution of business processes. Event logs, i.e., collections of events representing the execution of several instances of a process, are the starting point of process mining algorithms. An example event log is shown in Table 1. Each *event* that relates to a row in the table is related to specific activities of the underlying process. Furthermore, we refer to a collection of events related to a specific process instance of the process as a *case* (represented by the *Case-id* column). Both cases and events may have different attributes. An event log that is a collection of events and cases is defined as follows.

**Definition 1** (Event Log) Let  $\mathcal{E}$  be the universe of events,  $\mathcal{C}$  be the universe of cases,  $\mathcal{AT}$  be the universe of attributes, and  $\mathcal{U}$  be the universe of attribute values. Moreover, let  $C \subseteq \mathcal{C}$  be a non-empty set of cases, let  $E \subseteq \mathcal{E}$  be a non-empty set of events, and let  $AT \subseteq \mathcal{AT}$  be a set of attributes. We define  $(C, E, \pi_C, \pi_E)$  as an event log, where  $\pi_C: C \times \mathcal{AT} \rightarrow \mathcal{U}$  and  $\pi_E: E \times \mathcal{AT} \rightarrow \mathcal{U}$ . Any event in the event log has a case, therefore,  $\nexists_{e \in E} (\pi_E(e, \text{case}) \notin C)$  and  $\bigcup_{e \in E} (\pi_E(e, \text{case})) = C$ .

Furthermore, let  $\mathcal{A} \subseteq \mathcal{U}$  be the universe of activities and let  $\mathcal{V} \subseteq \mathcal{A}^*$  be the universe of sequences of activities. For any  $e \in E$ , function  $\pi_E(e, \text{activity}) \in \mathcal{A}$ , which means

## 6 Performance-Preserving Event Log Sampling for Predictive Monitoring

**Table 1** Simple example of an event log. Rows capture *events* recorded in the context of the execution of the process. An event describes at what point in time an activity was performed. Other data attributes may be available as well.

| Case-id | Event-id | Activity name     | Starting time    | Finishing time ...   |
|---------|----------|-------------------|------------------|----------------------|
| ⋮       | ⋮        | ⋮                 | ⋮                | ⋮                    |
| 7       | 35       | Register(a)       | 2021-01-02 12:23 | 2021-01-02 12:25 ... |
| 7       | 36       | Analyze Defect(b) | 2021-01-02 12:30 | 2021-01-02 12:40 ... |
| 7       | 37       | Inform User(g)    | 2021-01-02 12:45 | 2021-01-02 12:47 ... |
| 8       | 39       | Register(a)       | 2021-01-02 12:23 | 2021-01-02 13:15 ... |
| 7       | 40       | Test Repair(e)    | 2021-01-02 13:05 | 2021-01-02 13:20 ... |
| 7       | 41       | Archive Repair(h) | 2021-01-02 13:21 | 2021-01-02 13:22 ... |
| ⋮       | ⋮        | ⋮                 | ⋮                | ⋮                    |

that any event in the event log has an activity. Moreover, for any  $c \in C$  function  $\pi_C(c, \text{variant}) \in \mathcal{A}^* \setminus \{\langle \rangle\}$  that means any case in the event log has a variant.

Therefore, there are some mandatory attributes that are *case* and *activity* for events and *variants* for cases. Variant information plays an important role an in some process mining applications, e.g., process discovery and conformance checking, just this information is considered. In this regard, event logs are considered as a multiset of sequences of activities. In the following, a simple event log is defined.

**Definition 2** (Simple event log) Let  $\mathcal{A}$  be the universe of activities and let the universe of multisets over a set  $X$  be denoted by  $\mathcal{B}(X)$ . A simple event log is  $L \in \mathcal{B}(\mathcal{A}^*)$ . Moreover, let  $\mathcal{EL}$  be the universe of event logs and  $EL = (C, E, \pi_C, \pi_E) \in \mathcal{EL}$  be an event log. We define function  $sl: \mathcal{EL} \rightarrow \mathcal{B}(\{\pi_E(e, \text{activity}) : e \in E\}^*)$  returns the simple event log of an event log. The set of unique variants in the event log is denoted by  $sl(EL)$ .

Therefore,  $sl$  returns the multiset of variants in the event logs. Note that the size of a simple event log equals the number of cases in the event logs, i.e.,  $sl(EL) = C$

In this paper, we use sampling techniques to reduce the size of event logs. An event log sampling method is defined as follows.

**Definition 3** (Event log sampling) Let  $\mathcal{EL}$  be the universe of event logs and  $\mathcal{A}$  be the universe of activities. Moreover, let  $EL = (C, E, \pi_C, \pi_E) \in \mathcal{EL}$  be an event log, we define function  $\delta: \mathcal{EL} \rightarrow \mathcal{EL}$  that returns the sampled event log where if  $(C', E', \pi'_C, \pi'_E) = \delta(EL)$ , then  $C' \subseteq C$ ,  $E' \subseteq E$ ,  $\pi'_E \subseteq \pi_E$ ,  $\pi'_C \subseteq \pi_C$ , and consequently,  $sl(\delta(EL)) \subseteq sl(EL)$ . We define that  $\delta$  is a variant-preserving sampling if  $sl(\delta(EL)) = sl(EL)$ .

In other words, a sampling method is variant-preserving if and only if all the variants of the original event log are presented in the sampled event log.




**Fig. 1** A schematic view of the proposed sampling procedure

To use machine learning methods for prediction, we usually need to transfer each case to one or more features. The feature is defined as follows.

**Definition 4** (Feature) Let  $\mathcal{AT}$  be the universe of attributes,  $\mathcal{U}$  be the universe of attribute values, and  $\mathcal{C}$  be the universe of cases. Moreover, let  $AT \subseteq \mathcal{AT}$  be a set of attributes. A feature is a relation between a sequence of attributes' values for  $AT$  and the target attribute value, i.e.,  $f \in (\mathcal{U}^{AT} \times \mathcal{U})$ . We define  $fe: \mathcal{C} \times \mathcal{EL} \rightarrow \mathcal{B}(\mathcal{U}^{AT} \times \mathcal{U})$  is a function that receives a case and an event log, and returns a multiset of features.

For the next and final activity prediction, the target attribute value should be an activity. However, for the remaining time prediction, the target attribute value is a numerical value. Moreover, a case in the event log may have different features. For example, suppose that we only consider the activities. For the case  $\langle a, b, c, d \rangle$ , we may have  $(\langle a \rangle, b)$ ,  $(\langle a, b \rangle, c)$ , and  $(\langle a, b, c \rangle, d)$  as features. Furthermore,  $\sum_{c \in \mathcal{C}} fe(c, EL)$  are the corresponding features of event log  $EL = (C, E, \pi_C, \pi_E)$  that could be given to different machine learning algorithms. For more details on how to extract features from event logs please refer to [34].

## 4 Proposed Sampling Methods

In this section, we propose an event log preprocessing procedure that helps prediction algorithms to perform faster while maintaining reasonable accuracy. The schematic view of the proposed instance selection approach is presented in Figure 1. First, we need to traverse the event log and find the variants and corresponding traces of each variant in the event log. Moreover, different distributions of data attributes in each variant will be computed. Afterward, using different sorting and instance selection strategies, we are able to select some of the cases and return the sample event log. In the following, each of these steps is explained in more detail.

1. *Traversing the event log*: In this step, the unique variants of the event log and the corresponding traces of each variant are determined. In other words, consider event log  $EL$  that  $\overline{sl}(EL) = \{\sigma_1, \dots, \sigma_n\}$  where  $n = \|\overline{sl}(EL)\|$ , we aim to split  $EL$  to  $EL_1, \dots, EL_n$  where  $EL_i$  only contains all the cases that  $C_i = \{c \in \mathcal{C} : \pi_C(c, variant) = \sigma_i\}$  and  $E_i = \{e \in E : \pi_E(e, case) \in C_i\}$ . Obviously,  $\bigcup_{1 \leq i \leq n} (C_i) = \mathcal{C}$  and  $\bigcap_{1 \leq i \leq n} (C_i) = \emptyset$ .

8 *Performance-Preserving Event Log Sampling for Predictive Monitoring*

2. *Distribution Computation*: In this step, for each variant of the event log, we compute the distribution of different data attributes  $a \in AT$ . It would be more practical if the interesting attributes are chosen by an expert. Both event and case attributes can be considered. A simple approach is to compute the frequency of categorical data values. For numerical data attributes, it is possible to consider the average or the median of values for all cases of each variant.
3. *Sorting the cases of each variant*: In this step, we aim to sort the traces of each variant. We need to sort the traces to give a higher priority to those traces that can represent the variant better. One way is to sort the traces based on the frequency of the existence of the most occurred data values of the variant. For example, we can give a higher priority to the traces that have more frequent resources of each variant. It is also possible to sort the traces based on their arrival time or randomly.
4. *Returning sample event logs*: Finally, depending on the setting of the sampling function, we return some of the traces with the highest priority for all variants. The most important point about this step is to know how many traces of each variant should be selected. In the following, some possibilities will be introduced.
  - *Unique selection*: In this approach, we select only one trace with the highest priority. In other words, suppose that  $L' = sl(\delta(EL))$ ,  $\forall_{\sigma \in L'} L'(\sigma) = 1$ . Therefore, using this approach we will have  $\|sl(\delta(EL))\| = \|sl(EL)\|$ . It is expected that using this approach, the distribution of frequency of variants will be changed and consequently the resulted prediction model will be less accurate.
  - *Logarithmic distribution*: In this approach, we reduce the number of traces in each variant in a logarithmic way. If  $L = sl(EL)$  and  $L' = sl(\delta(EL))$ ,  $\forall_{\sigma \in L'} L'(\sigma) = \lceil \log_k(L(\sigma)) \rceil$ . Using this approach, the infrequent variants will not have any trace in the sampled event log and consequently it is not variant-preserving. According the above formula, by using a higher base for the logarithm (i.e.,  $k$ ), the size of the sampled event log is reduced more.
  - *Division*: This approach performs similar to the previous one, however, instead of using logarithmic scale, we apply the division operator. In this approach,  $\forall_{\sigma \in L'} L'(\sigma) = \lceil \frac{L(\sigma)}{k} \rceil$ . A higher  $k$  results in fewer cases in the sample event log. Note that as  $\lceil \cdot \rceil$  considered in the above formula, using this approach all the variants have at least one trace in the sampled event log and it is variant-preserving.

There is also a possibility to consider other selection methods. For example, we can select the traces completely randomly from the original event log.

By choosing different data attributes in Step 2 and different sorting algorithms in Step 3, we are able to lead the sampling of the method on *which* cases should be chosen. Moreover, by choosing the type of distribution in Step 4, we determine *how many* cases should be chosen. To compute how sampling

method  $\delta$  reduces the size of the given event log  $EL$ , we use the following equation:

$$R_S = \frac{\|sl(EL)\|}{\|sl(\delta(EL))\|} \quad (1)$$

The higher  $R_S$  value means, the sampling method reduces more the size of the training log. By choosing different distribution methods and different  $k$ -values, we are able to control the size of the sampled event log. It should be noted that the proposed method will apply just to the training event log. In other words, we do not sample event logs for development and test datasets.

## 5 Evaluation

In this section, we aim at designing some experiments to answer the research question, i.e., "Is it possible to have computational performance improvement of prediction methods by using the sampled event logs, while maintaining a similar accuracy?". It should be noted that the focus of the experiments is not on prediction model tuning to have higher accuracy. Conversely, we aim to analyze the effect of using sampled event logs (instead of the whole datasets) on the required time and the accuracy of prediction models.

In the following, we first explain evaluation settings and event logs that are used. Afterward, we provide some information about the implementation of sampling methods, and finally, we show the experimental results.

### 5.1 Evaluation Setting

Here we first explain the prediction methods and parameters that are used in the evaluation. Afterward, the used metrics and event logs are discussed.

#### 5.1.1 Evaluation Parameters

We have used to machine learning methods to train the prediction models, i.e., *LSTM* and *XGBoost*. For predicting the next activity, our LSTM network consisted of an input layer, two LSTM layers with *dropout* rates of 10%, and a dense output layer with the *SoftMax* activation function. We used "categorical cross-entropy" to calculate the loss and adopted *ADAM* as an optimizer. we built the same architecture of *LSTM* for the remaining time predicting with some differences. We employed "mean absolute error" as a loss function and "Root Mean Squared Propagation" (*RMSprop*) as an optimizer. We used *gbtree* with a max depth of 6 as a booster in our XGBoost model for both of the next activity and remaining time prediction tasks. Uniform distribution is used as the sampling method inside our XGBoost model. To avoid overfitting in both models, the training set is further divided into 90% training set and 10% validation set to stop training once the model performance on the validation set stops improving. We used the same parameter setting of both models for original event logs and sampled event logs. To access our implementations of these methods and the feature generation please refer to <https://github.com>.

[com/gyunamister/pm-prediction/](https://com/gyunamister/pm-prediction/). It should be noted that for the remaining time prediction part of the experiments the settings and codes of the survey presented in [35] are used and just the sample of event logs are considered for training the methods. For details of the feature generation and feature encoding steps, for the next activity please refer to [19], and for the remaining time prediction refer to [35].

To sample the event logs, we use three distributions that are *log distribution*, *division*, and *unique variants*. For the *log* distribution method, we have used 2, 3, 5, and 10 (i.e.,  $\log_2, \log_3, \log_5$ , and  $\log_{10}$ ). For the division method, we have used 2, 3, 5, and 10 (i.e.,  $d2, d3, d5$ , and  $d10$ ). For each event log and each sampling method, we have used a 5-fold cross-validation. Moreover, as the results of the experiments are non-deterministic, all the experiments have been repeated 5 times and the average values are represented. Moreover, to have a fair evaluation, in all the steps, one CPU thread has been used.

### 5.1.2 Metrics

To evaluate the correctness of prediction methods for predicting next activities, we have considered two metrics, i.e., *Accuracy* and *F1-score*. The F1-score is used because in some cases we have imbalanced data [36]. For remaining time prediction, we consider Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) measures as they were used in [35] that are computed as follows.

$$MAE = \frac{1}{n} \sum_{t=1}^n e_t \quad (2)$$

$$RMSE = \sqrt{\frac{1}{n} \sum_{t=1}^n e_t^2} \quad (3)$$

In the above equations,  $e_t$  indicates the prediction error for the  $t^{th}$  instance of validation data. In other words, we aim to compute the absolute difference between the predicted and real values (in days) for the remained time. For both measures, the lower value means higher accuracy. Note that, similar to [35], we considered seconds as the time unit to compute these two metrics.

To evaluate how accurate the prediction methods using the sampled event logs, we have used relative metrics that compared them with the case that whole event logs are used according to the following equations.

$$R_{Acc} = \frac{\text{Accuracy using the sampled training log}}{\text{Accuracy using the whole training log}} \quad (4)$$

$$R_{F1} = \frac{\text{F1-score using the sampled training log}}{\text{F1-score using the whole training log}} \quad (5)$$

In both above equations a value close to 1 means that using the sampling event logs, the prediction methods behave almost similar to the case that the

**Table 2** Event logs that are used in the evaluation and results of using them for the next activity and remaining time prediction

| Log        | Traces | Variants | Activities | Feature Extraction Time | Next Activity Prediction |      |          |      |               |       | Remaining Time Prediction |       |         |        |               |     |
|------------|--------|----------|------------|-------------------------|--------------------------|------|----------|------|---------------|-------|---------------------------|-------|---------|--------|---------------|-----|
|            |        |          |            |                         | Accuracy                 |      | F1_Score |      | Training Time |       | MAE                       |       | RMSE    |        | Training Time |     |
|            |        |          |            |                         | LSTM                     | XGB  | LSTM     | XGB  | LSTM          | XGB   | LSTM                      | XGB   | LSTM    | XGB    | LSTM          | XGB |
| BPIC2012-W | 9658   | 2643     | 6          | 400.2475                | .913                     | .698 | .646     | .641 | 690.5         | 28.8  | 11.111                    | 7.467 | 16.214  | 10.380 | 64            | 29  |
| RTFM       | 150370 | 231      | 11         | 25231.13                | .967                     | .849 | .741     | .763 | 607.6         | 178.9 | 315.811                   | 2.787 | 455.563 | 8.558  | 352           | 379 |
| Sepsis     | 1050   | 846      | 16         | 141.9163                | .957                     | .640 | .605     | .614 | 1628.5        | 310.5 | 0.014                     | 0.003 | 0.042   | 0.011  | 7             | 3   |

whole data is used for the training. Moreover, values higher than 1 indicate the accuracy/F1-score of prediction methods has improved.

Unlike previous metrics, for *MAE* and *RMSE*, a higher value means the prediction model is less accurate in predicting the remaining time. Therefore, we use the following measures.

$$R_{MAE} = \frac{MAE \text{ using the whole training log}}{MAE \text{ using the sampled training log}} \quad (6)$$

$$R_{RMSE} = \frac{RMSE \text{ using the whole training log}}{RMSE \text{ using the sampled training log}} \quad (7)$$

In both of the above measures, a higher value means that the applied instance selection method preserves higher accuracy. In case the values of the above measures are higher than 1, the instance selection methods improve the accuracy of prediction models compared to the case that the whole training data have been used.

To compute the improvement in the performance of feature extraction and training time, we will use the following equations.

$$R_{FE} = \frac{\text{Feature extraction time using whole data}}{\text{Feature extraction time using the sampled data}} \quad (8)$$

$$R_t = \frac{\text{Training time using whole data}}{\text{Training time using the sampled data}} \quad (9)$$

For both equations, the resulting values indicate how many times applying the sampled log is faster than using all data.

### 5.1.3 Event Logs

We have used three event logs widely used in the literature. In the *BPIC-2012-W* event log, relating to a process of an insurance company, the average of variant frequencies is low. In the *RTFM* event log, which corresponds to a road traffic management system, we have some high frequent variants and several infrequent variants. Moreover, the number of activities in this event log is high. In the *Sepsis* event log, relating to a health care process, there are several variants that most of them are unique. Some of the activities in the last two event logs are infrequent, which makes these event logs imbalanced. Some information about these event logs and the result of using prediction methods on them is presented in [Table 2](#).

**Table 3** Reduction in size of training event logs (i.e.,  $R_S$ ) and the improvement in the feature extraction process (i.e.,  $R_{FE}$ ) using different sampling methods

| Log        | $R_S$ |      |      |      |        |        |        |         |        | $R_{FE}$ |      |       |       |         |          |          |          |          |
|------------|-------|------|------|------|--------|--------|--------|---------|--------|----------|------|-------|-------|---------|----------|----------|----------|----------|
|            | d2    | d3   | d5   | d10  | log2   | log3   | log5   | log10   | Unique | d2       | d3   | d5    | d10   | log2    | log3     | log5     | log10    | Unique   |
| BPIC2012-W | 1.50  | 1.93 | 2.42 | 2.92 | 6.67   | 10.13  | 18.31  | 28.62   | 3.48   | 1.39     | 2.20 | 2.66  | 2.64  | 13.25   | 35.03    | 60.26    | 146.88   | 2.03     |
| RTFM       | 2.00  | 2.99 | 4.97 | 9.87 | 273.40 | 429.63 | 653.78 | 1002.47 | 546.80 | 3.97     | 6.79 | 15.57 | 55.63 | 8569.65 | 17053.64 | 33072.59 | 55915.15 | 18886.04 |
| Sepsis     | 1.07  | 1.13 | 1.18 | 1.19 | 11.20  | 15.85  | 35.00  | 70.00   | 1.20   | 1.17     | 1.34 | 1.39  | 1.39  | 11.37   | 23.42    | 66.08    | 128.77   | 1.40     |

According to [Table 2](#), using the whole event data we usually have high accuracy for the next activity prediction. However, the F1-score is not that high which is mainly because the event logs are imbalanced (specifically *RTFM* and *Sepsis*). Moreover, the MAE and RMSE values are very high. Specifically for the *RTFM* event log. It is mainly because process instances' durations are very long in this event log and consequently the  $e_t$  values are higher. Finally, there is a direct relation between the size of event logs and the required time for extracting features and training the models.

## 5.2 Evaluation Results

Here, we provide the results of using sampled training event logs instead of whole training event logs. First we show how by using sampling the size of training data is reduced in [Table 3](#). As it is expected the highest reduction is occurred when  $\log_{10}$  is used. Using this sampling, the size of the *RTFM* event log is more than 1000 times reduced. However, for the *Sepsis* event log as most variants are unique, sampling the training event logs using divide distribution could not result in high  $R_S$ . Moreover, this table shows how using the sampling event logs can reduce the required time to extract features of the event data, i.e.,  $R_{FE}$ . As it is expected, there is a correlation between the size reduction of the sampled event logs and the improvement in the  $R_{FE}$ .

In the following, we show that how using the sampling event logs affects on the next activity and remaining time prediction.

### 5.2.1 Next Activity Prediction

The accuracy of both LSTM and XGboost methods that are trained using sampled training data is presented in [Table 4](#). Results indicate that in most cases when the division sampling methods are used, we can achieve a similar accuracy, i.e.,  $R_{Acc}$  close to 1, compared to the case that the whole training data is used. In some cases like using  $d2$  for the *Sepsis* event log and LSTM method, the accuracy of the training method is even improved. However, for the *RTFM* using sampling with logarithmic or unique distribution highly change the frequency distribution of variants and consequently cause to have higher reduction in the accuracy of predicting models. Moreover, the accuracy reduction is higher when the were used.

In [Table 5](#),  $R_{F1}$  of trained models are depicted. The results again indicate that using the sampling method with divide distribution in most cases lead to having similar (and some times higher) F1-score.

**Table 4**  $R_{Acc}$  of different event logs when different sampling methods are used.

| Log        | LSTM  |       |       |       |      |      |      |       |        | XGB   |       |       |       |      |      |      |       |        |
|------------|-------|-------|-------|-------|------|------|------|-------|--------|-------|-------|-------|-------|------|------|------|-------|--------|
|            | d2    | d3    | d5    | d10   | log2 | log3 | log5 | log10 | Unique | d2    | d3    | d5    | d10   | log2 | log3 | log5 | log10 | Unique |
| BPIC2012-W | 1.001 | .997  | 1.001 | 1.000 | .992 | .994 | .990 | .970  | 1.000  | 1.000 | 1.000 | .998  | .998  | .932 | .929 | .891 | .887  | .999   |
| RTFM       | 1.000 | .999  | .999  | .999  | .955 | .955 | .964 | .965  | .943   | 1.000 | 1.000 | 1.000 | 1.000 | .682 | .687 | .691 | .775  | .580   |
| Sepsis     | 1.001 | 1.001 | 1.001 | 1.001 | .982 | .980 | .959 | .948  | 1.001  | 1.000 | .996  | .995  | .993  | .827 | .844 | .704 | .633  | .996   |

**Table 5**  $R_{F1}$  of different event logs when different sampling methods are used.

| Log        | LSTM  |       |       |       |      |      |      |       |        | XGB   |       |       |       |      |      |      |       |        |
|------------|-------|-------|-------|-------|------|------|------|-------|--------|-------|-------|-------|-------|------|------|------|-------|--------|
|            | d2    | d3    | d5    | d10   | log2 | log3 | log5 | log10 | Unique | d2    | d3    | d5    | d10   | log2 | log3 | log5 | log10 | Unique |
| BPIC2012-W | .976  | .969  | .972  | .973  | .958 | .962 | .943 | .854  | .973   | .988  | .985  | .976  | .974  | .950 | .948 | .913 | .908  | .975   |
| RTFM       | 1.002 | 1.004 | .999  | .989  | .724 | .730 | .784 | .786  | .637   | 1.000 | 1.000 | 1.000 | 1.000 | .679 | .689 | .689 | .772  | .582   |
| Sepsis     | 1.009 | 1.002 | 1.004 | 1.010 | .794 | .751 | .439 | .273  | 1.009  | .999  | .994  | .992  | .991  | .839 | .848 | .710 | .628  | .993   |

**Table 6**  $R_t$  of different event logs when different sampling methods are used.

| Log        | LSTM |     |     |      |      |      |      |       |        | XGB |     |     |      |      |       |       |       |        |
|------------|------|-----|-----|------|------|------|------|-------|--------|-----|-----|-----|------|------|-------|-------|-------|--------|
|            | d2   | d3  | d5  | d10  | log2 | log3 | log5 | log10 | Unique | d2  | d3  | d5  | d10  | log2 | log3  | log5  | log10 | Unique |
| BPIC2012-W | 1.3  | 1.8 | 2.1 | 2.6  | 4.8  | 4.5  | 13.3 | 13.9  | 2.8    | 1.9 | 2.0 | 2.4 | 2.2  | 2.8  | 4.7   | 5.6   | 6.3   | 3.0    |
| RTFM       | 2.2  | 3.2 | 5.4 | 10.1 | 33.5 | 39.8 | 45.6 | 49.0  | 35.5   | 2.3 | 4.2 | 6.6 | 12.9 | 84.3 | 197.4 | 176.4 | 238.8 | 205.5  |
| Sepsis     | 2.3  | 2.8 | 2.8 | 2.9  | 17.9 | 26.5 | 77.6 | 116.0 | 2.9    | 3.6 | 4.4 | 4.1 | 4.6  | 22.8 | 90.0  | 143.2 | 239.8 | 4.5    |

By considering the results of these two tables, we found that specifically when the LSTM method is used we will have similar accuracy and F1-score. Moreover, for the *Sepsis* and *BPIC2012-W* that variants have similar frequency having all the variants (i.e., using divide and unique distributions) can help the prediction method to have results similar to the case that we use the whole training data. However, for the *RTFM* event log that has some high frequent variants, using the unique distribution results in lower accuracy and F1-score.

**Table 6** shows how much training time is faster using the sampled training data instead of using the whole event data. There is a direct relation between the size reduction and  $R_t$  (refer to the results in **Table 3**). However, in most cases, the performance improvement is higher for the XGboost method. Considering the results in this table and **Table 5**, we found that using the sampling method we are able to improve the performance of the next activity prediction methods while they provide similar results. However, if oversampling (e.g., applying *log10* for *RTFM*) will result lower accuracy.

### 5.2.2 Remaining Time Prediction

In **Table 7** and **Table 8**, we show how by using the sampled event logs, MAE and RMSE of different remaining time prediction methods are changed. The results indicate that for the LSTM method independent of the sampling method for all event logs, we are able to provide the prediction similar to the case that whole event logs are used. It seems that the settings that are used for training the prediction method are not good. In other words, the trained model is not accurate enough. We repeat this experiment for LSTM with several different parameters, but we have almost the same results.

**Table 7**  $R_{MAE}$  of different event logs when different sampling methods are used.

| Log        | LSTM  |      |      |      |      |      |      |       |        | XGB  |      |      |      |      |      |      |       |        |
|------------|-------|------|------|------|------|------|------|-------|--------|------|------|------|------|------|------|------|-------|--------|
|            | d2    | d3   | d5   | d10  | log2 | log3 | log5 | log10 | Unique | d2   | d3   | d5   | d10  | log2 | log3 | log5 | log10 | Unique |
| BPIC2012-W | 1.000 | .999 | .999 | .995 | .999 | .999 | .999 | .999  | .999   | .962 | .938 | .911 | .896 | .928 | .917 | .836 | .789  | .888   |
| RTFM       | .999  | .998 | .998 | .998 | .998 | .998 | .998 | .998  | .998   | .899 | .695 | .523 | .367 | .175 | .160 | .147 | .098  | .122   |
| Sepsis     | .996  | .995 | .994 | .994 | .994 | .994 | .994 | .994  | .994   | .917 | .842 | .734 | .639 | .653 | .609 | .513 | .352  | .468   |

**Table 8**  $R_{RMSE}$  of different event logs when different sampling methods are used.

| Log        | LSTM  |       |       |      |       |       |       |       |        | XGB  |      |      |      |      |      |      |       |        |
|------------|-------|-------|-------|------|-------|-------|-------|-------|--------|------|------|------|------|------|------|------|-------|--------|
|            | d2    | d3    | d5    | d10  | log2  | log3  | log5  | log10 | Unique | d2   | d3   | d5   | d10  | log2 | log3 | log5 | log10 | Unique |
| BPIC2012-W | 1.000 | 1.000 | 1.000 | .988 | 1.000 | 1.000 | 1.000 | 1.000 | 1.000  | .977 | .962 | .945 | .933 | .932 | .930 | .879 | .833  | .925   |
| RTFM       | .999  | .999  | .999  | .999 | .999  | .999  | .999  | .999  | .999   | .857 | .732 | .503 | .304 | .244 | .227 | .214 | .145  | .190   |
| Sepsis     | .999  | .999  | .999  | .999 | .999  | .999  | .999  | .999  | .999   | .978 | .908 | .857 | .813 | .630 | .624 | .572 | .463  | .731   |

**Table 9**  $R_t$  of different event logs when different sampling methods are used.

| Log        | LSTM |      |      |      |       |       |       |        |        | XGB  |      |      |       |        |        |        |        |        |
|------------|------|------|------|------|-------|-------|-------|--------|--------|------|------|------|-------|--------|--------|--------|--------|--------|
|            | d2   | d3   | d5   | d10  | log2  | log3  | log5  | log10  | Unique | d2   | d3   | d5   | d10   | log2   | log3   | log5   | log10  | Unique |
| BPIC2012-W | 1.47 | 1.89 | 2.47 | 3.28 | 5.24  | 7.62  | 11.17 | 7.23   | 2.50   | 1.30 | 1.54 | 1.44 | 1.67  | 3.71   | 4.76   | 8.07   | 9.93   | 1.91   |
| RTFM       | 1.93 | 2.79 | 3.64 | 9.25 | 65.17 | 73.17 | 80.73 | 102.47 | 69.11  | 2.34 | 3.71 | 5.96 | 13.04 | 151.28 | 206.34 | 269.74 | 178.07 | 288.12 |
| Sepsis     | 1.40 | 1.37 | 1.49 | 1.64 | 1.82  | 2.19  | 1.66  | 1.72   | 1.55   | 1.45 | 1.68 | 1.95 | 1.95  | 1.97   | 2.57   | 2.87   | 3.33   | 2.51   |

For the XGboost method, the results indicate that if we do not oversample, we can have high  $R_{MAE}$  and  $R_{RMSE}$ . In general, as the main attribute for the next activity prediction is the sequence of activities, it is less sensitive to sampling. However, to predict the remaining time, the other data attributes can be essential too. In other words, for the remaining prediction, we need larger sampled event logs.

In Table 9, it is shown how by sampling event logs, we are able to reduce the required training time and improve the performance of the remaining time prediction process. By considering the results in Table 9 and Table 6, as we have expected, by having higher  $R_S$  the  $R_t$  value is higher.

## 6 Discussion

In this section, we discuss the results that are illustrated in the previous section. The results indicate that we do not always have a typical trade-off between the accuracy of the trained model and the performance of the prediction procedure. For example, for the next activity prediction, there are some cases where the training process is much faster than the normal procedure, even though the trained model provides an almost similar or higher accuracy and F1-score. Thus, the proposed instance selection procedure can be applied when we aim to apply hyper-parameter optimization [37]. In this way, more settings can be analyzed in a limited time. Moreover, it is reasonable to use the proposed method when we aim to train an online prediction method or on naive hardware such as cell phones.

For example, for the *RTFM* event log—as there are some highly frequent variants—the division distribution may be more useful. In other words, independently of the used prediction method, if we change the distribution of

variants (e.g., using *unique* distribution), it is expected that the accuracy will sharply decrease. However, for event logs with a more uniform distribution, we can use *unique* distributions to sample event logs. Furthermore, the results indicate that the effect of the chosen distribution (i.e., *unique*, *division*, and *logarithmic*) is more important than the used *k-value*. It is mainly because the *logarithmic* distribution may remove some of the variants and the *unique* distribution change the frequency distribution of variants. Therefore, it would be interesting to investigate more on the characteristics of the given event log and suitable sampling parameters for such distribution. For example, if most variants of a given event log are unique (e.g., *Sepsis*), using the *logarithmic* distribution leads to have remarkable  $R_S$  and consequently,  $R_{FE}$  and  $R_t$  will be very high. However, we will lose most of the variants and the trained model might have poor predictions.

By analysis of the results, we have found that the infrequent activities can be ignored using some hyper-parameter settings. The significant difference between F1-score and Accuracy values in Table 2 indicates this problem too. Using the sampling methods that modify the distribution of the event logs such as the *unique* method can help the prediction methods to also consider these activities. However, as these activities are infrequent, improving in the prediction of them would not impact highly on the presented aggregated F1-score value.

Finally, in real scenarios, the process can change because of different reasons [38]. This phenomenon is usually called *concept drift*. By considering the whole event log for training the prediction model, it is most probable that these changes are not considered in the prediction. Using the proposed sampling procedure, and giving higher priorities to newer traces, it is expected that we are able to adapt to the changes faster, which may be critical for specific applications.

## Limitations

Comparing the results for the next activity and remaining time prediction, we found that predicting the remaining time of the process is more sensitive to sampling. In other words, this application requires more data to predict accurately. Therefore, for the cases that the target attribute depends more on other data attributes compared to variant information, we need to sample more data otherwise, the trained model might be inaccurate.

We have also found a critical problem in predictive monitoring. In some cases, specifically using LSTM for predicting the remaining time, the accuracy of the predictions is low. For the next activity prediction, it is possible that the prediction models almost ignore infrequent activities. In these cases, even if we use the training data for the evaluation, we do not have acceptable results. This problem in machine learning is called a high bias error [39]. In other words, the training is not efficient even using whole data and we need to change the prediction method (or its parameters).

Another important outcome of the results is that for different event logs, we should use different sampling methods to achieve the highest performance. Considering lots of different machine learning methods and their parameters can lead to an increase in the search space and the complexity for users. In other words, finding the right setting for sampling may be challenging in real scenarios. Therefore, it would be valuable to research the relationship between the event log characteristics and suitable sampling parameters that can be used for preprocessing the training event log.

## 7 Conclusion

In this paper, we proposed an instance selection approach to improve the performance predictive business process monitoring methods. We suggested that it is possible to use a sample of training event data instead of whole training event data. To evaluate the proposed approach, we consider two main applications of predictive business monitoring, i.e., the next activity and the remaining time prediction. Results of applying the proposed approaches on three real-life event logs and two widely used machine learning methods, i.e., LSTM and XGboost, indicate that in most cases, we are able to improve the performance of predictive monitoring algorithms while providing similar accuracy compared to the case that the whole training event logs are used. However, by oversampling, the accuracy of the trained model might be reduced. Moreover, we have found that the remaining time prediction application is more sensitive to sampling.

To continue this research, we aim to extend the experiments to study the relationship between the event log characteristics and the sampling parameters. In other words, we aim to help the end-user to adjust the sampling parameters based on the characteristics of the given event log. Moreover, it would be great to investigate how can we apply the proposed sampling procedure for streaming event data which is potentially one of the major advantages of the proposed method in a real-life setting. Finally, it is interesting to investigate more on feature selection methods for improving the performance of the predictive monitoring procedure. It is expected that similar to process instance sampling, feature selection methods are able to reduce the required training time. In other words, the training is not efficient even using whole data and we need to change the prediction method (or its parameters).

Another important outcome of the results is that for different event logs, we should use different sampling methods to achieve the highest performance. Considering lots of different machine learning methods and their parameters can lead to an increase in the search space and the complexity for users. In other words, finding the right setting for sampling may be challenging in real scenarios. Therefore, it would be valuable to research the relationship between the event log characteristics and suitable sampling parameters that can be used for preprocessing the training event log.

## Acknowledgment

The authors would like to thank the Alexander von Humboldt (AvH) Stiftung for funding this research.

## References

- [1] van der Aalst, W.M.P., Schonenberg, M.H., Song, M.: Time prediction based on process mining **36**(2), 450–475. <https://doi.org/10.1016/j.is.2010.09.001>
- [2] Hitfox Group, Breuker, D., Matzner, M., University of Muenster, Delfmann, P., University of Koblenz-Landau, Becker, J., University of Muenster: Comprehensible predictive models for business processes **40**(4), 1009–1034. <https://doi.org/10.25300/MISQ/2016/40.4.10>
- [3] Marquez-Chamorro, A.E., Resinas, M., Ruiz-Cortes, A.: Predictive monitoring of business processes: A survey **11**(6), 962–977. <https://doi.org/10.1109/TSC.2017.2772256>
- [4] Breiman, L.: Bagging predictors. *Machine learning* **24**(2), 123–140 (1996)
- [5] Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, August 13–17, 2016, pp. 785–794. ACM, ??? (2016)
- [6] Senderovich, A., Di Francescomarino, C., Ghidini, C., Jorbina, K., Maggi, F.M.: Intra and inter-case features in predictive process monitoring: A tale of two dimensions. In: *International Conference on Business Process Management*, pp. 306–323 (2017). Springer
- [7] Teinemaa, I., Dumas, M., Rosa, M.L., Maggi, F.M.: Outcome-oriented predictive process monitoring: Review and benchmark. *ACM Transactions on Knowledge Discovery from Data (TKDD)* **13**(2), 1–57 (2019)
- [8] Evermann, J., Rehse, J., Fettke, P.: Predicting process behaviour using deep learning. *Decis. Support Syst.* **100**, 129–140 (2017)
- [9] Zhou, L., Pan, S., Wang, J., Vasilakos, A.V.: Machine learning on big data: Opportunities and challenges. *Neurocomputing* **237**, 350–361 (2017). <https://doi.org/10.1016/j.neucom.2017.01.026>
- [10] Pourghassemi, B., Zhang, C., Lee, J.H., Chandramowlishwaran, A.: On the limits of parallelizing convolutional neural networks on gpus. In: *SPAA ’20: 32nd ACM Symposium on Parallelism in Algorithms and Architectures*, Virtual Event, USA, July 15–17, 2020, pp. 567–569. ACM, ???

- 18     *Performance-Preserving Event Log Sampling for Predictive Monitoring*  
(2020)
- [11] Garca, S., Luengo, J., Herrera, F.: Data Preprocessing in Data Mining. Springer, ??? (2014)
  - [12] Wilson, D.R., Martinez, T.R.: Reduction techniques for instance-based learning algorithms. *Mach. Learn.* **38**(3), 257–286 (2000). <https://doi.org/10.1023/A:1007626913721>
  - [13] Wilson, D.L.: Asymptotic properties of nearest neighbor rules using edited data. *Systems, Man and Cybernetics, IEEE Transactions on* **2**(3), 408–421 (1972). <https://doi.org/10.1109/TSMC.1972.4309137>
  - [14] de Leoni, M., van der Aalst, W.M.P., Dees, M.: A general process mining framework for correlating, predicting and clustering dynamic behavior based on event logs **56**, 235–257. <https://doi.org/10.1016/j.is.2015.07.003>
  - [15] van der Aalst, W.M.P.: *Process Mining - Data Science in Action*, Second Edition. Springer, ??? (2016)
  - [16] Huang, Z., Xu, W., Yu, K.: Bidirectional LSTM-CRF models for sequence tagging. *CoRR* **abs/1508.01991** (2015) [arXiv:1508.01991](https://arxiv.org/abs/1508.01991)
  - [17] Sani, M.F., Vazifehdoostirani, M., Park, G., Pegoraro, M., van Zelst, S.J., van der Aalst, W.M.: Event log sampling for predictive monitoring. In: *Machine Learning for Process Mining* (2021)
  - [18] Polato, M., Sperduti, A., Burattin, A., de Leoni, M.: Time and activity sequence prediction of business process instances. *Computing* **100**(9), 1005–1031 (2018)
  - [19] Park, G., Song, M.: Predicting performances in business processes using deep neural networks. *Decis. Support Syst.* **129** (2020)
  - [20] Rogge-Solti, A., Weske, M.: Prediction of remaining service execution time using stochastic petri nets with arbitrary firing delays. In: Basu, S., Pautasso, C., Zhang, L., Fu, X. (eds.) *Service-Oriented Computing - 11th International Conference, ICSOC 2013, Berlin, Germany, December 2-5, 2013, Proceedings*, vol. 8274, pp. 389–403. Springer, ??? (2013)
  - [21] Teinemaa, I., Dumas, M., Maggi, F.M., Di Francescomarino, C.: Predictive business process monitoring with structured and unstructured data. In: *International Conference on Business Process Management*, pp. 401–417 (2016). Springer
  - [22] Tax, N., Verenich, I., Rosa, M.L., Dumas, M.: Predictive business process monitoring with LSTM neural networks. In: Dubois, E., Pohl, K. (eds.)

- Advanced Information Systems Engineering - 29th International Conference, CAiSE 2017, Essen, Germany, June 12-16, 2017, Proceedings, vol. 10253, pp. 477–492. Springer, ??? (2017)
- [23] Navarin, N., Vincenzi, B., Polato, M., Sperduti, A.: LSTM networks for data-aware remaining time prediction of business process instances. In: 2017 IEEE Symposium Series on Computational Intelligence, SSCI 2017, Honolulu, HI, USA, November 27 - Dec. 1, 2017, pp. 1–7. IEEE, ??? (2017)
  - [24] Nguyen, A., Chatterjee, S., Weinzierl, S., Schwinn, L., Matzner, M., Eskofier, B.M.: Time matters: Time-aware lstms for predictive business process monitoring. In: Leemans, S.J.J., Leopold, H. (eds.) *Process Mining Workshops - ICPM 2020 International Workshops*, Padua, Italy, October 5-8, 2020, Revised Selected Papers, vol. 406, pp. 112–123. Springer, ??? (2020)
  - [25] Pegoraro, M., Uysal, M.S., Georgi, D.B., van der Aalst, W.M.P.: Text-aware predictive monitoring of business processes. In: Abramowicz, W., Auer, S., Lewanska, E. (eds.) *24th International Conference on Business Information Systems, BIS 2021, Hannover, Germany, June 15-17, 2021*, pp. 221–232 (2021)
  - [26] Verenich, I.: Explainable predictive monitoring of temporal measures of business processes. In: Depaire, B., Smedt, J.D., Dumas, M., Fahland, D., Kumar, A., Leopold, H., Reichert, M., Rinderle-Ma, S., Schulte, S., Seidel, S., van der Aalst, W.M.P. (eds.) *Proceedings of the Dissertation Award, Doctoral Consortium, and Demonstration Track at BPM 2019 Co-located with 17th International Conference on Business Process Management, BPM 2019, Vienna, Austria, September 1-6, 2019. CEUR Workshop Proceedings*, vol. 2420, pp. 26–30. CEUR-WS.org, ??? (2019). <http://ceur-ws.org/Vol-2420/paperDA6.pdf>
  - [27] Stierle, M., Brunk, J., Weinzierl, S., Zilker, S., Matzner, M., Becker, J.: Bringing light into the darkness - A systematic literature review on explainable predictive business process monitoring techniques. In: Rowe, F., Amrani, R.E., Limayem, M., Matook, S., Rosenkranz, C., Whitley, E.A., Quammah, A.E. (eds.) *28th European Conference on Information Systems - Liberty, Equality, and Fraternity in a Digitizing World*, ECIS 2020, Marrakech, Morocco, June 15-17, 2020 (2021). [https://aisel.aisnet.org/ecis2021\\_rip/8](https://aisel.aisnet.org/ecis2021_rip/8)
  - [28] Sindhgatta, R., Moreira, C., Ouyang, C., Barros, A.: Exploring interpretable predictive models for business processes. In: Fahland, D., Ghidini, C., Becker, J., Dumas, M. (eds.) *Business Process Management*

- 18th International Conference, BPM 2020, Seville, Spain, September 13-18, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12168, pp. 257–272. Springer, ??? (2020). [https://doi.org/10.1007/978-3-030-58666-9\\_15](https://doi.org/10.1007/978-3-030-58666-9_15). [https://doi.org/10.1007/978-3-030-58666-9\\_15](https://doi.org/10.1007/978-3-030-58666-9_15)
- [29] Galanti, R., Coma-Puig, B., de Leoni, M., Carmona, J., Navarin, N.: Explainable predictive process monitoring. In: van Dongen, B.F., Montali, M., Wynn, M.T. (eds.) 2nd International Conference on Process Mining, ICPM 2020, Padua, Italy, October 4-9, 2020, pp. 1–8. IEEE, ??? (2020). <https://doi.org/10.1109/ICPM49681.2020.00012>. <https://doi.org/10.1109/ICPM49681.2020.00012>
- [30] Pauwels, S., Calders, T.: Incremental predictive process monitoring: The next activity case. In: Polyvyanyy, A., Wynn, M.T., Looy, A.V., Reichert, M. (eds.) Business Process Management - 19th International Conference, BPM 2021, Rome, Italy, September 06-10, 2021, Proceedings. Lecture Notes in Computer Science, vol. 12875, pp. 123–140. Springer, ??? (2021). [https://doi.org/10.1007/978-3-030-85469-0\\_10](https://doi.org/10.1007/978-3-030-85469-0_10). [https://doi.org/10.1007/978-3-030-85469-0\\_10](https://doi.org/10.1007/978-3-030-85469-0_10)
- [31] Wang, T., Zhu, J.-Y., Torralba, A., Efros, A.A.: Dataset distillation. arXiv preprint arXiv:1811.10959 (2020) [arXiv:1811.10959](https://arxiv.org/abs/1811.10959) [cs.LG]
- [32] Fani Sani, M., van Zelst, S.J., van der Aalst, W.M.P.: The impact of biased sampling of event logs on the performance of process discovery. *Computing* **103**(6), 1085–1104 (2021)
- [33] Fani Sani, M., van Zelst, S.J., van der Aalst, W.M.P.: Conformance checking approximation using subset selection and edit distance. In: Advanced Information Systems Engineering - 32nd International Conference, CAISE 2020, Grenoble, France, June 8-12, 2020, Proceedings, vol. 12127, pp. 234–251. Springer, ??? (2020)
- [34] Qafari, M.S., van der Aalst, W.M.P.: Root cause analysis in process mining using structural equation models. In: Business Process Management Workshops - BPM 2020 International Workshops, Seville, Spain, September 13-18, 2020, Revised Selected Papers, vol. 397, pp. 155–167. Springer, ??? (2020)
- [35] Verenich, I., Dumas, M., Rosa, M.L., Maggi, F.M., Teinemaa, I.: Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring. *ACM Transactions on Intelligent Systems and Technology (TIST)* **10**(4), 1–34 (2019)
- [36] Luque, A., Carrasco, A., Martín, A., de las Heras, A.: The impact of class imbalance in classification performance metrics based on the binary confusion matrix. *Pattern Recognit.* **91**, 216–231 (2019). <https://doi.org/>

[10.1016/j.patcog.2019.02.023](https://doi.org/10.1016/j.patcog.2019.02.023)

- [37] Bergstra, J., Bardenet, R., Bengio, Y., Kégl, B.: Algorithms for hyperparameter optimization. In: *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a Meeting Held 12-14 December 2011, Granada, Spain*, pp. 2546–2554 (2011)
- [38] Carmona, J., Gavalda, R.: Online techniques for dealing with concept drift in process mining. In: *Advances in Intelligent Data Analysis XI - 11th International Symposium, IDA 2012, Helsinki, Finland, October 25-27, 2012. Proceedings*, vol. 7619, pp. 90–102. Springer, ??? (2012)
- [39] van der Putten, P., van Someren, M.: A bias-variance analysis of a real world learning problem: The coil challenge 2000. *Mach. Learn.* **57**(1-2), 177–195 (2004). <https://doi.org/10.1023/B:MACH.0000035476.95130.99>