<https://helda.helsinki.fi>

Exploiting Event Log Event Attributes in RNN Based Prediction

Hinkka, Markku

Springer Nature Switzerland
2019-09-01

Hinkka , M , Lehto , T & Heljanko , K 2019 , Exploiting Event Log Event Attributes in RNN
pö Based Prediction . in T Welzer , J Eder , V Podgorelec , R Wrembel , I
pö , M Morzy , T Tzouramanis , J Darmont & A Kamiaali Latifi (eds) , Ne
Databases and Information Systems : ADBIS 2019 Short Papers, Workshops BBIGAP,
QAUCA, SemBDM, SIMPDA, M2P, MADEISD, and Doctoral Consortium, Bled, Slovenia,
pö September 8 11, 2019, Proceedings . Communications in Computer and
Science , vol. 1064 , Springer Nature Switzerland , Cham , pp. 405-416 , European
Conference on Advances in Databases and Information Systems , Bled , Slovenia ,
08/09/2019 . https://doi.org/10.1007/978-3-030-30278-8_40

<http://hdl.handle.net/10138/339400>

https://doi.org/10.1007/978-3-030-30278-8_40

cc_by_nd

acceptedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Exploiting Event Log Event Attributes in RNN Based Prediction

Markku Hinkka^{1,3}[0000-0002-3679-3677], Teemu Lehto^{1,3}[0000-0002-1332-1853],
and Keijo Heljanko^{2,4}[0000-0002-4547-2701]

¹ Aalto University, School of Science, Department of Computer Science, Finland

² University of Helsinki, Department of Computer Science, Finland,

³ QPR Software Plc, Finland

⁴ HIIT Helsinki Institute for Information Technology

markku.hinkka@aalto.fi, teemu.lehto@qpr.com, keijo.heljanko@helsinki.fi

Abstract. In predictive process analytics, current and historical process data in event logs is used to predict the future, e.g., to predict the next activity or how long a process will still require to complete. Recurrent neural networks (RNN) and its subclasses have been demonstrated to be well suited for creating prediction models. Thus far, event attributes have not been fully utilized in these models. The biggest challenge in exploiting them in prediction models is the potentially large amount of event attributes and attribute values. We present a novel clustering technique that allows for trade-offs between prediction accuracy and the time needed for model training and prediction. As an additional finding, we also find that this clustering method combined with having raw event attribute values in some cases provides even better prediction accuracy at the cost of additional time required for training and prediction.

Keywords: process mining, predictive process analytics, prediction, recurrent neural networks, gated recurrent unit

1 Introduction

Event logs generated by systems in business processes are used in Process Mining to automatically build real-life process definitions and as-is models behind those event logs. There is a growing number of applications for predicting the properties of newly added event log cases, or process instances, based on case data imported earlier into the system [4][5][13][19]. The more the users start to understand their processes, the more they want to optimize them. This optimization can be facilitated by performing predictions. To be able to predict properties of new and ongoing cases, as much information as possible should be collected that is related to the event log traces and relevant to the properties to be predicted. Based on this information, a model of the system creating the event logs can be created. In our approach, the model creation is performed using supervised machine learning techniques.

In our previous work [9] we have explored the possibility to use machine learning techniques for classification and root cause analysis for a process mining-related classification task. In the paper, experiments were performed on the efficiency of several feature selection techniques and sets of structural features (a.k.a. activity patterns) based on process paths in process mining models in the context of a classification task. One of the biggest problems with the approach is the finding of the structural features having the most impact on the classification result. E.g., whether to use only activity occurrences, transitions between two activities, activity orders, or other even more complicated types of structural features such as detecting subprocesses or repeats. For this purpose, we have proposed another approach in [10], where we have examined the use of recurrent neural network techniques for classification and prediction. These techniques are capable of automatically learning more complicated causal relationships between activity occurrences in *activity sequences*. We have evaluated several different approaches and parameters for the recurrent neural network techniques and have compared the results with the results we collected in our work. In both the previous publications [9][10], focusing on boolean -type classification tasks based on the *activity sequences* only.

In this work we build on our previous work to further improve the prediction accuracy of prediction models by exploiting additional event attributes that are often available in the event logs while also taking into account the scalability of the approach to allow users to precisely specify the event attribute detail level suitable for the prediction task ahead. Our goal is to develop a technique that would allow the creation of a tool that is, based on a relatively simple set of parameters and training data, able to efficiently produce a prediction model for any case-level prediction task, such as predicting the next activity or the final duration of a running case. Fast model rebuilding is also required in order for a tool to be able to also support, e.g., interactive event and case filtering capabilities. Thus, the performance of the system under study is measured using four different metrics: Success rate, input vector length that gives a rough indication of the memory usage, the time required for training a model and the time required for performing predictions using the already trained model.

To answer these requirements, we introduce a novel method of exploiting event attributes into RNN prediction models by clustering events by their event attribute values and using the cluster labels in the RNN input vectors instead of the raw event data. This makes it easy to manage the input RNN vector size no matter how many event attributes there are in the data set. E.g., users can configure the absolute maximum length of the one-hot vector used for the event attribute data which will not be exceeded, no matter how many actual attributes the dataset has. RNN is an ideal choice for process mining-related prediction tasks since it learns temporal dynamic behavior by using its internal state to process sequences of inputs. Since predictions are usually made based on sequences of events, this makes it a more natural machine learning technique in process mining context than more traditional approaches, such as random forest and gradient boosting.

Our prediction engine source code is available in GitHub ⁵.

The rest of this paper is structured as follows: Section 2 is a summary of the latest developments around the subject. In Section 3, we present the problem statement and the related concepts. Section 4 presents our solution for the problem. In Section 5 we present our test framework used to test our solution. Section 6 describes the used datasets as well as performed prediction scenarios. Section 7 presents the experiments and their results validating our solution. Finally Section 8 draws the final conclusions.

2 Related Work

Lately, there has been a lot of interest in the academic world on predictive process monitoring which can clearly be seen, e.g., in [6] where the authors have collected a survey of 55 accepted academic papers on the subject. In [18], the authors have compared several approaches spanning three different research fields: Machine learning, process mining and grammar inference. As a result, they have found that overall, the techniques from machine learning field generate more accurate predictions than grammar inference and process mining fields.

In [19] the authors used Long Short-Term Memory (LSTM) recurrent neural networks to predict the next activity and its timestamp. They use one-hot encoded activity labels and three numerical time-based features: duration between the current activity and the previous activity, time within the day and time within the week. Event attributes were not considered at all. In [4] the authors trained LSTM networks to predict the next activity. In this case, however, network inputs are created by concatenating categorical, character string-valued event attributes and then encoding these attributes via an embedding space. They also note that this approach is feasible only because of the small number of unique values each attribute had in their test datasets. Similarly, in [17], the authors take a very similar approach based on LSTM networks, but this time also incorporate both discrete and continuous event attribute values. Discrete values are one-hot encoded, whereas continuous values are normalized using min-max normalization and added to the input vectors as single values.

In [14] the authors use Gated Recurrent Unit (GRU) recurrent neural networks to detect *anomalies* in event logs. One one-hot encoded vector is created for activity labels and one for each of the included string-valued event attributes. These vectors are then concatenated in a similar fashion to our solution into one vector representing one event, which is then given as input to the network. We use this approach for benchmarking our own clustering-based approach (labeled as *Raw* feature in the text below). The system proposed in their paper is able to predict both the next activity and the next values of event attributes. Specifically, it does not take case attributes and temporal attributes into account.

In [21] the authors train a RNN to predict the most likely future activity sequence of a running process based only on the sequence of activity labels.

⁵ <https://github.com/mhinkka/articles>

Similarly our earlier publication [9] used sequences of activity labels to train a LSTM network to perform a boolean classification of cases.

Also, process mining models obtained using process mining techniques themselves can be used as a model for prediction. In [1] the authors first generate a process tree using the inductive miner algorithm, after which this process tree is used to predict the future path of ongoing cases. This approach does not take any additional event- or case attributes into account.

None of the mentioned earlier works present a solution that is scalable for datasets having lots of event- or case attributes and unique attribute values.

3 Problem

Using RNN to perform case-level predictions on event logs has lately been studied a lot. However, there has not been any scalable approach to handling event attributes in the RNN setting. Instead, e.g., in [14] authors used separate one-hot encoded vector for each attribute value. Having this kind of an approach when you have, e.g., 10 different attributes, each having 10 unique values would already require a vector of 100 elements to be added as input for every event. The longer the input vectors become, the more time and memory it gets for the model to create accurate models from them. This increases the time and memory required to use the model for predictions.

4 Solution

Since in addition to having activity labels in the input vectors, we need to store also event attribute-related information, we decided to include several feature types into the input vectors of the RNN. Input vectors are formatted as shown in Table 1, where each column represents one feature vector element f_{ab} , where a is the index of the feature and b is the index of the element of that feature. In the table, n represents the number of feature types used in the feature vector and m_k represents the number of elements required in the input vector for feature type k . Thus, each feature type produces one or more numeric elements into the input vector, which are then concatenated together into one actual input vector passed to RNN both in training and in prediction phases. Table 2 shows an example input vector having three different feature types: activity label, raw event attribute values (only single event attribute named *food* having four unique values) and the event attribute cluster where clustering has been performed separately for each unique activity.

For this paper, we encoded only event activity labels and event attributes into the input vectors. However, this mechanism can easily incorporate also other types of features not described here. The only requirement for added features is that it needs to be able to be encoded into a numeric vector as shown in Table 1 whose length must be the same for each event.

Table 1: Feature input vector structure

f_{11}	f_{12}	...	f_{1m_1}	f_{21}	...	f_{2m_2}	...	f_{n1}	...	f_{nm_n}
----------	----------	-----	------------	----------	-----	------------	-----	----------	-----	------------

Table 2: Feature input vector example content

row	$activity_{eat}$	$activity_{drink}$	$food_{salad}$	$food_{pizza}$	$food_{water}$	$food_{soda}$	$cluster_1$	$cluster_2$
1	1	0	1	0	0	0	1	0
2	0	1	0	0	1	0	1	0
3	1	0	0	1	0	0	0	1
4	1	0	0	1	0	0	0	1
5	0	1	0	0	0	1	0	1

4.1 Event Attributes

Our primary solution for incorporating information in event attributes into input vectors is to cluster all the event attribute values in the training set and then use a one-hot encoded cluster identifier to represent all the attribute values of the element. The used clustering algorithm must be such that it tries to automatically find the optimal number of clusters for the given data set within the range of 0 to N clusters, where N can be configured by the user. By changing N, the user can easily configure the maximum length of the one-hot -vector as well as the precision of how detailed attribute information will be tracked. For this paper, we experimented with a slightly modified version of Xmeans -algorithm [15]. Another option could have been a method where silhouette scoring [16] is used to determine the optimal number of clusters for k-means [8], but based on our tests, this approach did not perform fast enough to be applied as our selected approach.

It is very common that different activities get processed by different resources yielding a completely different set of possible attribute values. E.g., different departments in a hospital have different people, materials and processes. Also, in the example feature vector shown in Table 2, *food* -event attribute has completely different set of possible values depending on the *activity* since it is forbidden by, e.g., the external system to not allow activity of type *eat* to have *food* event attribute value of *water*. If we cluster all the event attributes using single clustering, we would easily lose this activity-type-specific information.

In order to retain this activity-specific information, we used separate clustering for each unique activity type. All the event attribute clusters are encoded into one one-hot encoded vector representing only the resulting cluster label for that event, no matter what its activity is. This is shown in the example table as $cluster_N$, which represents the row having N as a clustering label. E.g., in the example case, $cluster_1$ is 1 in both rows 1 and 2. However, row 1 is in that cluster because it is in the 1st cluster of the $activity_{eat}$ activity, whereas row 2 is in that cluster because it is in the 1st cluster of the $activity_{drink}$ activity. Thus, in order to identify the actual cluster, one would require both the activity label and the cluster label. For RNN to be able to properly learn about the actual

event attribute values, it needs to be given both the activity label and the cluster label in the input vector. Below, this approach is labeled as *ClustN*, where N is the maximum cluster count.

For benchmarking, we also experimented with a *raw* implementation where event attributes were used so that every event attribute is encoded into its own one-hot encoded vector and then concatenated into the actual input vectors. This method is lossless since every unique event attribute value has its own indicator in the input vector. Below, this approach is referred to as *Raw*. Finally, we experimented also using both *Raw* and *Clustered* event attribute values. Below, this approach is referred to as *BothN*, where N is the maximum cluster count.

4.2 Formal problem definition

Basically, the problem we are solving in this paper is that we encode as much information as possible from event log event attributes into a single numeric RNN input vector whose length is user-configurable. In this section, we will present the formal definitions required to describe our clustering-based approach for solving this problem. We will build our formal definitions on the basis of the definitions given in Process Mining [20] -book Chapter 5 describing the basic concepts in process mining.

First, we define relationships between event log and events as follows:

Definition 1. *Event log L is a set of cases $c_i \in L$. Each case c_i is a sequence of events $c_i = \langle e_1, e_2, \dots, e_n \rangle$. We denote $e \in L$ iff $\exists c_i = \langle e_1, e_2, \dots, e_n \rangle \in L$, such that $e_j = e$, for some $1 \leq j \leq n$.*

Next, we define a function for accessing event attributes of an event:

Definition 2. *Let $\#_n : e \mapsto v$, where $(n \in AN) \wedge (e \in L)$, and AN is the finite set of all the attribute labels for the events in the given event log. v is the value of that attribute and can be of any arbitrary type.*

Using this function, we can refer to any of the *standard attributes* (in this paper, only *activity*, *time* and *caseid* are considered as standard attributes), as well as event log-specific attributes. For this paper, we need to define these attribute label sets as follows:

Definition 3. *Let $AN_{std} \subseteq AN$, be the set of standard attribute labels: $\langle activity, time, caseid \rangle$. Also, let $(AN_L \cap AN_{std} = \emptyset) \wedge (AN_{std} \cup AN_L = AN)$, where AN_L is the set of event attributes in event log L , which are not part of standard attributes.*

The standard attributes listed above have the following meanings:

Definition 4. *$\#_{activity}(e)$ is the activity label associated with the event e . This describes what has occurred. Activity labels in this paper are considered to be textual descriptions of the performed task.*

$\#_{time}(e)$ is the timestamp of the event e . This describes when something has

occurred. In this paper, timestamps include both time and date of the occurred event.

$\#_{caseid}(e)$ is the identifier of the case associated with the event e . This describes the (case) context of the occurrence. Case identifiers in this paper are considered to be short textual or numeric identifiers identifying the case.

Using these, we can formally define a set of all the activity labels in the event log as follows:

Definition 5. Let $\mathcal{A}_L$ be the set of activity labels so that $\forall e \in L, \#_{activity}(e) \in \mathcal{A}_L$.

Next, we split the event log into two disjoint sets: Training set and test set. Formally this can be expressed as:

Definition 6. Training set is $L_{tr} \subset L$, where $L_{tr} \neq \emptyset$. Similarly, test set is $L_t \subset L$, where $(L_t \neq \emptyset) \wedge (L_t \cap L_{tr} = \emptyset) \wedge (L_t \cup L_{tr} = L)$.

We also formally define separate subsets of activity labels for both the test and the training set as follows:

Definition 7. $\mathcal{A}_{tr} \subseteq \mathcal{A}_L$ is used to denote all the activity labels available in the training set, whereas $\mathcal{A}_t \subseteq \mathcal{A}_L$ denotes those in the test set.

Similarly, we denote sets of available attributes in both the training set and the test set as follows:

Definition 8. $AN_{tr} \subseteq AN_L$ is used to denote all the attribute names available in the training set, whereas $AN_t \subseteq AN_L$ denotes those in the test set.

Next, we define a function used to concatenate multiple vectors to each other.

Definition 9. Let $concat : \langle X_0, \dots, X_n \rangle \mapsto Y$ be a function that returns a single numeric vector Y consisting of the concatenated contents of numeric vectors $X_0, \dots, X_n$ in the specified order.

Then we define a function that maps each unique value into unique integer value as follows:

Definition 10. Let $codify : x \times X \mapsto y$, where X is a finite set of all the possible values for $x \in X$, and $y \in 1, \dots, |X|$ is the value x being mapped bijectively into an integer. Thus, $codify$ creates a bijection between values and an integer representation of that value.

Now we can give a formal definition for the one-hot encoding function as follows:

Table 3: One-hot encoding example

Original	After <i>codify</i>	After <i>onehot</i>
$\langle a, b, c, d \rangle$	$\langle 1, 2, 3, 4 \rangle$	$\langle \langle 1, 0, 0, 0 \rangle, \langle 0, 1, 0, 0 \rangle, \langle 0, 0, 1, 0 \rangle, \langle 0, 0, 0, 1 \rangle \rangle$

Definition 11. Let $onehot : x \times X \mapsto Y$, where X is a finite set of all the possible values and $x \in X$ is the value being encoded, be an *onehot* encoding function that transforms x into a numeric vector Y of length $|X|$. Vector $Y = \langle y_1, \dots, y_{|X|} \rangle$, where

$$y_k = \begin{cases} 1, & \text{iff } k = codify(x, X) \\ 0, & \text{otherwise} \end{cases}$$

Thus, every unique value in X returns an unique numeric vector.

An example of a one-hot encoding process is shown in Table 3, where the set of all possible values X is $\langle a, b, c, d \rangle$, and we are applying *onehot* function to each of the four possible values separately.

Next, since we need to be able to iterate through all the attributes, we need to specify an unambiguous way to map iteration index to an attribute name. For this purpose we define the following function:

Definition 12. Let $attname : n \mapsto an$, where $an \in AN_L$, and $n \in 1, \dots, |AN_L|$ be a bijective function for mapping an positive integer iteration index to an attribute name.

Using the previous definition, we can refer to n th event attribute of e by writing $\#_{attname(n)}(e)$. Next, we define a method for creating subsets of events in a way that each subset will have all the events having one specific activity label. Formally we express these sets as:

Definition 13. Let B_{act} , where $act \in \mathcal{A}_{tr}$ be the set of events $e \in L$ such that $\#_{activity}(e) = act$.

Next, we will define a function for retrieving a set of all attribute values of a given set of events:

Definition 14. Let $\#_n : \langle e_1, \dots, e_n \rangle \mapsto Y$, where $e_i \in L \forall (1 \leq i \leq n)$. This function returns a set of all the attribute values of attribute having index n for given events.

Using these definitions, the set of all the attribute values of all the events having a specific activity label act can be referred to using $\#_{attname(n)}(B_{act})$, which yields a set of attribute values of attribute having index n for all the events having activity label act . Next, In order to perform clustering for events in activity buckets, we need to first one-hot encode event attribute values.

Definition 15. Let $onehot_{B_{act}} : e \times n \mapsto Y$, where $e \in L$, $act \in \mathcal{A}_{tr}$, and $n \in 1, \dots, |AN_L|$, be a function that performs one-hot encoding for attribute having attribute iteration index of n for event e . Thus,
 $Y = onehot(\#_{attname(n)}(e), \#_{attname(n)}(B_{\#_{activity}(e)}))$.

Using this definition, we can transform all event attribute values into a single numeric vector using the following additional function definition.

Definition 16. *One-hot encoding function for all event attributes of given event:*

$$\begin{aligned} \text{onehot}_{\#} : e \mapsto & \text{concat}(\text{onehot}_{B_{\# \text{activity}(e)}}(e, 1), \\ & \dots, \\ & \text{onehot}_{B_{\# \text{activity}(e)}}(e, |AN_L|)) \end{aligned}$$

, where $e \in L$.

Next, we define the actual clustering function that uses the number vectors created by $\text{onehot}_{\#}$ function as follows.

Definition 17. $\#_{\text{cluster}} : e \times \max_{cc} \mapsto y$, where $e \in L_{tr}$, and $\max_{cc}$ denotes the configured maximum cluster count, and $1 \leq y \leq \max_{cc}$, and $y \in \mathbb{Z}$, denotes the assigned cluster label among the set of all the possible cluster labels. Clustering is performed separately for each activity label B_{act} , where $act \in \mathcal{A}_{tr}$ using suitable clustering algorithm. Every clustered event is translated into clustering input vector using $\text{onehot}_{\#}(e)$ function. These input vectors are then used as data points for the clustering algorithm.

Thus, every activity will have its own independent clustering having $\max_{cc}$ as the maximum cluster count. In the testing phase, the clusterings created from the training data will be used to fit the input vectors created from the events in the testing data.

Finally, we can specify the vector used as input vector in RNN training as follows:

Definition 18. *Generating input vector for one event in training is performed using:*

$$\begin{aligned} \text{inputvector} : e \times \max_{cc} \mapsto & \text{concat}(\text{onehot}_{\# \text{activity}(e)}(e, ACT_{tr}), \\ & \text{onehot}_{\# \text{cluster}(e, \max_{cc})}(e, cl)) \end{aligned}$$

, where $e \in L_{tr}$, and $cl = \langle 1, \dots, n \rangle$, where $n \leq \max_{cc}$ represents the set of all the possible cluster labels. The value of $n \in \mathbb{Z}$ depends on the clustering algorithm and represents the actual maximum number of clusters that were created for event attribute data of any single activity in L_{tr} .

These input vectors are then passed to the RNN training as ordered sequences of event input vectors, where each sequence represents all the events of a single case in the L_{tr} in the order determined by their ascending timestamps.

As a result, when training, every event is preprocessed by performing the input vector generation using inputvector function. At the testing phase, the same encoding functions are used, however, if the event used in testing has some attributes, attribute values or activities that were not part of the training data set, those will just be ignored. Also, event attribute clustering in the testing phase is performed using the clustering models created in the training phase. Thus, in order to store the trained model, also all the trained clustering models must be stored.

5 Test Framework

We have performed our test runs using an extended Python-based prediction engine that was used in our earlier work [9]. The engine is still capable of supporting most of the hyperparameters that we experimented with in our earlier work, such as used RNN unit type, number of RNN layers and the used batch size. The prediction engine we built for this work takes a single JSON configuration file as input and outputs test result rows into a CSV file.

Tests were performed using a commonly used 3-fold cross-validation technique to measure the generalization characteristics of the trained models. In 3-fold cross-validation, the input data is split into three subsets of equal size. Each of the subsets is tested one by one against models trained using the other two subsets.

5.1 Training

Training begins by loading the event log data contained in the two of the three event log subsections. After this, the event log is split into actual training data and validation data that used to find the best performing model out of all the model states during all the test iterations. For this, we picked 75% of the cases for the training and the rest for the validation dataset. After this, we initialize event attribute clusters as described in Section 4.1.

The actual prediction model and the data used to generate the actual input vectors is performed next. This data initialization involves splitting cases into prefixes and also taking a random sample of the actual available data if the amount of data exceeds the configured maximum amount of prefixes. To avoid running out of memory during any of our tests, these limits were set to 75000 for training data and 25000 for validation data. We also had to filter out all the cases having more than 100 events.

Finally, after the model is initialized, we start the actual training in which we concatenate all the requested feature vectors as well as the expected outcome into the RNN model repeatedly for the whole training set until 100 test iterations have passed. The number of actual epochs trained in each iteration is configurable. In our experiments, the total number of epochs was set to be 10. After every test iteration, the model is validated against the validation set. To improve validation performance, if the size of the validation set is larger than a separately specified limit (10000), a random sample of the whole validation set is used. These test results, including additional status and timing-related information, are written into resulting test result CSV file. If the prediction accuracy of the model against the validation set is found to be better than the accuracy of any of the models found thus far, then the network state is stored for that model. Finally, after all the training, the model having the best validation test accuracy is picked as the actual result of the model training.

Table 4: Used Event logs and their relevant statistics

Event log	# Cases	# Activities	# Events	# Attributes	# Unique values
BPIC12 ⁶	13087	24	262200	1	3
BPIC13, incidents ⁷	7554	13	65533	8	2890
BPIC14 ⁸	46616	39	466737	1	242
BPIC17 ⁹	31509	26	1202267	4	164
BPIC18 ¹⁰	43809	41	2514266	5	360

5.2 Testing

In the testing phase, the third subset of cross-validation folding is tested against the model built in the previous step. After initializing the event log following similar steps as in the training phase, the model is asked for a prediction for each input vector built from the test data. To prevent running out of memory and to ensure tests are not taking an exceedingly long time to run, we limited the number of final test traces to 100000 traces and used random sampling when needed. The prediction result accuracy, as well as other required statistics, are written to the resulting CSV file.

6 Test Setup

We performed our tests using several different data sets. Some details of the used data sets can be found in the Table 4. The table lists the number of cases, activities, events and event attributes for each event log. *# Unique values* column shows the sum of all the unique attribute values for each of the selected attributes.

The criteria for selecting or not selecting an event attribute in a model were based on the maximum usage of any unique value that the attribute has in the event log. If a value of an attribute was used in more than 4% of all the events in the event log, then that attribute was included in the test. Besides, we did not select any attributes that had just one unique attribute value. Names of all the selected event attributes are listed in the Table 5.

For each dataset, we performed the next activity prediction where we wanted to predict the next activity of any ongoing case. This was accomplished by splitting every input case into possibly multiple *virtual* cases depending on the number of events the case had. If the length of the case was shorter than 4, the whole case was ignored. If the length was equal or higher, then a separate *virtual* case was created for all prefixes at least of length 4. Thus, for a case of length 6, 3 cases were created: One with length 4, one with 5 and one with 6. For all

⁶ <https://doi.org/10.4121/uuid:3926db30-f712-4394-aebc-75976070e91f>

⁷ <https://doi.org/10.4121/uuid:500573e6-acc-4b0c-9576-aa5468b10cee>

⁸ <https://doi.org/10.4121/uuid:c3e5d162-0cfd-4bb0-bd82-af5268819c35>

⁹ <https://doi.org/10.4121/uuid:5f3067df-f10b-45da-b98b-86ae4c7a310b>

¹⁰ <https://doi.org/10.4121/uuid:3301445f-95e8-4ff0-98a4-901f1f204972>

Table 5: Included event attributes

Event log	Attribute names
BPIC12	lifecycle:transition
BPIC13, incidents	impact lifecycle:transition org:group org:resource organization country organization involved product resource country
BPIC14	Assignment Group
BPIC17	Action EventOrigin lifecycle:transition org:resource
BPIC18	activity doctype note org:resource subprocess

these prefixes, the next activity label was used as the expected outcome. For the full-length case, the expected outcome was a special *finished*-token.

7 Experiments

For experiments, we have used the same system that we used already in our previous work [9]. The system had Windows 10 operating system and its hardware consisted of 3.5 GHz Intel Core i5-6600K CPU with 32 GB of main memory and NVIDIA GeForce GTX 960 GPU having 4 GB of memory. Out of those 4 GB, we reserved 3 GB for the tests. The testing framework was built on the test system using the Python programming language. The actual recurrent neural networks were built using Lasagne ¹¹ library that works on top of Theano ¹². Theano was configured to use GPU via CUDA for expression evaluation.

We used one layer GRU [2] as the RNN type. Adam [11] was used as gradient descent optimizer with parameters of $\beta_1 = 0.9$ and $\beta_2 = 0.999$. 256 was used as the hidden dimension size as well as the mini-batch size and 0.01 as the learning rate. Even though it is quite probable that more accurate results could have been achieved by selecting, e.g., different hidden dimension sizes and learning rates depending on the size of the input vectors, we decided to use the fixed values. This decision was made in order to make the interpretation of

¹¹ <https://lasagne.readthedocs.io/>

¹² <http://deeplearning.net/software/theano/>

the test results easier by minimizing the number of variables affecting the test results.

We performed next activity predictions using all the four combinations of features, five data sets and three different maximum cluster counts: 20, 40, and 80 clusters. The results of these runs are shown in Table 6. In the table, *Features*-column shows the used set of features. *S.rate* shows the achieved prediction success rate. *In.v.s.* shows the size of the input vector. This column can be used to give some kind of indication on the memory usage of using that configuration. Finally, *Tra.t.* and *Pred.t.* columns tell us the time required for performing the training and the prediction for all the cases in the test dataset. In both cases, this time includes the time for setting up the neural network, clusterings and preparing the dataset from JSON format. Sample standard deviation has been included in both *S.rate* and *Tra.t* in parentheses to indicate how spread out the measurements are within all the three test runs. Each row in the table represents three cross-validation runs with a unique combination of dataset and feature that was tested. Rows having the best prediction accuracy within a dataset are shown using bold font. *None* -feature represents the case in which there were no event attribute information at all in the input vector, *ClustN* represents a test with one-hot encoded cluster labels of event attributes clustered into maximum of N clusters, *Raw* represents having all one-hot encoded attribute values individually in the input vector, and finally, *BothN* represents having both one-hot encoded attribute values and one-hot encoded cluster labels in the input vector.

We also aggregated some of these results over all the datasets using the maximum cluster size of 80 clusters. Figure 1 shows the average success rates of different event attribute encoding techniques over all the tested datasets. Figure 2 shows the average input vector lengths. Figure 3 and Figure 4 shows the averaged training and prediction times respectively. Finally, Figure 5 shows the average success rates of different event attribute encoding techniques over all the tested datasets in the case where the maximum cluster count was set to be 80.

Next, we measured the statistical significance of the results. First, we measured whether we could reject the null hypothesis: "the best success rate results could have been achieved by using *Raw* features". By performing a one-tailed t -test, while assuming equal variances, for the results, we find out that this hypothesis can not be completely rejected in any of the test datasets. However, when assessing the null hypothesis: "we can achieve better success rates without taking event attributes into account at all than by taking them into account as clustered attribute values", we can reject it in all the datasets. Similarly, we can easily reject null hypothesis: "training a model using *Raw* features can be as fast as using the most accurate tested clustered features" in all the other cases, except in BPIC12, where the used input vectors were essentially identical due to the small number of unique attribute values in that model.

Based on all of the test results and statistical significance analysis, we can see that having event attribute values included improved the prediction accuracy over not having them included at all in all datasets. The effect ranged from 0.5%

Table 6: Statistics of next activity prediction using different sets of input features

Dataset	Features	S.rate (σ)	In.v.s.	Tra.t. (σ)	Pred.t.
BPIC12	None	85.8% (0.3%)	25.7	489.0s (7.0s)	35.1s
	Clust20	86.0% (0.4%)	30.0	500.6s (2.5s)	31.6s
	Clust40	85.8% (0.3%)	30.0	499.7s (1.3s)	31.9s
	Clust80	86.2% (0.1%)	30.0	502.1s (2.3s)	7.5s
	Raw	85.9% (0.3%)	29	504.3s (0.5s)	38.9s
	Both20	86.0% (0.2%)	33	515.3s (2.6s)	40.4s
	Both40	86.0% (0.4%)	33	517.7s (3.6s)	40.4s
	Both80	86.3% (0.1%)	33	518.2s (4.0s)	40.7s
BPIC13	None	62.9% (0.3%)	13.7	165.6s (21.2s)	3.5s
	Clust20	66.8% (0.3%)	34.7	188.0s (22.4s)	4.7s
	Clust40	67.2% (0.7%)	54.7	214.8s (3.1s)	5.4s
	Clust80	67.0% (0.6%)	94.7	258.4s (4.7s)	6.0s
	Raw	68.2% (1.1%)	2353.7	2611.7s (44.7s)	74.8s
	Both20	69.1% (0.6%)	2359.3	2464.6s (309.0s)	94.4s
	Both40	68.9% (0.5%)	2395.7	2687.1s (227.3s)	106.6s
	Both80	68.4% (0.7%)	2429.3	2821.8s (33.5s)	194.3s
BPIC14	None	37.8% (1.5%)	40.3	488.1s (5.3s)	36.1s
	Clust20	39.9% (0.5%)	61.7	523.3s (3.5s)	40.4s
	Clust40	40.0% (0.3%)	80.3	553.5s (3.8s)	43.6s
	Clust80	40.2% (0.1%)	84.7	556.8s (10.5s)	43.6s
	Raw	39.7% (1.4%)	272.0	825.7s (2.8s)	68.0s
	Both20	40.6% (0.6%)	292.3	907.1s (7.5s)	78.6s
	Both40	40.6% (0.6%)	309.3	943.3s (10.6s)	82.0s
	Both80	37.3% (4.2%)	305.0	935.1s (26.9s)	156.7s
BPIC17	None	86.4% (0.4%)	27.7	518.7s (2.8s)	107.7s
	Clust20	90.8% (0.3%)	48.7	556.3s (3.7s)	132.4s
	Clust40	90.2% (1.4%)	68.3	637.5s (58.3s)	143.9s
	Clust80	90.2% (0.4%)	108.7	647.3s (3.7s)	142.8s
	Raw	89.9% (0.5%)	190	816.4s (5.2s)	164.9s
	Both20	89.9% (0.5%)	211.0	867.8s (3.5s)	188.0s
	Both40	90.2% (0.3%)	230.3	910.9s (19.3s)	193.7s
	Both80	89.6% (0.6%)	271.3	986.5 (4.4s)	197.7s
BPIC18	None	71.3% (9.3%)	43	516.0s (9.5s)	197.0s
	Clust20	79.0% (0.9%)	64.0	588.7s (13.7s)	268.7s
	Clust40	79.9% (0.2%)	84.0	628.4s (2.8s)	286.1s
	Clust80	79.5% (0.1%)	124.0	701.3s (7.4s)	306.9s
	Raw	79.3% (0.4%)	349.7	1173.7s (83.1s)	381.2s
	Both20	79.7% (0.5%)	377.7	1213.1s (48.1s)	463.2s
	Both40	79.9% (0.5%)	401.0	1301.9s (82.9s)	540.3s
	Both80	79.3% (0.5%)	425.7	1405.4s (87.2s)	619.9s

in BPIC12 model to 8.5% in BPIC18. As shown in Figure 1, very similar success rates were achieved using *ClustN* features as with *Raw*. However, model training and the actual prediction can be performed faster using *ClustN* approaches than either *Raw* or *BothN*. This effect is the most prominently visible in BPIC13

Fig. 1: Average prediction success rate over all the datasets

Fig. 2: Average length of the input vector over all the datasets

Fig. 3: Average training time over all the datasets

Fig. 4: Average prediction time over all the datasets

Fig. 5: Average prediction success rate over all the datasets separately

results, where, due to the model having a large amount of unique attribute values, the size of the input vector is almost 68 times bigger and the training time almost 14 times longer using *Raw* feature than *Clust20*. At the same time, the accuracy is still better than not having event attributes at all (about 3.9% better) and only slightly worse (about 1.4%) than when using *Raw* feature. This indicates that clustering can be a really powerful technique for minimizing

the time required for training especially when there are a lot of unique event attribute values in the used event log. Even when using the maximum cluster count of 20, prediction results will be either not affected or improved with a relatively small impact on the training and prediction time. Input vector size and memory usage are affected by clustered features in a similar fashion with the training and testing time: The more there are unique attribute values in the event log, the greater the difference between the input vector sizes needed in *ClustN* and *Raw* features.

In all the datasets, the best prediction accuracy is always achieved either by using only clustering or by using both clustering and raw attributes at the same time.

7.1 Threats to validity

As threats to the validity of the results in this paper, it is clear that there are a lot of variables involved. As the initial set of parameter values, we used parameters that were found good enough in our earlier work and did some improvement attempts based on the results we got. It is most probable that the set of parameters we used were not optimal ones in each test run. We also did not test all the parameter combinations and the ones we did, we tested often only once, even though there was some randomness involved, e.g., selecting the initial cluster centers in the XMeans algorithm. However, we think that since we tested the results in several different datasets using a 3-fold cross-validation technique, our results can be used at least as a baseline for further studies. All the results generated by the test runs, as well as all the source data and the test framework itself, are available in support materials ¹³.

Also, we did not test with datasets having a huge number of event attribute values, the maximum amount tested being 2890. However, it can be seen that since the size of the input vectors is completely user-configurable when performing event attribute clustering, the user him/herself can easily set limits to the input vector length which should take the burden off from the RNN and move the burden to the clustering algorithms, which are usually more efficient in handling lots of features and feature values. When evaluating the results of the performed tests and comparing them with other similar works, it should be taken into account that data sampling was used in several phases of the testing process.

8 Conclusions

Clustering can be applied to attribute values to improve the accuracy of predictions performed on running cases. In four of the five experimented data sets, having event attribute clusters encoded into the input vectors outperforms having the actual attribute values in the input vector. Also, due to raw attribute

¹³ <https://github.com/mhinkka/articles>

values having direct effect to input vector lengths, the training and prediction time will be directly affected by the number of unique event attribute values. Clustering does not have this problem: The number of elements reserved in the input vector for clustered event attribute values can be adjusted freely. The memory usage is directly affected by the length of the input vector. In the tested cases, the number of clusters to use to get the best prediction accuracy seemed to depend very much on the used datasets, when the tested cluster sizes were 20, 40 and 80. In some cases, having more clusters improved the performance, whereas, in others, it did not have any significant impact, or even made the accuracy worse. We also found out that in some cases, having attribute cluster indicators in the input vectors improved the prediction even if the input vectors also included all the actual attribute values.

As future work, it would be interesting to test this clustering approach also with other machine learning model types such as more traditional random forest and gradient boosting machines. Similarly, it could be interesting to first filter out some of the most rarely occurring attribute values before clustering the values. This could potentially reduce the amount of noise added to the clustered data and make it easier for the clustering algorithm to not be affected by noisy data. Another idea that we leave for future study is whether it would be a good idea to first perform some kind of a feature selection algorithm such as influence analysis [12], recursive feature elimination [7] or mRMR [3] to find the attribute values that correlate the most with the prediction results and have those attribute values added into the input vectors as raw one-hot encoded attribute values in addition to having the one-hot encoded cluster labels. More work is also required to understand exactly what properties of the event log affect the optimal number of clusters to use. Finally, more study is required to understand whether a similar clustering approach performed for event attributes in this work could be applicable also for encoding case attributes.

9 Acknowledgments

We want to thank QPR Software Plc for funding our research. Financial support of Academy of Finland project 313469 is acknowledged.

References

1. G. Bernard and P. Andritsos. Accurate and transparent path prediction using process mining. In T. Welzer, J. Eder, V. Podgorelec, and A. K. Latific, editors, *Advances in Databases and Information Systems - 23rd European Conference, ADBIS 2019, Bled, Slovenia, September 8-11, 2019, Proceedings*, volume 11695 of *Lecture Notes in Computer Science*, pages 235–250. Springer, 2019.
2. K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio. On the properties of neural machine translation: Encoder-decoder approaches. In D. Wu, M. Carpuat, X. Carreras, and E. M. Vecchi, editors, *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Doha,*

- Qatar, 25 October 2014*, pages 103–111. Association for Computational Linguistics, 2014.
3. C. H. Q. Ding and H. Peng. Minimum redundancy feature selection from microarray gene expression data. *J. Bioinformatics and Computational Biology*, 3(2):185–206, 2005.
 4. J. Evermann, J. Rehse, and P. Fettke. Predicting process behaviour using deep learning. *Decision Support Systems*, 100:129–140, 2017.
 5. C. D. Francescomarino, M. Dumas, F. M. Maggi, and I. Teinemaa. Clustering-based predictive process monitoring. *CoRR*, abs/1506.01428, 2015.
 6. C. D. Francescomarino, C. Ghidini, F. M. Maggi, and F. Milani. Predictive process monitoring methods: Which one suits me best? In Weske et al. [22], pages 462–479.
 7. P. M. Granitto, C. Furlanello, F. Biasioli, and F. Gasperi. Recursive feature elimination with random forest for PTR-MS analysis of agroindustrial products. *Chemo-metrics and Intelligent Laboratory Systems*, 83(2):83–90, 2006.
 8. J. A. Hartigan and M. A. Wong. Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108, 1979.
 9. M. Hinkka, T. Lehto, K. Heljanko, and A. Jung. Structural feature selection for event logs. In E. Teniente and M. Weidlich, editors, *Business Process Management Workshops - BPM 2017 International Workshops, Barcelona, Spain, September 10-11, 2017, Revised Papers*, volume 308 of *Lecture Notes in Business Information Processing*, pages 20–35. Springer, 2017.
 10. M. Hinkka, T. Lehto, K. Heljanko, and A. Jung. Classifying process instances using recurrent neural networks. In F. Daniel, Q. Z. Sheng, and H. Motahari, editors, *Business Process Management Workshops - BPM 2018 International Workshops, Sydney, NSW, Australia, September 9-14, 2018, Revised Papers*, volume 342 of *Lecture Notes in Business Information Processing*, pages 313–324. Springer, 2018.
 11. D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
 12. T. Lehto, M. Hinkka, and J. Hollmén. Focusing business improvements using process mining based influence analysis. In M. L. Rosa, P. Loos, and O. Pastor, editors, *Business Process Management Forum - BPM Forum 2016, Rio de Janeiro, Brazil, September 18-22, 2016, Proceedings*, volume 260 of *Lecture Notes in Business Information Processing*, pages 177–192. Springer, 2016.
 13. N. Navarin, B. Vincenzi, M. Polato, and A. Sperduti. 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*, pages 1–7. IEEE, 2017.
 14. T. Nolle, A. Seeliger, and M. Mühlhäuser. Binet: Multivariate business process anomaly detection using deep learning. In Weske et al. [22], pages 271–287.
 15. D. Pelleg and A. W. Moore. X-means: Extending k-means with efficient estimation of the number of clusters. In P. Langley, editor, *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000), Stanford University, Stanford, CA, USA, June 29 - July 2, 2000*, pages 727–734. Morgan Kaufmann, 2000.
 16. P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
 17. S. Schöning, R. Jasinski, L. Ackermann, and S. Jablonski. Deep learning process prediction with discrete and continuous data features. In E. Damiani, G. Spanoudakis,

- and L. A. Maciaszek, editors, *Proceedings of the 13th International Conference on Evaluation of Novel Approaches to Software Engineering, ENASE 2018, Funchal, Madeira, Portugal, March 23-24, 2018.*, pages 314–319. SciTePress, 2018.
18. N. Tax, I. Teinemaa, and S. J. van Zelst. An interdisciplinary comparison of sequence modeling methods for next-element prediction. *CoRR*, abs/1811.00062, 2018.
 19. N. Tax, I. Verenich, M. L. Rosa, and M. Dumas. Predictive business process monitoring with LSTM neural networks. In E. Dubois and K. Pohl, editors, *Advanced Information Systems Engineering - 29th International Conference, CAiSE 2017, Essen, Germany, June 12-16, 2017, Proceedings*, volume 10253 of *Lecture Notes in Computer Science*, pages 477–492. Springer, 2017.
 20. W. M. P. van der Aalst. *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011.
 21. I. Verenich, M. Dumas, M. L. Rosa, F. M. Maggi, D. Chasovskyi, and A. Rozumnyi. Tell me what’s ahead? predicting remaining activity sequences of business process instances. Unpublished, June 2016.
 22. M. Weske, M. Montali, I. Weber, and J. vom Brocke, editors. *Business Process Management - 16th International Conference, BPM 2018, Sydney, NSW, Australia, September 9-14, 2018, Proceedings*, volume 11080 of *Lecture Notes in Computer Science*. Springer, 2018.