<https://helda.helsinki.fi>

Structural Feature Selection for Event Logs

Hinkka, Markku

Springer-Verlag
2018

Hinkka , M , Lehto , T , Heljanko , K & Jung , A 2018 , Structural Feature Selection for Event Logs . in E Teniente & M Weidlich (eds) , BUSINESS PROCESS MANAGEMENT WORKSHOPS (BPM 2017) . Lecture Notes in Business Information Processing , vol. 308 , Springer-Verlag , pp. 20-35 , 15th International Conference on Business Process Management (BPM) , Barcelona , Spain , 10/09/2017 . https://doi.org/10.1007/978-3-319-74030-0_2

<http://hdl.handle.net/10138/307222>

https://doi.org/10.1007/978-3-319-74030-0_2

other

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.

Structural Feature Selection for Event Logs

Markku Hinkka^{1,2}, Teemu Lehto^{1,2}, Keijo Heljanko^{1,3}, and Alexander Jung¹

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

² QPR Software Plc, Finland

³ HIIT Helsinki Institute for Information Technology

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

Abstract. We consider the problem of classifying business process instances based on structural features derived from event logs. The main motivation is to provide machine learning based techniques with quick response times for interactive computer assisted root cause analysis. In particular, we create structural features from process mining such as activity and transition occurrence counts, and ordering of activities to be evaluated as potential features for classification. We show that adding such structural features increases the amount of information thus potentially increasing classification accuracy. However, there is an inherent trade-off as using too many features leads to too long run-times for machine learning classification models. One way to improve the machine learning algorithms' run-time is to only select a small number of features by a feature selection algorithm. However, the run-time required by the feature selection algorithm must also be taken into account. Also, the classification accuracy should not suffer too much from the feature selection. The main contributions of this paper are as follows: First, we propose and compare nine different feature selection algorithms by means of an experimental setup comparing their classification accuracy and achievable response times. Second, we discuss the potential use of feature selection results for computer assisted root cause analysis as well as the properties of different types of structural features in the context of feature selection.

Keywords: automatic business process discovery, process mining, prediction, classification, machine learning, clustering, feature selection

1 Introduction

In Process Mining, unstructured *event logs* generated by systems in business processes are used to automatically build real-life process definitions and as-is models behind those event logs. There is a growing need to be able to predict properties of newly added event log cases, or process instances, based on case data imported earlier into the system. In order to be able to predict properties of the new cases, as much information as possible should be collected that is related to the event log cases 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 machine learning techniques.

One good source of additional case related features is the information stored into the *sequence of activities* visited by cases. This information includes, e.g., number of times an event log case has visited a certain *activity* and the number of times a case has *transitioned* between two specific activities. Features collected in this fashion are often highly dependent on each other. E.g., a patient whose visit to hospital takes long time (outcome) has quite often been in surgery from which he/she has moved into a ward. In this example, we already can easily find three structural features of which any can be used to predict whether the visit lasted long or not: visited surgery, transitioned from surgery to ward, visited ward. However, depending on all the other patients in the data set, it may be that there are no cases where only a subset of these three features occurs, thus making it redundant to have all three features taken into account when building a model for prediction purposes. Thus, one feature could well be enough to give as accurate predictions as having them all.

Another important aspect in Process Mining is that it is often desired to be able to show dependencies between features. Thus, selecting a feature selection algorithm that produces also this information for minimal extra cost is often tempting. For this purpose, the list of the most relevant features and the extent of their contribution should somehow be returned. One example of this kind of root cause analysis technique is influence analysis described in [19].

The primary motivation for this paper is the need to perform classification based on structural features originating from *activity sequences* in event logs as accurately as possible and using a minimum amount of computing resources and maximizing the throughput in order to be able to use the method even in some interactive scenarios. This motivation comes from the need to build a system that can do classification and root cause analysis activities accurately on user configurable phenomena based on huge event logs collected and analyzed, e.g., using *Big Data processing frameworks* and methods such as those discussed in our earlier paper [17]. The response time of this classification system should be good enough to be used as part of a web browser based interactive process mining tool where user wants to perform classifications and expects classification results to be shown within a couple of seconds.

As an extension to the original paper presented in BPAI 2017 conference, we have increased the amount of feature selection methods by examining 3 new feature selection methods and compared their performance with the original set of algorithms. The first new tested algorithm is Fisher scoring based approach, whereas the two other algorithms are hybrid algorithms mixing Clustering separately with mRMR and Fisher scoring algorithms. Some additional details were also added into the discussion of relevant concepts including a more detailed study on related works by adding a new Section 2.3 discussing structural features that are conceptually and also computationally more complex structural features often referred to in related literature. In the experiment sections we

included three new publicly available BPI Challenge data sets to be used in the tests.

The rest of this paper is structured as follows: Section 2 introduces main concepts related to this paper. Section 3 discusses the feature selection concept and gives brief introduction to the methods used in this paper. Section 4 will then present a framework used for comparing performance of the selected feature selection approaches. The results of the tests will be presented in Section 5 after which we will discuss related work in Section 6. Finally, Section 7 draws the conclusions from the test results.

2 Problem Setup

The concepts and terminology used throughout this paper mostly follow those commonly used in process mining and machine learning communities. However, the following subsections will provide some short introduction to the most important concepts related to this paper.

2.1 Concepts

An *event log* can be any ordered list of records known as *events*. Every event has at least a *case identifier*, an *activity identifier* and some additional property such as a time stamp that can be used to put the events into some deterministic order. A *case identifier* is used to group *events* belonging somehow into some common contexts. This could be, e.g., session identifier in a web server log, a customer or order identifier in a customer and order handling system such as CRM. Thus, for every event, an unambiguous *case* can be identified which represents a collection of *events* belonging to the same process. The events for a case are represented in the form of a *trace*, i.e., a sequence of unique events [31]. In addition to the properties listed above, every event can also include any number of additional *event attributes*. Similarly, every case can include any number of additional *case attributes*. A *transition* represents the transition between two successive activities within a case.

In this paper, the term *activity sequence* is used to represent the ordered set of activity labels within cases. *Activity sequences* are represented as strings using a notation used by Aalst et al. [31] augmented by "virtual" start and end *activities* making it possible to deal with case start and end events and transitions leading to them from outside of the process using similar notation as for normal activity to activity transitions. S is used to denote this virtual *start event* and E is used for *end event*.

E.g., a activity pattern representation $\langle S, a, b, c \rangle$ represents all the *activity sequences* which start from *activity* labeled **a**, since there is **S**-activity in the beginning of the representation. **a** is then followed by **b** and **c** activities. Thus, this pattern can match both of the following activity sequences: $\langle S, a, b, c, E \rangle$, $\langle S, a, b, c, d, E \rangle$.

Pattern / predictor type	Example sequences(s)
Activity	$\langle a \rangle, \langle b \rangle$
Transition / 2-grams	$\langle S, a \rangle, \langle a, b \rangle, \langle b, b \rangle, \langle b, E \rangle$
Starter	$\langle S, a \rangle$
Finisher	$\langle b, E \rangle$
Ordering	$\langle a \rangle \rightarrow \langle b \rangle$

Table 1: Structural Feature Types

2.2 Structural Features

As opposed to normal case attributes added to cases in event logs, structural feature term in this paper is used for representing properties of activity sequences of cases. Thus, they can be derived directly from actual activity sequences without need to include any additional custom properties. Having a case identifier, activity identifier and order information such as a time stamp for each event occurrence, is enough. In order to simplify the tests and keep requirements for applying the results of this paper to its minimum, we decided to only concentrate on structural features as predictors in this paper. However, in real use cases, the best results are achieved by including also all the available additional case attributes such as duration, age, etc. into features from which the feature selection is performed [20].

There are several different types of structural features to select from. In this paper, we use notations similar to those used in regular expressions [29] combined with notation commonly used for activity sequences [31]. The patterns we focused in this paper are listed in Table 1 with examples of matches when the sequence of activities is illustrated as $\langle S, a, b, b, E \rangle$.

For every predictor type listed in the Table 1, there can be several possible implementations. In this paper, we consider structural features of *activity* and *2-gram* predictor types to be such that their values correspond to the number of occurrences of that pattern within each activity sequence. *Starter* and *finisher* predictor types however are boolean values indicating whether that pattern is valid for an activity sequence. Order feature type is considered to be a boolean value such that it is true only if the first occurrences of both ends of the order relation are in the specified order.

The difference between 2-gram and order pattern is that order allows any number of activities to be between the activities of the ordering relation, whereas in 2-gram, the activities of the relation must be successive in the whole sequence of activities. The importance of predictor types also depends very heavily on the type of the data set and the scenario being predicted.

One more factor to take into account when selecting the actual features is how to handle situations where a feature has more than two different values. E.g., a patient may have visited surgery multiple times while visiting a hospital. In some cases, depending on what we are trying to predict and what kind of prediction models are to be built, it could be better to split these kinds of features into several boolean features. Thus, e.g., we could have a feature for a patient having

visited surgery 4 times. However, one has to be careful when to split features like this in order to avoid an explosion in the number of features created for each original feature.

One final step before passing features to the actual model training is to identify and remove any duplicate features that have behaved identically through the whole training set. Some training methods do this automatically, but some do not.

2.3 More Complex Structural Features

As we were interested in minimizing the response time, we decided to consider only the patterns listed in the previous chapter because having more complex patterns would have made the combinations of different predictor types and features to become too big to be able to perform exhaustive tests for and leading to the problem of *curse of dimensionality* [14]. Also, the extraction of all the feature types presented in the literature would have required much more computation time than the selected relatively simple features used in this work. However, in this chapter we briefly explore some of the other widely used feature types.

Some of the most studied patterns in the process mining community have been related to discovering repeats such as tandem- and maximal repeats [3][2][24] within the sequences of activities and discriminative patterns [24][22][5].

For example, maximal repeat α is defined as being a subsequence of a sequence of activities S where α occurs at least twice within S [2] and extending α to include an additional character to left or right of the sequence would break the inequality. Similarly, a tandem repeat α is defined as being a subsequence of a sequence of activities S where α occurs successively multiple times S . Both of these repeat types can be extracted efficiently using suffix trees [13]. However, due to relative complexity of feature extraction from the source event log, we decided not to explore these feature types in this paper.

Discriminative patterns are referred by Lo et. al. in their article[22] are capturing repetitive iterative patterns within the sequences of activities in traces. After this, a feature selection algorithm based on Fisher scoring [9] is used to select the most discriminating features to be used for actual classification. This method picks the most important repetitive iterative patterns and uses only them as the set of features. The biggest difference this algorithm has to the methods we are using in this paper is that they concentrated only on iterative patterns and that they use only Fisher scoring -based algorithm. Due to the complexities involved in extracting iterative patterns, we decided to include in our paper just the portion where Fisher scoring is used for feature selection.

One solution for overcoming the problem of selecting the features in classification is also to use a feature selection technique that does not require specifying the value of each feature separately as input to the model training. One such solution is to use neural networks and deep learning based techniques such as Long short-term memory [18][10]. The problem with this approach for our purposes is that the resulting model can not be easily used for root cause analysis and to find out what exact structural features may be causing the observed effect.

2.4 Classification

Since all the tests performed in this paper are performed on data sets consisting of already completed cases, we are performing *classification* using machine learning prediction algorithms. Classification in machine learning usually consists of two phases: training a model and performing the actual classifications using the trained model. In this paper, we concentrate on building the classification model using *supervised learning methods*, where algorithms are trained using *predictors*, together with their *outcomes*. A core part of the model building is the selection of *features* to be used as predictors. Often the more you have independent features that may have effect in the outcome, the better. As shown in Section 2.2, a lot of features can be created directly from the activity sequences of the cases themselves.

Another important factor that has direct effect to the prediction performance of the model is the algorithm that is used for building the model and making the predictions. In this paper, we focus on the feature selection part. However, we need to also validate the performance of the feature selection using a set of algorithms. In order to minimize the skew in the results caused by the validation algorithm itself, we decided to compare the efficiency of the selected features using two different approaches. First, for a given set of selected features, we determined the prediction accuracy obtained by a particular supervised learning method, i.e., the gradient boosting machine (GBM). GBM was selected due to its good reputation [25, 14] and performance in both accuracy and response times in our own tests. The second method was to approximate the mutual information score [23] between each of the selected set of features and two different data sets. The first data set consisted of all the available predictors without any feature selection. The second data set consisted only of the outcomes to be predicted.

As we are concentrating on features originating from process mining, at the granularity level of a case, the prediction inputs that are usually used in the field are actually custom case properties such as the customer name or an identifier of the owner of the case. The outcomes that we want to predict are usually values of some custom case properties or some calculated case content dependent values such as durations, some kind of cost of the case or some other metrics measuring the quality of the case. In this paper, we concentrate only in features inherent to the activity sequences inside cases and measure how well certain outcomes can be predicted only based on those features.

The used data set is split into two parts: training and test. Training data set is used for two purposes. First, features are selected from the whole training data set. After this, a model is trained using all the cases in the training data, but only using the selected features as predictors. Finally, once the model has been built, the model is tested against the test data and its performance is estimated using accuracy and mutual information metrics.

3 Feature Selection Methods

The aim of feature selection is to reduce the dimensionality of the structural features constructed from the raw data. Reducing the dimensionality not only reduces the computational complexity of the subsequent prediction methods, it may also lead to an improved prediction accuracy. Indeed, learning algorithms based on a smaller set of features are less prone to overfitting, i.e., the effect of erratic statistical variations of a particular observed dataset is reduced. Finally, feature selection also enhances the interpretability (visualization) of the features and understanding classifications based on them (e.g., if only two numerical features are selected, we can illustrate them by means of a scatterplot).

Initially we also considered testing a couple of feature extraction algorithms. Feature extraction differs from feature selection in that they create new features that will be used instead of the original features. The newly created features try to maximize the variance and expressive power of the features by combining several original features into one new feature. This has a drawback that it hides the original features and makes it harder to understand the properties of the created model. E.g., in root cause analysis, it is often desirable to understand how much the outcome depends of certain features and also to understand which features have an effect to the outcome. Due to this shortcoming, we decided to not include any feature extraction algorithms into this paper.

No additional parallelization techniques were used, thus if the algorithm did not support parallelism out of the box, it was not run in parallel. The following subsections briefly describe the basics of each of the feature selection methods tested in this paper including information on the used R programming language packages and their configurations. We also briefly tested an algorithm based on *Support Vector Machine (SVM)* [1] [36] using radial kernel, but decided to leave it out of the paper due to very poor results and extremely long response times, which were order of magnitude slower than with any of the other tested algorithms. The following subsections will briefly describe the details of all the remaining tested algorithms.

3.1 Random Selection

The most trivial of all the tested algorithms was a randomized selection where the desired number of features were just randomly selected from all the available features. This method was used as a baseline in order to gain a better understanding on the quality of other used selection algorithms when compared with an algorithm that does not in any way take any properties of the selected features themselves into account. This serves as a baseline selection algorithm. There should not be any algorithm that performs consistently worse than this. However, in order to alleviate the effect of inherently noisy random selections, median of three separate test runs was used in the experiments. Thus, only the test which yielded the median prediction accuracy was used as the actual result. In the graphs and analysis below, this algorithm is labeled as *Random*.

3.2 Fisher Scoring

Fisher scoring is based on measuring a Fisher score [9] for each available feature after which N features that produced the highest Fisher score are used as the selected features. Fisher scores behave in such a way that in order for a feature to have a high value, it must have very similar values within one classification value but very dissimilar values between different classification values. In the sections below, this algorithm is labeled as *Fisher*.

3.3 Feature Clustering

This method is influenced by the idea provided by Covoes et al. [7]. In the algorithm developed for this paper, the training data is first clustered so that every structural feature in the training set constitutes one clustering data point. Each activity sequence in the training data represents one dimension for clustering data points with values equaling the number of times that structural feature occurs within that activity sequence. K-means algorithm is then used to generate K clusters using `kmeans` R-function which is based on algorithm by Hartigan et al. [15]. For each K clusters, the feature having the minimum distance to the mean of that cluster will be selected as the representative for all the features in that cluster. It should be noted also that as a side product of applying this method for feature selection, every selected feature will actually represent all the features within the same cluster. Thus, for every original feature, you have one cluster it belongs into and exactly one feature that is representing that feature in that cluster. This could be useful, e.g., in some root cause analysis scenarios.

It should be noted that K-means feature clustering, being an unsupervised learning method, does not take outcomes into account in any way and thus divides the data point space evenly without any kind of weighting or prioritization. However, clustering method, as any other feature selection method used in our paper, can easily be used in combination with other feature selection algorithms. This makes it possible to implement *hybrid feature selection methods* where more than one method is used to select the final set of selected features. We used this technique to combine clustering method with several other methods to combine the properties of methods.

Four different versions of this algorithm that are covered in this paper are briefly discussed in the following subsections.

Clustering Only For this paper we explored two different approaches One that first removed all the features having exactly the same occurrence pattern within all the cases thus removing duplicate vectors before the actual clustering step. The other variation of this algorithm does not perform this preprocessing step. The results of different variations being nearly the same except for the processing time, which was clearly faster with the algorithm that first dropped out all the features having exactly the same values for all the cases in the training set. Thus, we decided to limit our tests only to this algorithm variant. In the graphs shown below, this algorithm is labeled as *Cluster*.

Clustering with Variable Importance In variable importance based feature selection, some Machine learning algorithm capable of building variable importance information, such as random forest [21], is first performed. After this, the results of the algorithm are used to pick N variables having the greatest effect on the outcome. These N variables are then used as the selected features. Since the performance of variable importance algorithm itself was found to be very poor when using predictor types having hundreds of features we decided to use a hybrid approach where we first use the clustering approach described above to remove about 75% of all the features, after which variable importance is calculated for each feature using random forest algorithm and from there, the desired number of the most important target features is picked. In this paper, `randomForest` -R package is first used to generate a model after which `varImp` -R function in `Caret`-package is used to extract the most important features based on the information gathered by the random forest algorithm. This algorithm is labeled as *ClustImportance* in the graphs and analysis below.

Clustering with Fisher Scoring This is a hybrid feature selection method where the normal clustering method described in Section 3.3 is performed in a way that the amount of features is brought down from the original full set of features to twice the number of features that are to be selected. If there are less than twice the number of features to pick in the first place, then this clustering step is skipped. After clustering has been performed, Fisher method is used to pick exactly the desired number of features as the final set of selected features. This method is referred to as *ClustFisher* in the sections below.

3.4 Minimum Redundancy Maximum Relevancy

This is a mutual information based approach [8], which uses mutual information as a proxy for computing relevance and redundancy among features. The implementation used in this paper was provided by `mRMRe` -R package which claims to provide a highly efficient implementation of the mRMR feature selection via parallelization and lazy evaluation of mutual information matrix. We used ensemble method both with `solution_count` set to 1, which provides results resembling classic mRMR, and also with 5, which does 5 separate runs and combines the results in the end. This time the results were also otherwise quite the same, except the 5-run version provided clearly better mutual information scores. Thus, in the graphs and analysis below, we use only 5 run version labeled as *mRMREns5*.

After finding out that both mRMR and Cluster seemed to be quite efficient methods for feature selection, it was decided to include also a hybrid of these two methods in a similar fashion to *ClustFisher*. In this case, Clustering is used to select twice the number of desired features, after which `mRMREns5` is used to make the final selection out of them. This hybrid feature selection is referred to as *ClustmRMR* in the sections below.

3.5 Least Absolute Shrinkage and Selection Operator (LASSO)

This is a regression analysis method that can be used for feature selection [30]. It is related to least squares regression where the solution minimizes the sum of the squares of the errors made. The unique property for this regression technique is the usage of additional regularization that enables discarding irrelevant features and forces usage of simpler models that do not include them. Since the LASSO implementation in *glmnet* R-package in itself did not provide means of sorting features by their importance and since it was not possible to directly adjust the desired number of target features, the actual used algorithm first performed 10 iterations of LASSO algorithm each yielding slightly different results. After this, all the results were collected into a single list with each feature weighted by the number of occurrences of that feature within all the LASSO results. Finally, this list was sorted from the largest height to smallest and the desired number of features were picked from the beginning of this sorted list. Two different variations of this algorithm were tested: one using λ_{1se} as the prediction penalty parameter and the other using λ_{min} . Due to the results being almost identical in both the cases, we selected the one using λ_{1se} prediction penalty parameter. In the results below, this algorithm is labeled as *LASSO1se*.

3.6 Markov Blanket

Markov blanket of a variable X is a minimal variable subset conditioned on which all other variables are probabilistically independent of X . [37] For Bayesian networks, Markov blanket of X consists of the union of the following three types of neighbors: the direct parents of X , the direct successors of X , and all direct parents of X 's direct successors. Bayesian networks can be inferred from the training data after which Markov blanket for the created network is calculated by selecting the outcome as X . The result is the set of features to select. *bnlearn* -R package was used to perform Markov Blanket based feature selection. *Hill-Climbing* algorithm is first used to construct a Bayesian network structure out of the training data. After this, Markov Blanket is extracted out of the network for the outcome feature. Finally, out of these results, the desired number of features are selected from the beginning of the returned list, or if the result does not have all the required features, only the returned features are selected. In the results below, this algorithm is labeled as *Blanket*

3.7 Recursive Feature Elimination

Recursive feature elimination [12] starts with estimating the variable importances of all the features in the training data as in the Variable Importance -technique. After this, a smaller subset of the most important features is selected and variable importances are estimated again. This is repeated until the desired feature subset size is reached after which the resulting features can be picked. In this paper, three different variations of this method were tested: a test with only one iteration, another with two iterations and the third one with

Event Log	# Cases	# Selected	# Act	# SF	# 2-Gram	# Order	# Total
BPIC14-4k [34]	4000	1441 / 581	39	20	772	1033	1864
BPIC14-40k [34]	40000	8108 / 7473	39	20	772	1033	1864
BPIC12 [33]	13087	3330	36	21	161	866	1084
BPIC13, incidents [27]	7554	1579	12	8	75	129	224
BPIC17 [35]	31509	11584	26	14	165	459	664
Hospital [32]	1143	372	624	36	4272	79571	84503

Table 2: Used Event logs and numbers of features by predictor types

four iterations. Caret R-package’s *rfe* algorithm was used for these tests, with the default method based on random forests. After the initial tests, it was found out that while the accuracy and mutual information of all the cases were very close to each other, the average processing time of the 2-step algorithm was clearly better than the others. Thus, in the graphs and analysis below, we only concentrate on this algorithm labeled as *Rec2S*.

4 Test Setup

Testing was performed on a single system using Microsoft R Open version 3.3, Windows 10 operating system. The used hardware consisted of 3.5 GHz Intel Core i5-6600K CPU with 8 GB of memory. Tests were performed using five publicly available data sets. All the required structural features were extracted from actual event logs using the query interface of QPR ProcessAnalyzer [26].

Initial tests were performed on 4000 and 40000 case sample of BPI Challenge 2014 dataset [34]. These tests were used to make the number of tested feature selection methods smaller by dropping all the methods that clearly do not fit the interactive performance requirements that were part of the requirements. After this, all the rest of the datasets were used. Table 2 shows additional details of each dataset including exact numbers of extracted structural features of different types and the number of cases that belonged into the classification being predicted, which is shown in “# Selected” -column. The table has two rows BPIC14 case, one for 4000 case sample and one for 40000 case samples. Also for both of these version, two numbers are shown in the “# Selected” -column. The former is the number of cases having long durations and the latter the number of cases that represent a “request for information”. The total number of cases in BPIC14, without any sampling, is 46616 cases.

All the test runs were performed using an R function that ran all the desired test runs in sequence. At the beginning of every test run, random seed was initialized. Thus, the random case samples, used in BPIC14 data set, and other random values generated within the used algorithms behaved the same way in every run, provided that the algorithm used random -methods that support setting the seed using *set.seed* -R function.

In the tests, the training data was first extracted so that 25% of the provided data rows were randomly selected. This training data was first used by the feature selection algorithm to be tested, after which it was used to build the classification *GBM* model for predicting given phenomenon. This classification model was then used for measuring the performance of the feature selection. Mutual information metrics were approximated also at this final phase.

The first run of tests was performed using test data having 4000 cases extracted from the full BPI Challenge 2014 data set. For this first run, all the algorithms were tested so that the number of selected features were 10 and 30. For each of these combinations, 13 different sets of feature patterns were selected. The selected structural feature patterns were different combinations of the following patterns described in Table 1: activity, starter and finisher, 2-grams and ordering. In BPIC14-event logs, activity and 2-grams features included occurrence counts, while all the other feature types were just boolean values indicating whether the feature occurs at least once in a case. In all the other datasets, all the features are occurrence counts.

The combinations of used feature types were created in a way that all the possible combinations of the patterns were tested where activity pattern was present. We also included some other more interesting combinations, thus generating 11 different pattern combinations. In addition to these, we also tested 2-grams and ordering separately, as well as having both 2-grams and ordering. As discussed in the previous sections, we did not want to include any other patterns due to the number of potential new features that would have been needed in order to cover all the possible cases. E.g., adding 3-grams would potentially have generated N^3 additional features where N equals the number of different activities in the training data, which in this case is 39 yielding the maximum of 60000 new features.

All the tests performed on the first data set were run to predict two different outcomes, which are later in this paper referred to as scenarios. The first scenario was whether the case duration is longer than 7 days. In this case, nearly 36% of all the cases in the small test set had this outcome. This is an example of a prediction that can be trained directly from the event information without any need for additional case or event attributes. The second scenario that was tested is based on additional case-level information provided with the event data: Does the case represent a "request for information" or something else such as an "incident"? In this case, nearly 15% of all the test cases in the small sample had this outcome. For all the tests performed on the same sample size, the actual used cases and their predictors were always the same.

For all the other data sets, a classification was made based on the duration of cases. In BPIC12 and BPIC13, duration threshold was set to 2 weeks. In BPIC17, 4 weeks was used as threshold. In Hospital data set, 20 weeks was used.

5 Test Results

We began the actual analysis of our first round of tests by estimating the average classification accuracy of all the tested algorithms for all the tested feature counts and all the tested structural feature patterns using both the test scenarios in 4000 case sample. After the first test run, a second run was made with the same BPIC14 data set, but this time with a sample of 40000 cases. At this point we also decided to not include Blanket method into this test run due to its very slow performance and often did not manage to finish at all. The results of these test runs are shown in Figure 1. The first column on the chart labeled *None* shows the accuracy achieved by not performing any feature selection. Each column in the chart represent an average of 44 test run results, except for *None* column which represents the average of 22 test run results, since selecting different number of features is not applicable for it. Based on these results, we can immediately see that increasing the sample and also training set size increases also classification accuracy. We can also see that for 4000 case sample, the top three feature selection algorithms ordered in the descending order of accuracy are: Recursive, Cluster and Cluster with mRMR. However, when the sample size is increased to 40000 cases, the order is changed to: Cluster, Recursive, Cluster with Fisher while mRMR dropped to 5th place.

With the same test data, we also measured the average mutual information score of the data set filtered using feature selection algorithm. This is shown in Figure 2. In this chart, *None* shows the absolute maximum score achievable by any feature selection algorithm. The top three ranking of algorithms when ranked by mutual information is: Recursive, Cluster and mRMR. With 40000 case dataset, the ranking order is Cluster, Cluster with Fisher and LASSO. When the mutual information is calculated between the result of the feature selection algorithm and the expected outcome, the ranking is for smaller sample: Recursive, Cluster and Cluster with mRMR. For bigger sample this becomes: Cluster, Cluster with Fisher and LASSO, while mRMR can be found at 7th place.

After this, we analyzed the response times for all the tested feature selection algorithms with the same test data. This time however, we did not include starter and finisher predictor types since they would have made the readability of the figure much worse and also would not have provided much additional information due to the small effect they have into the results in the tested scenarios.

As seen in the Figure 3, the time required to perform feature selection for the tested algorithms and predictor types varied very much. Each column in this chart represents an average time required by 4 test runs. In the worst cases, the difference between the slowest and the fastest algorithm was three orders of magnitudes, with Fisher, Cluster, Cluster hybrids and mRMR usually performing much faster than all the rest. Out of these, Clusters usually slightly outperformed mRMR. With the 40000 case dataset, the most notable change is that mRMR now comes very close to, and sometimes even outperforms, Fisher while both of them being almost one order of magnitude faster than Cluster algorithms which in turn were at least one order of magnitude faster than the

Fig. 1: Average accuracy of all the tested algorithms.

Fig. 2: Average mutual information of all the tested algorithms.

LASSO and Recursive algorithms. Surely better performance could have also been achieved by optimizing the used R -scripts or by using some tailor-built natively compiled software, especially for algorithms that had more scripted parts such as Fisher and Cluster algorithms.

The largest performance variations within one algorithm were measured using Blanket algorithm which performed in the fastest predictor types, almost as fast as the fastest algorithms, but in the slowest predictor types, it performed almost over four orders of magnitudes slower.

Based on this data, it was decided to drop Blanket, LASSO1se and Rec2s from any further examinations. We also decided to drop Fisher and Cluster

Fig. 3: Average feature selection response time of all the tested algorithms.

Importance due to their somewhat poor classification accuracy compared to other remaining algorithms.

Next we used all the remaining algorithms to run all the tested test parameter combinations on several other datasets. Figure 4 shows the accuracy achieved when using all the tested predictor types as source features for the feature selection using the narrowed down set of algorithms and all the tested datasets. mRMREns5 columns are missing in Hospital-data set in the figure because the used R implementation of mRMR algorithm supported only maximum of 46340 features that was exceeded in that case. Figure 6 shows otherwise the same information, except this time only structural features of *activity* predictor type are included as source features. Figure 7 shows the same information when only features of *2-gram* predictor type are available. Finally, Figure 8 shows how the accuracy of the tested algorithms change when the available predictor types are changed.

Based on these figures, it can be seen that mRMR and Clustering both provide nearly as accurate results. BPIC17 is the only tested dataset where mRMR provides slightly better average accuracy than Cluster over all the tested predictor types. However, also in this case, mixing both Cluster and Fisher provided even better accuracy. Hospital dataset, which has very few cases and lots of different types of structural features, was the only dataset where hybrid cluster algorithms provided clearly better results than the normal Cluster dataset. It also seems that when the amount of cases is low compared to the amount of features, mRMR does not manage to get as accurate results as the Cluster algorithms. Generally, based on all these tests, it can be said that the order from the most accurate to the least accurate algorithm is: Cluster, Cluster Fisher, mRMR and Cluster mRMR. However, the difference between the worst and the best, based on 72 experiment results for every algorithm is only 0,7 %, which is not much. One small detail to note is that despite mRMR getting slightly inferior accuracies, it managed to get the best accuracy classification of all the experiments for BPIC12 dataset using only *2-gram* predictors where Cluster managed to get its slightly worse value using *activity*, *starter* and *finisher* predictors. It seems that having hybrid algorithm only helps when the amount of available features

Fig. 4: Average accuracy of the algorithms using all the available predictor types.

Fig. 5: The maximum accuracy of the algorithms using all the available predictor types.

is very high and the percentage of available relevant features is low, as is the case in the Hospital dataset. This clearly indicates that Clustering performs the clustering without taking the outcome into account at all, while hybrid version manage to get better results since they finalize the feature selection by trying to select those clusters that have the greatest impact to the outcome. One interesting finding from Figure 8 is also that the accuracy of mRMR can be improved from the accuracy achieved by using only *activity* predictors by using also either *2-gram* predictors or both *starter* and *finisher* predictors. This does not work for Clustering at least in BPIC datasets, since it always achieves the best average accuracy by using only *activity* predictors.

Since Clustering, Clustering hybrids and mRMR performed so well in this analysis that they could be incorporated without changes into some interactive process mining systems preferring under ten second response times when the size of the event log used for training is close to 1000 cases, we took Clustering and mRMR for closer inspection now emphasizing especially on the performance of different predictor types.

First, we analyzed the classification accuracy of both the algorithms separately for both the tested scenarios in BPIC14 event log having only 4000 cases. Figure 9 shows this information for Cluster algorithm and Figure 10 for mRMR algorithm. From these results, it can be seen that there is a lot of variation be-

Fig. 6: Average accuracy of the algorithms using only *activity* predictor type.

Fig. 7: Average accuracy of the algorithms using both *activity* and *2-gram* predictor types.

Fig. 8: Average accuracy by algorithm and by predictor type for all the tested BPIC event logs.

tween the two tested scenarios in the BPIC14 dataset. In both the cases and in almost all the predictor types, predicting case duration produced clearly worse results than in the categorization scenario. It seems that mRMR algorithm was not able to get any additional accuracy into its predictions by including any additional predictor types on top of *activity* predictors, whereas Cluster algorithm managed in the categorization case to get better accuracy when adding *order* type predictors in addition to *activities*.

It should be also noted that the time required for building a classification model with a feature selection algorithm selecting 10 features was only about 1% - 3% of the total time required when building the model with all the 1864

Fig. 9: Average accuracy for Cluster algorithm by predictor types separately for each scenario for 4k case sample in BPIC14.

Fig. 10: Average accuracy for mRMR algorithm by predictor types separately for each scenario for 4k case sample in BPIC14.

tested structural features without any feature selection. When using 3000 cases to build a model, the total measured time difference in the test system was about 250 seconds. During this time, it would have been possible to run the clustering feature selection several times. Thus, it is clear that having a feature selection performed before model building, at least when GBM is used, is essential when trying to improve the time required for model building.

As additional notes about the Dutch academic hospital dataset [32], the tests performed in this dataset indicated that it is absolutely critical to perform some kind of feature selection before training the model since building the model without any selection failed when attempting to use all the 79571 ordering features valid in this event log. The amount of features was actually so huge that some of the tested algorithms failed completely when trying to select the relevant features.

Thus, in order to fulfill all the requirements, we have for the algorithm, based on the performed tests, the best option from the selected set of algorithms is Cluster based feature selection with only *activity* type predictors with values being the occurrence counts of *activities* within a *case*. In our experiments, it gave the best overall trade-off in performance considering mutual information, classification accuracy and response time. For best possible response times, mRMR might be a better alternative with only *activity* predictors. Cluster Fisher is the recommended hybrid selection algorithm that performs in average slightly worse than plain Cluster when only *activity* type predictors are used, but when more features are available, it seems to perform slightly better.

6 Related Work

In addition to the related work related to more complex structural features referred already in Section 2.3, several papers have been written on the subject of applying data mining and machine learning techniques into predicting outcomes of the business processes. In [4], the authors present a framework that is capable of automatically detecting "signatures" that can be used to discriminate between desired and undesired behavior within traces both seen or unseen. These signatures are essentially combinations of structural features similar to those described in Section 2.2. This paper does not in itself specify any automatic feature selection method. Instead, the user is required to specify manually the desired activity sequence patterns, referred to as sequence feature types. After this all the matching features will be used for signature detection. Thus, our research complements the research made in this paper by experimenting with different automatic feature selection methods that could be applied before this signature detection phase in order to reduce the computational cost of signature detection at the cost of some prediction accuracy.

In [24], the authors evaluate the accuracy achieved with three different classification methods using several combinations of more complex structural feature patterns discussed in Section 2.3 for three different datasets. As result, they find out that just having Activity frequencies often yield, if not the best, then at least almost as good results as the best tested structural feature pattern combination. This finding is visible also in our tests as shown in Figures 9 and 10.

In [11], the authors present a framework for predicting outcomes of user specified predicates for running cases using clustering based on activity sequence prefixes and classification using attributes associated to events. In [20], the authors have assessed the benefits of including case and event attributes when performing predictions based on sequences of activities. In [28], the authors present a predictive process monitoring framework that is also able to mine unstructured textual information embedded into attributes related to events. In [6], the authors propose a recommendation system that automatically determines the risk that a fault will occur if the input the user is giving to the system will be used to carry on a process instance.

Until now there has not been systematic testing of applying automatic feature selection algorithms after selecting structural feature patterns to use and before building models used for classification. The aim of this feature selection is to minimize the computational cost of the building of classification models. Creating such an approach is crucial for obtaining predictions with interactive response time requirements. This is the primary contribution of this paper.

7 Conclusions

In this paper, we have designed a system for assessing the performance and response times of selected feature selection algorithms specifically tuned into the context of selecting structural features extracted from properties of sequences

of activities derived from event logs. Using this system, we tested nine feature selection techniques.

Each algorithm was tested first using a publicly available real-life Rabobank Group ICT dataset and tuned for two different classification use cases: Predicting whether the duration of a case is longer than seven days, and classifying whether a case is of type request for information. Most of the tests were also run using two different sample sizes out of the full dataset. For sanity checking and benchmarking purposes, we also added test runs without any feature selection and also with randomized feature selection. Finally, for a smaller sub-set of algorithms, we ran additional duration based classification tests on four other publicly available data sets.

We also proposed a rough categorization method for some of the types of structural features that can be extracted from event logs. In this paper, we selected four of these types for closer inspection.

As summary for all the tests and their results, it can be clearly seen that structural features can provide additional means for improving the precision to classifications made for cases in event logs. When the number of selected features is small, the most efficient source of features is activities. Increasing the number of features improves the classification accuracy, but also while doing so, best results are achieved by adding features from other structural feature types such as event type orderings into the set of structural features from which the feature selection is made. However, there is a drawback that having a bigger pool of features to select from makes creating classification models as well as the feature selection slower. As our goal was also to find an algorithm that could perform feature selection and classification with interactive response times using the sample sizes used in this paper, we found out that only one feature selection algorithm of the tested algorithms provided both the speed and accuracy required for the task.

According to the tests, the most consistently well performing algorithm was *Cluster* algorithm we developed for this paper which first used *k-means* algorithm for clustering features into the desired number of clusters by having cases as clustering dimensions, after which the features closest to the center of each cluster were selected as the selected features.

This algorithm performed especially well when the available structural features did not have many redundant features that did not have an effect into the final classification or when the available training data did not cover very well all the available features. In those cases, hybrid algorithms, such as one mixing both the *Clustering* and *Fisher* scoring, seem to outperform *Cluster* algorithm. Our partially self implemented *Cluster* algorithm was not, especially with larger number of cases, as fast as another quite well performing *mRMR* component. Both of these algorithms lose, in average accuracy in some data sets, to *Recursive Feature Selection*, but due to its decades slower response time, it can not be recommended due to the interactive usage requirements set in this paper. For computer assisted root cause analysis, in addition to providing the list of the most important features, *Cluster* algorithm provides also a mapping from

each of the original structural features to one selected feature that most closely resembles the original feature in the set of selected features.

All the raw information gathered from over 2300 successfully performed test runs can be found in the support materials [16]. This raw information, some of which was not discussed nor explored in this paper in detail include also: Measured accuracies, mutual information scores, computation times, selected features and confusion matrices.

8 Acknowledgements

We want to thank QPR Software Plc for funding our research. Financial support of Academy of Finland projects 139402 and 277522 is acknowledged.

References

1. K. P. Bennett and C. Campbell. Support vector machines: Hype or hallelujah? *SIGKDD Explorations*, 2(2):1–13, 2000.
2. R. J. C. Bose and W. M. van der Aalst. Trace clustering based on conserved patterns: Towards achieving better process models. In *Business Process Management Workshops*, volume 43, pages 170–181. Springer, 2009.
3. R. P. J. C. Bose and W. M. P. van der Aalst. Abstractions in process mining: A taxonomy of patterns. In U. Dayal, J. Eder, J. Koehler, and H. A. Reijers, editors, *Business Process Management, 7th International Conference, BPM 2009, Ulm, Germany, September 8–10, 2009. Proceedings*, volume 5701 of *Lecture Notes in Computer Science*, pages 159–175. Springer, 2009.
4. R. P. J. C. Bose and W. M. P. van der Aalst. Discovering signature patterns from event logs. In *IEEE Symposium on Computational Intelligence and Data Mining, CIDM 2013, Singapore, 16–19 April, 2013*, pages 111–118. IEEE, 2013.
5. H. Cheng, X. Yan, J. Han, and C.-W. Hsu. Discriminative frequent pattern analysis for effective classification. In *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, pages 716–725. IEEE, 2007.
6. R. Conforti, M. de Leoni, M. L. Rosa, W. M. P. van der Aalst, and A. H. M. ter Hofstede. A recommendation system for predicting risks across multiple business process instances. *Decision Support Systems*, 69:1–19, 2015.
7. T. F. Covoos, E. R. Hruschka, L. N. de Castro, and Á. M. Santos. A cluster-based feature selection approach. In E. Corchado, X. Wu, E. Oja, Á. Herrero, and B. Baruaque, editors, *Hybrid Artificial Intelligence Systems, 4th International Conference, HAIS 2009, Salamanca, Spain, June 10–12, 2009. Proceedings*, volume 5572 of *Lecture Notes in Computer Science*, pages 169–176. Springer, 2009.
8. 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.
9. R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification, 2nd Edition*. Wiley, 2001.
10. J. Evermann, J. Rehse, and P. Fettke. Predicting process behaviour using deep learning. *CoRR*, abs/1612.04600, 2016.
11. C. D. Francescomarino, M. Dumas, F. M. Maggi, and I. Teinmaa. Clustering-based predictive process monitoring. *CoRR*, abs/1506.01428, 2015.

12. 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.
13. D. Gusfield. *Algorithms on strings, trees and sequences: computer science and computational biology*. Cambridge university press, 1997.
14. J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann, 2000.
15. 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.
16. M. Hinkka. Support materials for articles. <https://github.com/mhinkka/articles>, 2017. Online; accessed 13-March-2017.
17. M. Hinkka, T. Lehto, and K. Heljanko. Assessing big data SQL frameworks for analyzing event logs. In *24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2016, Heraklion, Crete, Greece, February 17-19, 2016*, pages 101–108. IEEE Computer Society, 2016.
18. S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
19. 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.
20. A. Leontjeva, R. Conforti, C. D. Francescomarino, M. Dumas, and F. M. Maggi. Complex symbolic sequence encodings for predictive monitoring of business processes. In H. R. Motahari-Nezhad, J. Recker, and M. Weidlich, editors, *Business Process Management - 13th International Conference, BPM 2015, Innsbruck, Austria, August 31 - September 3, 2015, Proceedings*, volume 9253 of *Lecture Notes in Computer Science*, pages 297–313. Springer, 2015.
21. A. Liaw and M. Wiener. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.
22. D. Lo, H. Cheng, J. Han, S.-C. Khoo, and C. Sun. Classification of software behaviors for failure detection: a discriminative pattern mining approach. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 557–566. ACM, 2009.
23. P. E. Meyer. Information-theoretic variable selection and network inference from microarray data. *Ph. D. Thesis. Université Libre de Bruxelles*, 2008.
24. H. Nguyen, M. Dumas, M. L. Rosa, F. M. Maggi, and S. Suriadi. Mining business process deviance: A quest for accuracy. In R. Meersman, H. Panetto, T. S. Dillon, M. Missikoff, L. Liu, O. Pastor, A. Cuzzocrea, and T. K. Sellis, editors, *On the Move to Meaningful Internet Systems: OTM 2014 Conferences - Confederated International Conferences: CoopIS, and ODBASE 2014, Amantea, Italy, October 27-31, 2014, Proceedings*, volume 8841 of *Lecture Notes in Computer Science*, pages 436–445. Springer, 2014.
25. J. O. Ogutu, H.-P. Piepho, and T. Schulz-Streeck. A comparison of random forests, boosting and support vector machines for genomic selection. *BMC Proceedings*, 5(3):S11, 2011.
26. QPR Software Plc. *QPR Software to Offer Business Process optimization with Automated Business Process Discovery Software QPR ProcessAnalyzer*. 2011.
27. W. Steeman. Bpi challenge 2013, incidents, 2013.

28. I. Teinemaa, M. Dumas, F. M. Maggi, and C. D. Francescomarino. Predictive business process monitoring with structured and unstructured data. In M. L. Rosa, P. Loos, and O. Pastor, editors, *Business Process Management - 14th International Conference, BPM 2016, Rio de Janeiro, Brazil, September 18-22, 2016. Proceedings*, volume 9850 of *Lecture Notes in Computer Science*, pages 401–417. Springer, 2016.
29. K. Thompson. Programming techniques: Regular expression search algorithm. *Communications of the ACM*, 11(6):419–422, 1968.
30. R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
31. W. M. P. van der Aalst. *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011.
32. B. Van Dongen. Real-life event logs - hospital log, 2011.
33. B. Van Dongen. Bpi challenge 2012, 2012.
34. B. Van Dongen. Bpi challenge 2014, 2014.
35. B. Van Dongen. Bpi challenge 2017, 2017.
36. J. Weston, S. Mukherjee, O. Chapelle, M. Pontil, T. A. Poggio, and V. Vapnik. Feature selection for SVMs. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13, Papers from Neural Information Processing Systems (NIPS) 2000, Denver, CO, USA*, pages 668–674. MIT Press, 2000.
37. Y. Zeng, J. Luo, and S. Lin. Classification using Markov blanket for feature selection. In *The 2009 IEEE International Conference on Granular Computing, GrC 2009, Lushan Mountain, Nanchang, China, 17-19 August 2009*, pages 743–747. IEEE Computer Society, 2009.