

# Recurrent Process Mining with Live Event Data

Alifah Syamsiyah, Boudewijn F. van Dongen, Wil M.P. van der Aalst

Eindhoven University of Technology

A.Syamsiyah@tue.nl, B.F.v.Dongen@tue.nl, W.M.P.v.d.Aalst@tue.nl

**Abstract.** In organizations, process mining activities are typically performed in a recurrent fashion, e.g. once a week, an event log is extracted from the information systems and a process mining tool is used to analyze the process' characteristics. Typically, process mining tools import the data from a file-based source in a pre-processing step, followed by an actual process discovery step over the pre-processed data in order to present results to the analyst. As the amount of event data grows over time, these tools take more and more time to do pre-processing and all this time, the business analyst has to wait for the tool to finish. In this paper, we consider the problem of recurrent process discovery in live environments, i.e. in environments where event data can be extracted from information systems near real time. We present a method that pre-processes each event when it is being generated, so that the business analyst has the pre-processed data at his/her disposal when starting the analysis. To this end, we define a notion of intermediate structure between the underlying data and the layer where the actual mining is performed. This intermediate structure is kept in a persistent storage and is kept live under updates. Using a state of the art process mining technique, we show the feasibility of our approach. Our work is implemented in the process mining tool ProM using a relational database system as our persistent storage. Experiments are presented on real-life event data to compare the performance of the proposed approach with the state of the art.

**Keywords:** Recurrent proses mining · Live event data · Incremental process discovery

## 1 Introduction

Process mining is a discipline where the aim is to improve an organization's processes given the information from the so called *event logs* [16]. Process mining techniques have been successfully demonstrated in various case studies such as health care, insurance, and finance [5, 10, 12, 23]. In many of these cases, a one-time study was performed, but in practice, process mining is typically a recurring activity, i.e. an activity that is performed on a routine basis.

As an illustration, suppose that each manager of an insurance company has the obligation to report her/his work to a director once in each month to see

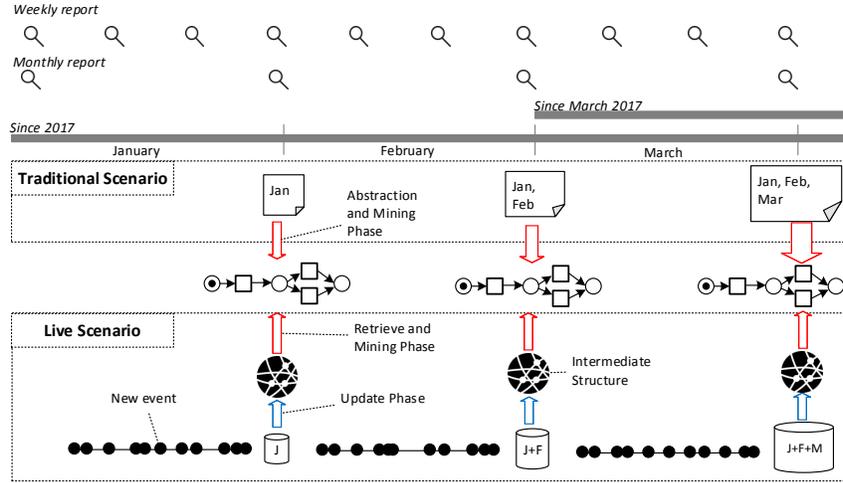


Fig. 1. Traditional vs live scenario in process discovery

the company’s progress since the beginning of a year. Typical analysis in such regular report incorporates the observations from previous month, last three months, or last year. This type of reporting requires an analyst to repeatedly produce analysis results from data that grows over time.

Existing process mining tools are not tailored towards such recurrent analyses. Instead, they require the analysts to export the event data from a running system, import it to the mining tool which pre-processes the data during importing and then use the tool on the pre-processed data. As the amount of data grows, the import and pre-processing phase takes longer and longer causing the analyst to waste time.

However, most information systems nowadays record real-time event data about business processes during their executions. This enables analysis techniques to import and pre-process the data when it “arrives”. By shifting the pre-processing time an analyst is able to do process mining on the pre-processed data instantly.

The idea of this live scenario is sketched in Figure 1. In the live scenario, we introduce a persistent storage for various structures that are kept as “an intermediate structure” by process mining algorithms. We then show how such intermediate structure can be updated without the need for full recomputation every time an event arrives. Using the intermediate structure of the state-of-the-art process mining technique, we show the feasibility of our approach in the general process mining setting and using experiments on real-life data, we show the added time-benefit for the analyst.

This paper is organized as follows. In Section 2 we introduce recurrent process mining and we focus on a traditional technique in this setting. In Section 3, we

show how recurrent process mining can be performed on live event data and we prove that the technique of Section 2 can be made incremental. We show the improvements of our work using real-life data in Section 4. Then, in Section 5 discusses some related work. Finally, we conclude the paper in Section 6.

## 2 Recurrent Process Mining

The challenge in process mining is to gain insights into an operational system in which activities are executed. The execution of these activities is reflected through events, i.e. events are occurrences of activities. Furthermore, in process mining, activities are being executed in the context of cases (also referred to as process instances) at a particular point in time.

### Definition 1 (Events, cases, and activities).

Let  $\mathcal{A}$  be the universe of activities, let  $\mathcal{C}$  be the universe of cases, and let  $\mathcal{Y}$  be the universe of events. We define  $\#_{act} : \mathcal{Y} \rightarrow \mathcal{A}$  a function labeling each event with an activity. We define  $\#_{cs} : \mathcal{Y} \rightarrow \mathcal{C}$  a function labeling each event with a case. We define  $\#_{tm} : \mathcal{Y} \rightarrow \mathbb{R}$  a function labeling each event with a timestamp.

In process mining, the general assumption is that event data is provided in the form of an event log. Such an event log is basically a collection of events occurring in a specific time period, in which the events are grouped by cases sequentialized based on their time of occurrence.

### Definition 2 (Event log and trace).

Let  $E \subseteq \mathcal{Y}$  be a collection of events and  $t_s, t_e \in \mathbb{R}$  two timestamps with  $t_s < t_e$  relating to the start and the end of the collection period.

A trace  $\sigma \in E^*$  is a sequence of events such that the same event occurs only once in  $\sigma$ , i.e.  $|\sigma| = |\{e \in \sigma\}|$ . Furthermore, each event in a trace refers to the same case  $c \in \mathcal{C}$ , i.e.  $\forall e \in \sigma \#_{cs}(e) = c$  and we assume all events within the given time period are included, i.e.  $\forall e \in \mathcal{Y} (\#_{cs}(e) = c \wedge t_s \leq \#_{tm}(e) \leq t_e) \implies e \in \sigma$ .

An event log  $L \in \mathcal{O}(E^*)^1$  is a set of traces.

Note that the time-period over which an event log is collected plays an important role in this paper. In most process mining literature, this time period is neglected and the event log is assumed to span eternity. However, in practice, analysis always consider event data pertaining to a limited period of time. Therefore, in this paper, we explicitly consider the following process mining scenarios (as depicted in Figure 1):

- Analysts perform process mining tasks on a recurrent schedule at regular points, e.g. once a week about the last week, or once a month about the last month, or
- Analysts perform process mining on the data since a pre-determined point in time, e.g. since 2017, or since March 2017.

<sup>1</sup>  $\mathcal{O}(E^*)$  denotes a powerset of sequences, i.e.  $L \subseteq E^*$

## 2.1 Traditional Recurrent Process Mining

To execute the two process mining scenarios, a multitude of process mining techniques is available. However, all have two things in common, namely (1) that the starting point for analysis is an event log and (2) that the analysis is performed in three phases, namely: *loading*, *abstraction*, and *mining*.

While the details may differ, all process mining techniques build an *intermediate structure* in memory during the abstraction phase. Typically, the time needed to execute this phase is linear in the number of events in the event log. This intermediate structure (of which the size is generally polynomial in the number of activities in the log, but not related to the number of events anymore) is then used to perform the actual mining task. In this paper, we consider the state-of-the-art process discovery algorithm, namely the Inductive Miner [6, 7], which is known to be flexible, to have formal guarantees, and to be scalable. For the Inductive Miner, the intermediate structure is the so-called *direct succession relation* (Definition 3) which counts the frequencies of direct successions.

### Definition 3 (Direct succession [16]).

Let  $L$  be an event log over  $E \subseteq \mathcal{Y}$ . The direct succession relation  $>_L : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{N}$  counts the number of times activity  $a$  is directly followed by activity  $b$  in some cases in  $L$  as follows:

$$>_L(a, b) = \sum_{\sigma \in L} \sum_{i=1}^{|\sigma|-1} \begin{cases} 1, & \text{if } \#_{act}(\sigma(i)) = a \wedge \#_{act}(\sigma(i+1)) = b \\ 0, & \text{otherwise.} \end{cases}$$

To illustrate how the Inductive Miner algorithm uses the direct succession relation described above, we refer to Table 1. Here, we show (by example), how the mining technique uses the relation as an intermediate structure to come to a result. In order to apply this traditional technique in a recurrent setting, the loading of the data, the abstraction phase, and the mining phase have to be repeated. When over time the event data grows, the time to execute the three phases also grows, hence performing the recurrent mining task considering one year of data will take 52 times longer than considering one week of data.

In Table 2 we show an example of the Inductive Miner applied to a real-life dataset of the BPI Challenge 2017 [17] which contains data of a full year. We record the times to perform the three phases of importing, abstraction, and mining on this dataset after the first month, at the middle of the year, and at the end of the year. It is clear that indeed the importing and abstraction times grow considerably, while the actual mining phase is orders of magnitude faster.

×	An exclusive-choice cut of $L$ is a cut $(\times, A_1, \dots, A_n)$ such that $\forall i, j \in \{1, \dots, n\} \forall a \in A_i \forall b \in A_j \ i \neq j \Rightarrow \not>_L(a, b)$
→	A sequence cut of $L$ is a cut $(\rightarrow, A_1, \dots, A_n)$ such that $\forall i, j \in \{1, \dots, n\} \forall a \in A_i \forall b \in A_j \ i < j \Rightarrow (>_L^+(a, b) \wedge \not>_L^+(b, a))$

Table 1: Examples of the use of the intermediate structure in Inductive Miner

Week	Inductive Miner		
	Loading	Abstraction	Mining
5	0.8520	2.8531	0.0254
26	3.6319	30.9528	0.0257
52	9.5854	93.5118	0.0291

Table 2: Process mining times (in seconds) in the traditional setting on data of the BPI Challenge 2017 [17]

Figure 2 shows the result of the Inductive Miner after the first 5 weeks of data, i.e. if an analyst has produced this picture on January 29<sup>th</sup> 2016, it would have taken 3.7305 seconds ( $= 0.8520 + 2.8531 + 0.0254$ ) in total to load the event log, build the abstraction, and do the mining in order to produce this picture using the Inductive Miner in ProM.

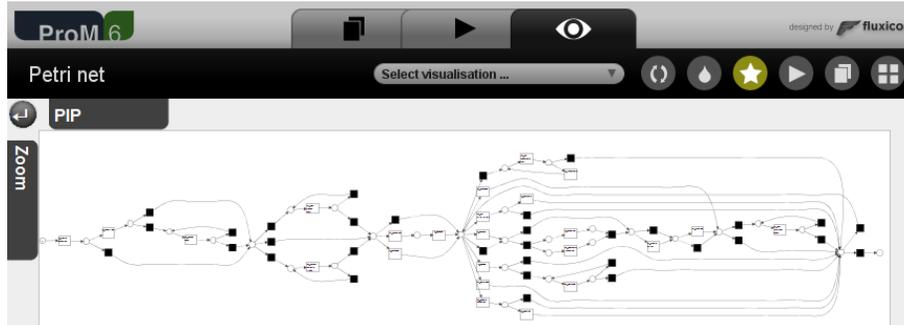


Fig. 2. Screenshot of ProM showing the result of the Inductive Miner on the BPI Challenge 2017 data considering the first 5 weeks of data.

In Section 3, we present a method to store the intermediate structure in a persistent storage and to keep it live under updates, i.e. every time an event is generated, the intermediate structure is updated (in constant time). This way, when a process mining task is performed, the time spent by the analyst is limited to the time to retrieve the intermediate structure and to do the actual mining.

### 3 Recurrent Process Mining with Live Event Data

It is easy to see that given an event log  $L$ , the relation in Definition 3 can be computed during a single linear pass over the event log, visiting each event exactly once. In this section, we present a live process mining technique based on the Inductive Miner which does not require the analyst to repeat the import and analysis phases every time a process mining task is performed.

In order to enable live process mining, we use a persistent event storage, called DB-XES [13], which uses a relational database to store event data. The

full structure of this relational database is beyond the scope of this paper, but what is important is that given a trace, it is possible to quickly retrieve the last event recorded for that trace.

**Definition 4 (Last event in a trace).**

Let  $E \subseteq \mathcal{Y}$  be a collection of events and let  $c \in C$  be a case. The function  $\lambda : C \rightarrow E \cup \{\perp\}$  is a function that returns the last event in  $E$  belonging to case  $c$ , i.e.

$$\lambda(c) = \begin{cases} \perp, & \text{if } \forall_{e \in E} \#_{cs}(e) \neq c, \\ e \in E, & \text{if } \#_{cs}(e) = c \wedge \nexists e' \in E (\#_{cs}(e') = c \wedge \#_{tm}(e') > \#_{tm}(e)). \end{cases}$$

Using DB-XES as a persistent storage and making use of the ability to query for the last event in a trace, we present the Incremental Inductive Miner in Section 3.1.

### 3.1 Incremental Inductive Miner

The Inductive Miner uses only the frequency of direct successions between activities as input as defined in Definition 3. Therefore, to enable an incremental version of the Inductive Miner, we present an update strategy that, given the relation  $>_L$  for some log  $L$  and a new event  $e$ , we can derive the relation  $>_{L'}$  where  $L'$  is the log  $L$  with the additional event  $e$ .

**Theorem 1 (Updating relation  $>$  is possible).**

Let  $E \subseteq \mathcal{Y}$  be a set of events and  $L$  a log over  $E$ . Let  $e \in \mathcal{Y} \setminus E$  be a fresh event to be added such that for all  $e' \in E$  holds  $\#_{ts}(e') < \#_{ts}(e)$  and let  $E' = E \cup \{e\}$  be the new set of events with  $L'$  the corresponding log over  $E'$ . We know that for all  $a, b \in \mathcal{A}$  holds that:

$$>_{L'}(a, b) = >_L(a, b) + \begin{cases} 0 & \text{if } \lambda(\#_{cs}(e)) = \perp, \\ 1 & \text{if } \#_{act}(\lambda(\#_{cs}(e))) = a \wedge \#_{act}(e) = b, \\ 0 & \text{otherwise.} \end{cases}$$

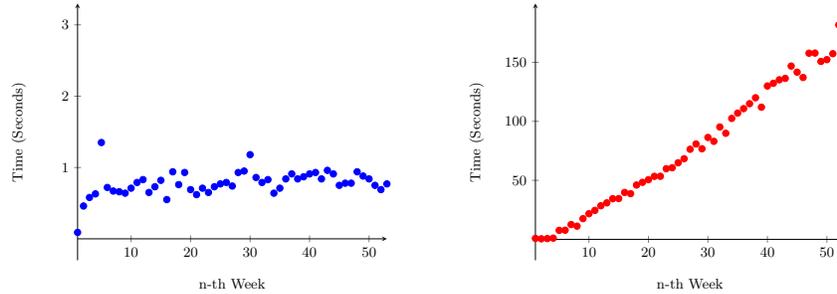
*Proof.* Let  $c = \#_{cs}(e) \in C$  be the case to which the fresh event belongs.

If for all  $e' \in E$  holds that  $\#_{cs}(e') \neq c$ , then this is the first event in case  $c$  and we know that  $L' = L \cup \langle e \rangle$ . Hence relation  $>$  does not change from  $L$  to  $L'$  as indicated by case 1.

If there exists  $e' = \lambda(\#_{cs}(e)) \in E$  with  $\#_{cs}(e') = c$ , then we know that there is a trace  $\sigma_c \in L$ . Furthermore, we know that  $L' = (L \setminus \{\sigma_c\}) \cup \{\sigma_c \cdot \langle e \rangle\}$ , i.e. event  $e$  gets added to trace  $\sigma_c$  of  $L$ . This implies that  $>_{L'}(\#_{act}(e'), \#_{act}(e)) = >_L(\#_{act}(e'), \#_{act}(e)) + 1$  as indicated by case 2.

In all other cases, the number of direct successions of two activities is not affected. ■

Using the simple update procedure indicated in Theorem 1 the Incremental Inductive Miner allows for recurrent process mining under live updates. In Section 4 we show how the effect of keeping relation  $>$  lives on the total time needed to perform the recurrent process mining task.



**Fig. 3.** The comparison of recurrent process discovery using DIIM (left) vs traditional Inductive Miner (right)

## 4 Experimental Results

We implemented the algorithm presented as a ProM<sup>2</sup> plug-in called *Database-Incremental Inductive Miner (DIIM)*<sup>3</sup>. DIIM is designed for recurrent process discovery based on live event data. It uses DB-XES as the back-end storage and the Inductive Miner as the mining algorithm. In this section, we show the experimental results of applying DIIM in a live scenario and we compare it to traditional Inductive Miner.

For the experiment, we used a real dataset from BPI Challenge 2017 [17]. This dataset pertains to the loan applications of a company from January 2016 until February 2017. In total, there are 1,202,267 events and 26 different activities which pertain to 31,509 loan applications.

In this experiment, we looked into some weekly reports where we were interested to see process models of the collective data since 2016. The last working day, i.e. Friday, was chosen as the day when we performed the process discovery to have a progress report for that week. In live scenario, we assumed that each event was inserted to the DB-XES database precisely at the time stated in the timestamp attribute of the event log. Then, the DB-XES system immediately processed each new event data as it arrived using triggers in the relational database, implementing the update procedures detailed in Section 3, thus keeping the relations live under updates. In traditional scenario, we split the dataset into several logs such that each log contained data for one week. For the  $n$ -th report, we combined the log from the first week until the  $n$ -th week, loaded it into ProM, and discovered a process model.

Figure 3 shows the experimental results of recurrent process discovery using DIIM and the Inductive Miner. The x-axis represents the  $n$ -th week, while the y-axis represents the time spent by user (in seconds) to discover process models. The blue dots are the experiment using DIIM which includes the total times to insert new events, update the intermediate structure, retrieve the values from the

<sup>2</sup> See <http://www.processmining.org> and <http://www.promtools.org>

<sup>3</sup> <https://svn.win.tue.nl/repos/prom/Packages/DatabaseInductiveMiner/Trunk/>

DB-XES, and mine the process models, while the red dots are the experiment using traditional Inductive Miner which includes the time to load the XES event logs, build the intermediate structure, and mine the process models.

As shown in Figure 3, our incremental technique is much faster, even when considering the time needed to insert events in the relational database, a process that is typically executed in real time and without the business analyst being present. More important however, is the trendlines of both approaches.

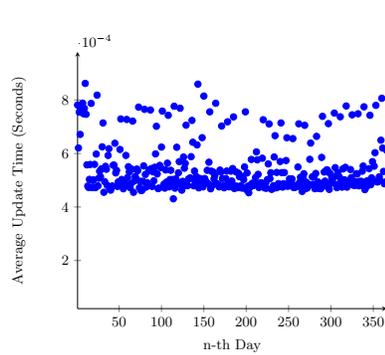
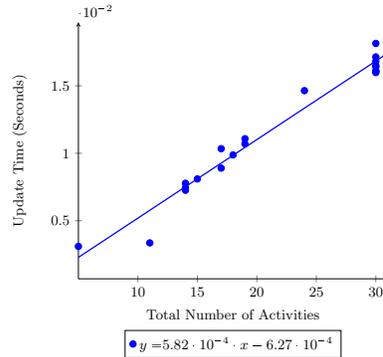
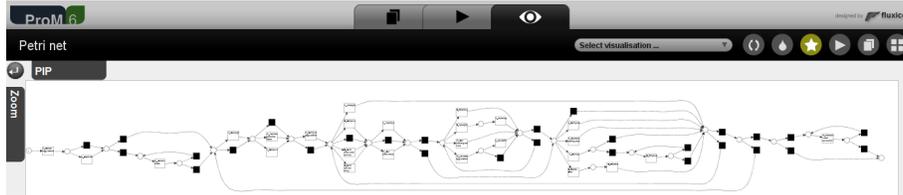
As expected, the time to perform the process mining task in the traditional setting is growing linear in the size of the event data (the arrival rate of events in this dataset is approximately constant during the entire year). This is due to the fact that the first two phases of loading the data and doing the abstraction into the intermediate structure scales linearly in the number of events, whereas the mining scales in the number of activities. The latter is considerably smaller than the former in most practical cases as well as in this example. Our incremental algorithms are more stable over time as the update time only depends on the number of *newly inserted* events and both the retrieval and mining times depend on the number of activities rather than the number of events.

The variations in the recorded values of the DIIM are therefore explained by the number of inserted events in a day. The higher the number of newly inserted events, the longer it takes to do the update in the relational database system of the intermediate structure. However, the total update time remains limited to around 1.4 seconds per day.

In order to see the average time for doing an update for a single event, we normalized the total update time with the number of events inserted in each day as shown in Figure 4. The x-axis represents the  $n$ -th day, while the y-axis represents the update time per event. As shown from the Figure 4, the average time to conduct an update for a single event stabilizes around 0.000545 seconds, i.e. the database system can handle around 1800 events per second and this includes the insertion of the actual event data in the underlying DB-XES tables.

To validate the fact that the update time scales linearly in the number of activities, we refer to Figure 5. For this experiment, we used a different dataset with 31 different activities and eleven thousands of events, provided to us by Xerox Services, India. The x-axis represents the total number of different activities which has been inserted to the database, while the y-axis represents the time in seconds to update an event. From the figure, it is clear that the update time indeed depends linearly on the number of activities.

It is important to realize that the results of the process mining techniques in both the traditional and the live setting are not different, i.e. the process models are identical. Figure 6 shows a screenshot of a process model in ProM, produced by the Incremental Inductive Miner considering all the data in the original file. In the traditional setting, it would have taken an analyst 103.1263 seconds to load the event log, build the abstraction and do the mining in order to get this picture on December 30<sup>th</sup> 2016. Due to the availability of the intermediate structure in the database, it would take the analyst only 0.0392 seconds to produce the same result using the Incremental Inductive Miner.


**Fig. 4.** Average update time per event

**Fig. 5.** The influence of number of activities to update time

**Fig. 6.** Screenshot of ProM showing the result of the Inductive Miner on the BPI Challenge 2017 data considering all data.

## 5 Related Work

For a detailed explanation about process mining, we refer to [16]. Here we primarily focus on the work related to recurrent process discovery in process mining and its applications on live event data.

In the current setting of process discovery, event data from a file-based system is imported to a process mining tool. This technique potentially creates redundancy of data reloading in environments which necessitate some repetitions in the discovery. Therefore, some researches have been looking to the area of databases, Hadoop, and other ways to store event data in a persistent manner.

A study in [22] examined a tool called XESame. To access the data in XESame, one needs to materialize the data by selecting and matching it with XES [4] elements. It does not provide a direct access to the database. A more advanced technique using ontology was proposed in [1, 2]. In this work, data can be accessed on-demand using query unfolding and rewriting techniques, called ontology-based data access. However, performance issues make this approach unsuitable for large event logs.

Relational XES, or RXES, was introduced in [18]. The RXES schema was designed with a focus on XES standard. Using experiments with real life data,

it was shown that RXES typically uses less memory compared to the file-based OpenXES and MapDB XESLite implementations [9]. As an improved version of RXES, DB-XES was introduced in [13] to enable process mining in the large. In [14] DB-XES basic schema is extended to allow instant social network mining, especially the handover of work networks.

Process mining not only covers procedural process discovery, but also declarative process discovery. The work in [11] deals with declarative process discovery using SQL databases. Building on top of RXES, the authors introduce a mining approach that directly works on relational event data by querying the log with conventional SQL. Queries can be customised and cover process perspective beyond the control flow, e.g., organisational aspects. However, none of these techniques handles live event data, the focus is often on static data that has been imported in a database.

Treating event data as a dynamic sequence of events has been explored in [19,20]. This work presented single-node stream extension of process mining tool ProM which enables researchers and business users to also apply process mining on streaming-data. The applicability of the framework on the cooperative aspects of process mining was demonstrated in [21]. Here the authors define the notion of case administration to store for each case the (partial) sequence of (activity, resource)-pairs seen so far.

Another online process discovery technique based on streaming technology was proposed in [8]. This work presented a novel framework for the discovery of LTL-based declarative process models from streaming event data. The framework continuously updates a set of valid business constraints based on the events occurring in the event stream. Moreover, it gives the user meaningful information about the most significant concept drifts occurring during process execution.

However, from the real-time techniques using streaming event data that we have seen so far, none of them deals with recurrent process discovery. Intermediate results are not stored after presenting the results. Therefore, all data (old and new) needs to be reloaded and processed each time a new analysis is needed. Building on from that concern, this work explores both in the ability to process live event data and to handle recurrent questions in process discovery.

## 6 Conclusion

Process mining aims to give insights into processes by discovering process models which adequately reflect the behavior seen in an event log. One of the challenges in process mining is the recurrent nature of mining tasks which, using traditional tools and techniques, require analysts to import and pre-process larger and larger datasets.

In this paper we focused on recurrent process discovery on live event data. Using the Inductive Miner technique as an example, we show how we can reduce the time needed for an analyst to do process mining by storing intermediate structure in a persistent storage. Then we show how to keep this intermediate structure alive under insertion of new events.

Using a concrete implementation, we use the relational database called DB-XES to store the event data and the intermediate structure. In this relational database, we added triggers to update the intermediate structure with the insertion of each event and we implemented the Incremental Inductive Miner as a version of the existing technique which is able to use this persistent intermediate structure directly.

We tested the performance of the proposed approach and compared it to the traditional techniques using real-life datasets. We show that loading and mining time of the traditional approach grows linearly as the event data grows. In contrast, our incremental implementation shows constant times for updating (per event) and the retrieval and mining times are independent of the size of the underlying data.

The core ideas in the paper are not limited to control flow. They are, for example, trivially extended to store intermediate structures keeping track of average times between activities or for social networks. Moreover, they are not restricted towards procedural process discovery. In [15] we show how the work extends into declarative process discovery, particularly using MINERful [3] as the discovery technique. We introduce a so-called controller function which we keep live under updates. Then we show that, using the controller function, we can keep all MINERful relations live under updates.

A more fundamental challenge for future work is the updating of the intermediate structures in batches of events, rather than for each event separately. Furthermore, we aim to enable these techniques to keep these structures live under removal of past events.

## References

1. D. Calvanese, T.E. Kalayci, M. Montali, and S. Tinella. Ontology-based Data Access for Extracting Event Logs from Legacy Data: The onprom Tool and Methodology. In *BIS 2017*, 2017.
2. D. Calvanese, M. Montali, A. Syamsiyah, and W.M.P. van der Aalst. Ontology-Driven Extraction of Event Logs from Relational Databases. In *BPI 2015*, pages 140–153, 2015.
3. C. Di Ciccio and M. Mecella. *Mining Constraints for Artful Processes*, pages 11–23. Springer Berlin Heidelberg, 2012.
4. C.W. Günther. XES Standard Definition. [www.xes-standard.org](http://www.xes-standard.org), 2014.
5. M.J. Jans, M. Alles, and M.A. Vasarhelyi. Process Mining of Event Logs in Auditing: Opportunities and Challenges. *Available at SSRN 2488737*, 2010.
6. S.J.J. Leemans, D. Fahland, and W.M.P. van der Aalst. Discovering Block-Structured Process Models from Event Logs - A Constructive Approach. In *Petri Nets 2013*, pages 311–329, 2013.
7. S.J.J. Leemans, D. Fahland, and W.M.P. van der Aalst. Discovering Block-Structured Process Models from Event Logs Containing Infrequent Behaviour. In *BPM Workshop 2013*, pages 66–78, 2013.
8. F.M. Maggi, A. Burattin, M. Cimitile, and A. Sperduti. *Online Process Discovery to Detect Concept Drifts in LTL-Based Declarative Process Models*, pages 94–111. Springer Berlin Heidelberg, 2013.

9. F. Mannhardt. XESLite Managing Large XES Event Logs in ProM. *BPM Center Report BPM-16-04*, 2016.
10. E. Rojas, J. Munoz-Gama, M. Sepúlveda, and D. Capurro. Process Mining in Healthcare: A Literature Review. *Journal of Biomedical Informatics*, 61:224–236, 2016.
11. S. Schönig, A. Rogge-Solti, C. Cabanillas, S. Jablonski, and J. Mendling. *Efficient and Customisable Declarative Process Mining with SQL*, pages 290–305. Springer International Publishing, Cham, 2016.
12. S. Suriadi, M.T. Wynn, C. Ouyang, A.H.M. ter Hofstede, and N.J. van Dijk. Understanding Process Behaviours in a Large Insurance Company in Australia: A Case Study. In *CAiSE 2013*, pages 449–464, 2013.
13. A. Syamsiyah, B.F. van Dongen, and W.M.P. van der Aalst. DB-XES: Enabling Process Mining in the Large. In *SIMPDA 2016*, pages 63–77, 2016.
14. A. Syamsiyah, B.F. van Dongen, and W.M.P. van der Aalst. Discovering Social Networks Instantly: Moving Process Mining Computations to the Database and Data Entry Time. In *BPMDS 2017*, 2017.
15. A. Syamsiyah, B.F. van Dongen, and W.M.P. van der Aalst. Recurrent Process Mining on Procedural and Declarative Approaches. *BPM Center Report BPM-17-03*, 2017.
16. W.M.P. van der Aalst. *Process Mining: Data Science in Action*. Springer, 2016.
17. B.F. van Dongen. BPI Challenge 2017, 2017.
18. B.F. van Dongen and S. Shabani. Relational XES: Data Management for Process Mining. In *CAiSE 2015*, pages 169–176, 2015.
19. S.J. van Zelst, A. Burattin, B.F. van Dongen, and H.M.W. Verbeek. Data Streams in ProM 6: A Single-Node Architecture. In *BPM Demo Session 2014*, page 81, 2014.
20. S.J. van Zelst, B.F. van Dongen, and W.M.P. van der Aalst. Know What You Stream: Generating Event Streams from CPN Models in ProM 6. In *BPM Demo Session 2015*, pages 85–89, 2015.
21. S.J. van Zelst, B.F. van Dongen, and W.M.P. van der Aalst. Online Discovery of Cooperative Structures in Business Processes. In *OTM Conferences 2016*, pages 210–228, 2016.
22. H.M.W. Verbeek, J.C.A.M. Buijs, B.F. van Dongen, and W.M.P. van der Aalst. XES, XESame, and ProM 6. In *Information Systems Evolution*, volume 72, pages 60–75, 2010.
23. Z. Zhou, Y. Wang, and L. Li. Process Mining Based Modeling and Analysis of Workflows in Clinical Care - A Case Study in a Chicago Outpatient Clinic. In *ICNSC 2014*, pages 590–595, 2014.