

Parallelism-based Session Creation to Identify High-Level Activities in Event Log Abstraction

Onur Dogan^{1,2} and Massimiliano de Leoni¹

¹ University of Padua, Department of Mathematics, 35121, Padua, Italy
onur.dogan@unipd.it, deleoni@math.unipd.it

² Izmir Bakircay University, Department of Management Information Systems, 35665, Izmir, Turkiye
onur.dogan@bakircay.edu.tr

Abstract. Process mining utilizes event data to gain insights into the execution of processes. While techniques are valuable, their effectiveness may be hindered when dealing with highly complex processes that have a vast number of variants. Additionally, because the recorded events in information systems are at a low-level, process mining techniques may not align with the higher-level concepts understood at the business level. Without abstracting event sequences to higher-level concepts, the outcomes of process mining, such as discovering a model, can become overly complex and challenging to interpret, rendering them less useful. Some research has been conducted on event abstraction, often requiring significant domain knowledge that may not be readily accessible. Alternatively, unsupervised abstraction techniques may yield less accurate results and rely on stronger assumptions. This paper introduces a technique that addresses the challenge of limited domain knowledge by utilizing a straightforward approach. The technique involves dividing traces into batch sessions, taking into account relationships between subsequent events. Each session is then abstracted as a single high-level activity execution. This abstraction process utilizes a combination of automatic clustering and visualization methods. The proposed technique was evaluated using a randomly generated process model with high variability. The results demonstrate the significant advantages of the proposed abstraction in effectively communicating accurate knowledge to stakeholders.

Keywords: Event log abstraction · Parallel relationships · Clustering · Visual analytics · Process discovery.

1 Introduction

In today's data-driven business landscape, organizations have access to vast amounts of event data generated by their information systems. This event data captures valuable insights into how processes are executed in the organization. Process mining techniques have emerged as powerful tools to analyze and leverage this event data to gain valuable insights into process execution. However, as processes become more complex and exhibit a large number of variants, the effectiveness of traditional process mining techniques is diminished [13]. Moreover, the low-level nature of the recorded events in information systems often fails to capture the higher-level concepts known at the business level, further complicating the analysis and interpretation of process mining results.

The need to abstract low-level event sequences into higher-level activities has been recognized as a crucial step in enhancing the usability and interpretability of process mining outcomes [10]. Traditional approaches to event abstraction often rely on significant domain knowledge, which may not be readily available or feasible to obtain. Thus, there is a pressing need to develop a technique that enables effective event abstraction without relying on extensive domain knowledge and with improved accuracy.

Existing approaches assume that each low-level sequence that is abstracted to one high-level activity does not interleave with other low-level sequences. In other words, there are no parallel high-level activities. This is a strong assumption because parallel activities are performed simultaneously, and the execution order changes in different traces.

This paper puts forward a technique that detects parallel low-level sequences by leveraging a combination of clustering and visualization methods. The proposed technique first identifies session creation rules by evaluating the time threshold and the parallel relationship between subsequent low-level events. Events having parallel relationships are assigned into different sessions to catch the parallelism. Then, it encodes the sessions as vectors, with each dimension representing an activity associated with a low-level event. A clustering algorithm groups the vectors of similar low-level sequences into clusters, where each cluster includes a set of parallel events. An abstracted model is generated using high-level event sequences. Then, each cluster, including sub-logs, is discovered for sub-models, and sub-models are replaced with the corresponding high-level activity on the abstracted model. After this replacement, the model's quality is measured with different parameter values. The parameter values giving the best quality are chosen to use for the evaluation of the abstraction technique.

The effectiveness of the proposed technique was evaluated through different synthetic processes and event logs. The results showed that the abstraction technique enables to mining process models with a better balance of fitness and precision.

2 Motivation

This section motivates the importance to consider the interleaving of low-level sequences. The discussion is based on the work by de Leoni et al. [7], but a similar reasoning applies when employing other techniques that overlook the parallelism of high-level activities [13]. The event-log abstraction technique in [7] divides traces into sessions, where each session represents the smallest subsequences within a trace. A session is determined based on a user-defined threshold, which is the time difference between the last event in a session and the first event in the subsequent session.

Let us consider the model in Figure 1 to motivate the importance of consider parallelism of high-level activities. The model consists of three parallel branches, each of which corresponds to one high-level activities. High-level activities A and B consists of the sequence of three low-level ones, in the model represented by the sequences $\langle A1, A2, A3 \rangle$, $\langle B1, B2, B3 \rangle$. Similarly, high-level activities C consists of the sequence of two low-level activities. The corresponding event log records the execution of these the low-level activities, and no information is stored about the high-level counterparts.




Fig. 1. An illustrative model including three parallel low-level sequences corresponding three different high-level activities




Fig. 2. Heatmap of clustering results for the technique in [7]. It highlights how each cluster has one or two evident dominant low-level activities (the cells of the darkest colors) according to the cluster centroids. There is only one cluster and the dominant low-level activities $D1$, $B3$, etc.

The timestamps of the low-level events are temporarily space out with a distance of one second from each other.

After applying the technique by de Leoni et al. [7] using a threshold of nine seconds, the outcome is one single cluster in which every low-level activity is added. This can be visualized through the heatmap in Figure 2, which illustrates the centroid of the created cluster. The most frequent low-level activities are given darker colours.

The expectation is to obtain four clusters, namely for the low-level activities related to the high-level A , B , C and D . However, since the low-level activities are given timestamps with distance of one second against a threshold session of nine seconds, the technique assigns every low-level activities into the same session, making it impossible to create the expected clusters. This can be solved by observing that low-level activities must not be in the same cluster if they can occur in any order, namely they are parallel. This idea motivates the research work reported in this paper. The remainder will detail the techniques. When parallelism is considered, we obtain the clusters depicted in Figure 3. The high-level activity A is covered via Cluster 0 where the most frequent low-level activities are indeed $A1$, $A2$ and $A3$. Clusters 1 and 3 incorporate the low-level activities related to the high-level activities C and B , respectively. Note the presence of clusters 2, 4 and 5, which are respectively related to high-level activities B , A and C : in these clusters, a predominant low-level activity is also $D1$, which indeed

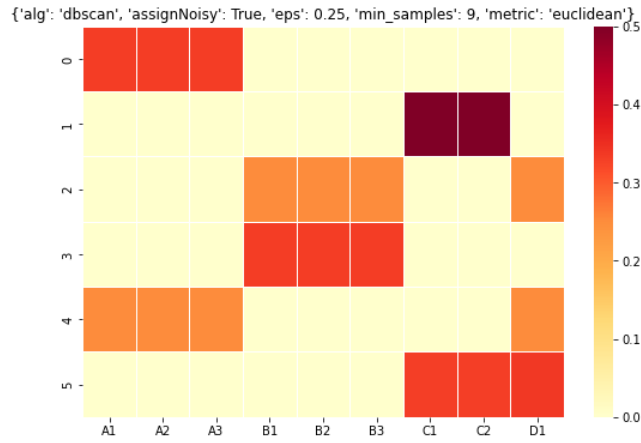


Fig. 3. Heatmap of clustering results for the proposed technique. Six clusters have different dominant low-level activities, and none of the clusters has low-level activities having a parallel relationship.

may follow $A3$, $B3$ or $C2$. In fact, the expected cluster for high-level activity D has been merged with those for A , B and C . The approach is thus not exactly achieving the four clusters, as expected, but yet is able to approximate the expected results.

3 Related Work

The problem of event-log abstract is currently gaining momentum, but most of the techniques do not consider the parallelism among high-level activities. Bakullari and van der Aalst [3] developed a framework for generating high-level events by dividing events into time windows. Van Eck et al. [11] presented an approach for creating high-level event logs from sensor data of smart products using segmentation and clustering techniques. Similarly, Brzychczy et al. [4] employed a combination of hierarchical clustering and domain knowledge to identify process stages. Michael Mayr et al. [9] extended the approach proposed by van Eck et al. [11] by incorporating time-series-based process-state and product-state data as additional attributes in the event log. Augusto et al. [1] applied the k -th order Markovian abstraction technique to create high-level activities and then discovered process models to compare the quality of the models based on a family of fitness and precision measures.

Some techniques support high-level activities parallelism, but they rely on manual effort and domain expertise. Mannhardt and Tax [8] suggested first discovering local process models from the event log to capture the behavior of specific subsets of events. The local process models were then used to lift the event log to a higher level of abstraction. Van Houdt et al. [12] employed a method that combines local process models with behavioral pattern mining to identify event abstraction patterns, leveraging on the event-log transformation by Mannhardt and Tax [8]. Baier et al. [2] leverage on domain knowledge to guide the event abstraction algorithm that also supports parallelism. By

incorporating domain knowledge, the algorithm achieves a more precise and contextually relevant event abstraction.

However, domain knowledge is often not readily available in real life. This paper presents a novel approach that considers parallelism, while not requiring domain knowledge.

4 Preliminaries

The initial step of an abstraction technique involves utilizing an *event log*, which comprises a collection of *traces*. Each trace represents a *process instance* (i.e., a *case*) manipulated in terms of *activities* performed and *process attributes* that are influenced.

Definition 1 (Events). Let A be the set of activity names, and T be the universe of timestamps, and I be the universe of event identifiers. An event is a triple $e = (a, t, i) \in A \times T \times I$ and represents that the event with identifier i refers to the execution of an instance of activity a for at timestamp t .

We use shortcuts $\lambda_A(e) = a$, $\lambda_T(e) = t$, $\lambda_I(e) = i$ to extract the activity name a , the timestamp t , and identifier i of any event e . Events may define several additional attributes, but they are not formalized here, to keep the notation simple. A trace refers to a sequence of events:

Definition 2 (Traces and Event Logs). A trace σ is a sequence of events, i.e., $\sigma \in \mathcal{E}^*$. An event log L is a set of traces, i.e., $\mathcal{L} \subset \mathcal{E}^*$.

The concept of parallelism used here is based on the dependency measured introduced within the Heuristic Miner [10]:

Definition 3 (Dependency Measure). Let A be the set of activity names, and T be the universe of timestamps, and I be the universe of event identifiers. Let $\mathcal{E} = A \times T \times I$ be the university of events. Let $\mathcal{L} \subset \mathcal{E}^*$ be an event log. For each pair $(a, b) \in A \times A$, let us denote $|a \succ_{\mathcal{L}} b|$ as the number of occurrences of events for activity a followed by events for b in \mathcal{L} . The dependency measure is defined as follows:

$$a \Rightarrow_{\mathcal{L}} b = \begin{cases} \frac{|a \succ_{\mathcal{L}} b| - |b \succ_{\mathcal{L}} a|}{|a \succ_{\mathcal{L}} b| + |b \succ_{\mathcal{L}} a| + 1} & \text{if } a \neq b \\ \frac{|a \succ_{\mathcal{L}} a|}{|a \succ_{\mathcal{L}} a| + 1} & \text{if } a = b \end{cases} \quad (1)$$

The value of $a \Rightarrow_{\mathcal{L}} b$ is always between -1 and 1. If the absolute value of $a \Rightarrow_{\mathcal{L}} b$ is lower than an user-defined threshold, activities a and b may be considered in parallel.

5 Abstraction Technique

Figure 4 depicts the steps of the proposed abstraction technique. The input consists of an event log, which is split into training and testing. The training event log is used to create sessions at step 1: the events of each trace of the event log are divided into sessions. A session is a sequence of events. At step 2, sessions are clustered, where each cluster




Fig. 4. Abstraction technique framework

represents a high-level activities. The sessions of each cluster are the sequence of low-level activities that are mapped to the high-level activity corresponding to that cluster. At step 3, the low-level event sequences are replaced by high-level events to obtain an abstracted event log: this abstracted event log can then be used, e.g., to mine a high-level model. The sessions of each cluster can be seen as a sublog that can be used as input for process discovery: this would discover a low-level process model of how each high-level activity is mapped to the executions of low-level activities. The low-level process models can then be merged with the abstract process model to in fact obtain a hierarchical model: this is step 4. At step 5, a testing event log is used to compute the harmonic mean of fitness and precision, and thus to assess the quality of the hierarchical model.

Very importantly, this procedure is repeated for different parameters of the clustering algorithm to tune the parameters and find the values that can guarantee higher quality of the hierarchical model.

Note that the technique is independent of the specific clustering algorithm, but we had to pick up, and we opted for the DBSCAN clustering algorithm in the implementation and experiments.

The remainder of this section provides further details of the different steps of the proposed techniques.

5.1 Parallelism-based Session Creation

Algorithm 1 illustrates how sessions are created. The input is the event log \mathcal{L} , a time threshold Δ and the dependency threshold Θ . In particular, Θ is used to set the minimum strength for a dependency $a \Rightarrow b$ to consider a and b not to be parallel: the ideal value would be one to indicate that a and b are parallel if even once it occurs a is followed by b and b by a . The reality is that the observed behavior can be incomplete and noisy, and hence a certain tolerance should be accepted. Values observed in [10] are usually between 0.7 and 0.9.

Algorithm 1: Session Creation based on Parallelism

Data: An event log \mathcal{L}
Data: A session time-threshold Δ , A dependency threshold Θ
Result: A set S of sessions

```

1  $S \leftarrow \emptyset$ 
2 foreach trace  $\sigma = \langle e_1, \dots, e_n \rangle \in \mathcal{L}$  do
3   Trace-session set  $S_\sigma \leftarrow \{\langle e_1 \rangle\}$ 
4   for event index  $i=2$  to  $n$  do
5      $event\_assigned \leftarrow False$ 
6     foreach session  $s = \langle f_1, \dots, f_m \rangle \in S_\sigma$  do
7       if  $(|\lambda_A(f_m) \Rightarrow_{\mathcal{L}} \lambda_A(e_i)| \geq \Theta)$  and  $(\lambda_T(e_i) - \lambda_T(f_m) < \Delta)$  then
8          $S_\sigma \leftarrow (S_\sigma \setminus s) \cup \{\langle f_1, \dots, f_m, e_i \rangle\}$ 
9          $event\_assigned \leftarrow True$ 
10        break
11      end
12    end
13    if  $event\_assigned = False$  then
14       $S_\sigma \leftarrow S_\sigma \cup \{\langle e_i \rangle\}$ 
15    end
16  end
17   $S \leftarrow S \cup S_\sigma$ 
18 end

```

As illustrated in Algorithm 1, the creation is done by trace (see line 2). For each trace $\sigma \in \mathcal{L}$, a set S_σ of the sessions of σ is created to include a session that consists of the sequence with only the first event (see line 3). In fact, a session for a trace σ is a sequence of events of σ . The loop from line 4 to 16 iterates over all events of the trace of length n , and adds each event e_i either to an existing session or to a new session. The inner loop from line 6 to 12 iterates over the existing sessions: if any session s concludes with an event f_m that is not parallel with e_i (i.e. the dependency metrics between the activities of events e_i and f_m is larger than Θ) and also the time difference between the two events is lower than the time threshold Δ (conditions at line 7), e_i is added to f_m (see line 8) and no more sessions are considered (see break of the inner loop at line 10). A special case is when the event is not added to any of the existing session because the conditions at line 7 are never both satisfied. In this case, the inner loop between 6 and 12 concludes. The boolean variable $event_assigned$ stays set to false: it would become true if and only if it is added to some session (cf. line 9 where the variable is set to true). The condition at line 13 refers to this case: if the event is assigned to no session, a new session is created and added to the existing set S_σ .

At the end of this step, every trace is broken down into the constituent sessions, which are then clustering as discussed in the next step.

5.2 Clustering

This study chooses the DBSCAN to cluster sessions, as previous studies [6, 7] showed that it yielded better results. Because the DBSCAN clustering method works with numerical variables, the sessions are required to be converted into an encoded numerical vector. The encoding creates vectors $\mathbb{S}_\sigma = (v_{x_1}, \dots, v_{x_n})$ with one dimension for each activity associated with low-level events. The set of activities for low-level events in the log is indicated as $\{x_1, \dots, x_n\}$. An event e exists in some log traces such as $\lambda_A(e) = x_i$ for each $i \in [1, n]$. The proposed algorithm implements a frequency-based algorithm. Dimension v_{x_k} indicates the frequency of events for activity x_k in \mathbb{S}_σ . In other words, the algorithm creates a one-dimensional vector sizing the number of unique events in the log. Then, it assigns 1 to the vector for the corresponding element and 0 for the others. When the event log includes start and end timestamps for an event, the value of v_{x_k} should be divided by two to avoid double counting.

Clusters are visualized on the heatmap for insights into how sessions were clustered. The x-axis refers to low-level activities, whereas the y-axis indicates clusters. The intersection of the axes is colored according to the value of the dimension a in cluster c . A value of 0 is represented as white, while 1 is given as the most intense red. The cluster centroids (c_1, \dots, c_k) are normalized by dividing the sum of the centroid values $(\frac{c_1}{\sum_{j=1}^k c_j}, \dots, \frac{c_k}{\sum_{j=1}^k c_k})$. Varying intensities of the color red represent the centroids on the heat maps.

5.3 Creation of Abstracted Model

'Step 2 Clustering' creates a set of clusters $C = \{C_1, \dots, C_k\}$ of sessions. For each trace σ , the sequence of the sessions $S_\sigma = \langle s_1, \dots, s_m \rangle$ are computed. Then, low-level activities are converted into high-level activities by abstracting traces $\sigma_H = \langle h_{\sigma_1}^s, h_{\sigma_1}^c, \dots, h_{\sigma_k}^s, h_{\sigma_k}^c \rangle$, where $h_{\sigma_i}^s$ and $h_{\sigma_i}^c$ are the start and complete, respectively, high-level (abstracted) events for the trace σ_i in the session s . Each cluster $c_j \in C$ is renamed as $NAME(c_j)$ in the abstracted event log by considering the most dominant activity in the cluster c_j .

5.4 Merging Sub-cluster Models

Merging sub-cluster models is required to create a hierarchical model. This step extracts the sub-processes associated with each high-level activity in the abstracted model and merges them to compare the original and abstracted models. Each cluster $c_j \in C$ represents an event log that serves as input for mining a model that captures the behavior of the sessions within that cluster.

Each high-level activity $NAME(c_j)$ is replaced by the model discovered from the low-level activities that comprise the corresponding high-level activity $NAME(c_j)$. The merging process is made by putting two invisible transitions before and after the model discovered from the low-level activities.

5.5 Computing Model Quality

Quality computing for a process model involves evaluating the discovered process model's accuracy, reliability, and usefulness. Fitness reflects how well the model aligns with the actual execution of the process, with higher fitness values indicating a closer match between the observed and modeled behavior [10]. Precision evaluates how effectively the model represents the actual process without oversimplifying or omitting essential details [10]. This study calculates the harmonic means of fitness and precision, particularly useful in situations where there are extreme values or outliers present in the data. Because models requiring to be abstracted have mainly low precision values, the harmonic mean balances the influence of both precision and fitness metrics.

5.6 Choosing Best Model

Choosing the best model is an iterative step that requires testing different parameter values. The time threshold was computed as the average time difference between two subsequent events in each trace. A time threshold set containing time threshold values up to three times the average time difference was created. ϵ set included 0.25, 0.5 and 0.75. *min_samples* values were computed by considering the number of activities in the event log. A *min_samples* set containing values up to three times the unique number of activities in the log was created. The combination of these three parameter values was tested in each experiment.

6 Evaluation

To evaluate the effectiveness of the proposed technique, a comprehensive assessment was conducted on a randomly generated process model characterized by high variability. Processes and Logs Generator 2 (PLG2) was used to create a model in this study [5]. PLG2 allows to generate random process models and to create accordant event logs. It also allows noise to be added to the event log, namely events can be randomly inserted, deleted or swapped, at the beginning, in the middle or at the end of the log's traces. A user-defined parameter enables to define the amount of noise (see [5] for further details). An implementation in Python, and the generated models and logs used in the evaluation are available at <https://github.com/onurdogannet/ParallelismbasedEventLogAbstraction>.

Because PLG2 generates a log with fixed 1-hour timestamps for each consecutive activity, we changed the timestamps randomly, considering relationships among activities. Then, the original models were discovered by Inductive Miner - infrequent (IMf) with a noise threshold of 0.2. Model 1 includes 109 activities. The average number of events per trace is 107.41, and the average trace duration is 18.9 hours. The average time between two subsequent activities is 17.34 minutes, which is considered to define the time threshold, Δ . Model 2 has 118 activities. The average number of events per trace is 105.01, and the average trace duration is 17.96 hours. The average time between two subsequent activities is 10.54 minutes. The logs with noises were generated by adding levels to the control flow of events. For example, a new log was generated with the following specifications related to the control flow of events: trace missing




Fig. 5. Quality of Model 1 in terms of the harmonic mean of fitness and precision




Fig. 6. Quality of Model 2 in terms of the harmonic mean of fitness and precision

head probability, trace missing tail probability, and trace missing episode probability. For two models, new logs with noises were generated with the noise probabilities of 5%, 10%, 15%, and 20% for the same specifications.

Figure 5 and Figure 6 provide an evaluation of process models based on their fitness, precision, and harmonic mean scores for the event log. The models include the baseline and abstractions with consecutiveness and parallelism. The original models show high fitness scores but low precision. On the other hand, abstraction techniques based on consecutiveness and parallelism result in lower fitness scores but improved precision, in general.

Across different session creation methods (Original, Consecutiveness, Parallelism), there are variations in fitness. For example, in Model 1, the Parallelism variation tends to have lower fitness compared to the Original or Consecutiveness, especially at higher noise levels. In Model 2, the fitness of the Consecutiveness variation varies significantly

with noise levels. It's important to note that Model 1 generally exhibits higher fitness scores compared to Model 2 under the same conditions and noise levels.

Precision varies across different noise levels and model names. In Model 1, the Parallelism variation tends to have higher precision compared to the other two variations across noise levels. In Model 2, the precision of the Parallelism variation is notably high, especially at 10% noise level, suggesting it is better at avoiding false positives under certain conditions.

The parallelism-based abstraction achieves a better balance between fitness and precision compared to the consecutiveness-based abstraction for both event logs, except the event logs with 20% noise. Because both abstraction methods consider noises by assigning them to the nearest cluster before creating the abstracted model, they gave better results until a certain noise level. However, after a certain level of noise, abstraction techniques may struggle to cope with the increased variation introduced by the noise, leading to reduced model quality. Intuitively, as the noise level continues to increase beyond 20%, the model quality is expected to deteriorate further. At 25% or 30% noise levels, the proportion of irrelevant events becomes even larger, making it more challenging for the abstraction techniques to identify relevant patterns and relationships. This can lead to a higher number of incorrect or misleading abstractions, resulting in a decline in both fitness and precision.

7 Conclusion and Future Research

This paper addresses the challenges traditional process mining techniques face in dealing with complex processes and low-level event data. By introducing a novel event abstraction technique that requires limited domain knowledge, the paper offers an approach to enhance the interpretability and usability of process mining outcomes. The proposed technique's effectiveness is validated through different noise levels, emphasizing its ability to provide accurate knowledge to stakeholders. The improved interpretability and usability of the abstracted process model enable users to gain valuable insights and make informed decisions regarding process analysis, optimization, and improvement.

For noise levels below 20%, the abstraction techniques perform better than the not abstracted model because they can effectively filter out some noise and separate the essential patterns, resulting in improved model quality compared to the raw, unprocessed data. The fact that parallelism is the best abstraction technique at noise levels less than 20% can be attributed to its ability to handle more diverse and complex event patterns. Parallelism captures relationships among events that occur simultaneously, which can be especially valuable when the noise level is low, and the patterns are more distinct. The choice of the best abstraction technique depends on the specific characteristics and noise level of the event log being analyzed.

For future work, we aim to apply different clustering techniques, such as hierarchical clustering, to obtain better clustering results, affecting the abstracted model's quality. Also, parameter values of the clustering technique and time threshold for session creation should be found algorithmically.

Acknowledgements. This research is financially supported by the Department of Mathematics of the University of Padua, through the BIRD projects “Data-driven Business Process Improvement” (code BIRD215924/21) and “Web-site Interaction Discovery” (code BIRD219730/21).

References

1. Augusto, A., Armas-Cervantes, A., Conforti, R., Dumas, M., La Rosa, M.: Measuring fitness and precision of automatically discovered process models: a principled and scalable approach. *IEEE Transactions on Knowledge and Data Engineering* **34**(4), 1870–1888 (2020)
2. Baier, T., Mendling, J., Weske, M.: Bridging abstraction layers in process mining. *Information Systems* **46**, 123–139 (2014)
3. Bakullari, B., van der Aalst, W.M.: High-level event mining: A framework. In: 2022 4th International Conference on Process Mining (ICPM). pp. 136–143. IEEE (2022)
4. Brzywczy, E., Trzcionkowska, A.: Process-oriented approach for analysis of sensor data from longwall monitoring system. In: *Intelligent Systems in Production Engineering and Maintenance*. pp. 611–621. Springer (2019)
5. Burattin, A.: Plg2: Multiperspective process randomization with online and offline simulations. In: *BPM (Demos)*. pp. 1–6. Citeseer (2016)
6. de Leoni, M., Dündar, S.: Event-log abstraction using batch session identification and clustering. In: *Proceedings of the 35th Annual ACM Symposium on Applied Computing*. pp. 36–44 (2020)
7. de Leoni, M., Pellattiero, L.: The benefits of sensor-measurement aggregation in discovering iot process models: a smart-house case study. In: *International Conference on Business Process Management*. pp. 403–415. Springer (2021)
8. Mannhardt, F., Tax, N.: Unsupervised event abstraction using pattern abstraction and local process models. *arXiv preprint arXiv:1704.03520* (2017)
9. Mayr, M., Luftensteiner, S., Chasparis, G.C.: Abstracting process mining event logs from process-state data to monitor control-flow of industrial manufacturing processes. *Procedia Computer Science* **200**, 1442–1450 (2022)
10. Van Der Aalst, W.M.P.: *Data science in action*. Springer (2016)
11. Van Eck, M.L., Sidorova, N., Van der Aalst, W.M.: Enabling process mining on sensor data from smart products. In: 2016 IEEE Tenth International Conference on Research Challenges in Information Science (RCIS). pp. 1–12. IEEE (2016)
12. Van Houdt, G., Depaire, B., Martin, N.: Unsupervised event abstraction in a process mining context: A benchmark study. In: *Process Mining Workshops: ICPM 2020 International Workshops, Padua, Italy, October 5–8, 2020, Revised Selected Papers 2*. pp. 82–93. Springer (2021)
13. van Zelst, S.J., Mannhardt, F., de Leoni, M., Koschmider, A.: Event abstraction in process mining: literature review and taxonomy. *Granular Computing* **6**, 719–736 (2021)