

# Detecting and Explaining Drifts in Yearly Grant Applications

Stephen Pauwels and Toon Calders

University of Antwerp  
Antwerp, Belgium

**Abstract.** During the lifetime of a Business Process changes can be made to the workflow, the required resources, required documents, . . . . Different traces from the same Business Process within a single log file can thus differ substantially due to these changes. We propose a method that is able to detect concept drift in multivariate log files with a dozen attributes. We test our approach on the BPI Challenge 2018 data consisting of applications for EU direct payment from farmers in Germany where we use it to detect Concept Drift. In contrast to other methods our algorithm does not require the manual selection of the features used to detect drift. Our method first creates a model that captures the relations between attributes and between events of different time steps. This model is then used to score every event and trace. These scores can be used to detect outlying cases and concept drift. Thanks to the decomposability of the score we are able to perform detailed root-cause analysis.

## 1 Introduction

Log files contain valuable information with regard to the execution of different Business Processes in an organization [1]. By applying different mining techniques on these logs we can extract the workflow, have a better understanding about the resource management, show the differences between departments, . . . . Most Process Mining techniques require a stable state in order to perform well. Business Processes, however, are prone to different kind of changes, both expected, documented changes as well as unexpected changes may occur. These changes can be induced by new company policies, new regulations, new activities, . . . . Concept Drift Detection tries to identify the possible change points (I.e. the points in time where changes to the process happen) and tries to explain what caused the drift. After detecting these points it is possible to split the log in smaller logs that do represent the same execution of a Business Process without changes and use standard process mining techniques to further examine the different processes [2].

In [3] we introduced a model that is able to score traces in order to test for anomalies within the log file. In this report we show that the same model can be used to detect drifts within a Process log. Rather than providing a fully-automated solution we offer visual aids that people can use to analyze data.

Data that is properly visualized can easily be investigated by people with often better results than both fully-automated systems or fully-manual systems [4].

We use the BPI Challenge 2018 dataset [5] to show the capabilities of our drift detection algorithm. The dataset consists of applications for EU direct payments from farmers in Germany. Due to changes in regulations the process changes every year making it a suitable dataset for testing our method.

The goal of this report is to provide visuals that allow us to easily detect concept drifts and helps in the investigation of the main causes of these drift. These visual aids are designed to be easily understood by both business experts and technical people.

Our report is structured as follows: first we describe existing work in the field of concept drift in Section 2. The dataset itself is described in Section 3 together with the few preprocessing steps we have made. In Section 4 we explain the model used for scoring the events and we introduce a general workflow that uses these scores to determine concept drift. Section 5 shows the results and analysis of the BPIC 2018 data, showing that our method does provide useful tools for detecting concept drift.

## 2 Related Work

In Bose et al. [6] a concept drift detection method is proposed that uses statistical tests over feature vectors. These feature vectors are based on the order in which activities occur in the workflow. They record for example how many activities always, sometimes, or never occur after a given activity. Other measures can be used to add extra features that may incorporate the other attributes in the log. The method does require manual selection of the exact features to be used in the statistical tests.

Maaradji et al. [7] propose an automated method for detecting drifts. They focus only on the activities within the process, and transform the traces in the log into partial order runs. These runs summarize the different workflows possible for the process. In order to detect drifts they search for statistically significant changes in the distribution of the most recent runs.

Accorsi et al. [8] propose to use clustering of traces to detect drifts in the process. The clustering is done by using the number of intermediate activities between activities within different windows.

Carmona et al. [9] create an abstract representation of the process in the form of a polyhedron, which is afterwards used to test how well the representation includes the traces in from log. When the traces differ from the representation a drift is detected. To find other drifts the entire method needs to be repeated. This method is able to work in an online environment where we can check for drifts during the execution of the process, whereas other methods require a log file containing completed traces.

case	activity	doctype	subprocess	success	year	department
8b99873a6136cfa6	mail income	Payment application	Application	true	2015	e7
8b99873a6136cfa6	mail valid	Payment application	Application	true	2015	e7
8b99873a6136cfa6	mail valid	Entitlement application	Main	true	2015	e7
8b99873a6136cfa6	initialize	Parcel document	Main	true	2015	e7
8b99873a6136cfa6	initialize	Control summary	Main	true	2015	e7
8b99873a6136cfa6	begin editing	Control summary	Main	true	2015	e7
8b99873a6136cfa6	finish editing	Control summary	Main	true	2015	e7
⋮	⋮	⋮	⋮	⋮	⋮	⋮
8b99873a6136cfa6	begin payment	Payment application	Application	true	2015	e7
8b99873a6136cfa6	finish payment	Payment application	Application	true	2015	e7

Table 1: Example trace from the dataset

### 3 Dataset

The dataset [5] consists of applications for EU agricultural grants from the last three years. It contains both attributes that were known when the application was submitted and attributes that contain the outcome of the application after the trace finished. We call the former type of attributes the *basic attributes* and the latter type the *derived attributes*.

The basic attributes contain information about the activity, the applicant, his parcels, the document used, etc, whereas the derived attributes contain whether or not the grant was assigned, possible penalties on the requested amount and the actual amount granted.

Unless otherwise stated we only use the basic attributes for building and testing the model, making it possible to create a model that can be used to test traces while they are still going on and detect drift as soon as possible.

The data consists of 2,514,266 events grouped in 43,809 traces each representing a single application. The shortest trace contains only 24 events, the longest 2,973 and on average there are 57 events per trace.

*Example 1.* A typical snippet of an example trace can be found in Table 1, which is taken from the BPIC dataset. Some attributes remain the same within a case (I.e. case, year and department) while the others may have different values.

## 4 Concept Drift Detection

In this section we first explain the model we use for scoring the events and introduce a workflow for detecting drifts and finding differences in the data.

### 4.1 The Model

Our model (extended Dynamic Bayesian Network or eDBN) is based on Dynamic Bayesian Networks (DBNs) [10], an extension of the well-known Bayesian Networks. These networks are capable of capturing Conditional Dependencies between various variables both within one single time step as over different time

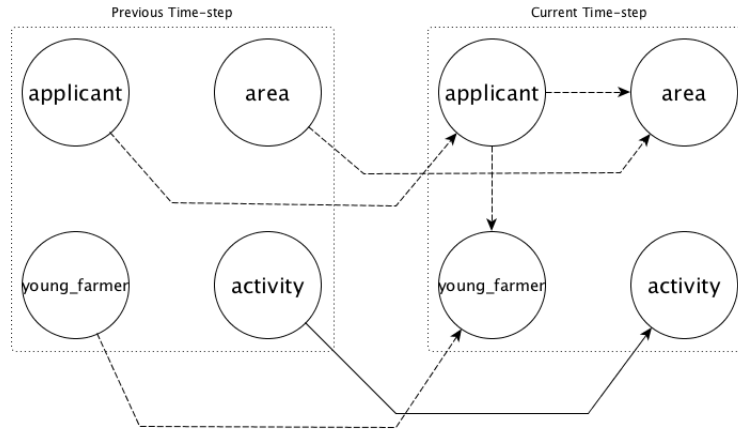


Fig.1: eDBN example with conditional (full) and functional dependencies (dashed) based on the dataset.

steps. As we have indicated in [3], the DBNs contain useful elements but lack some expressiveness in order to be useful in the context of Business Process logs. We introduce a second type of dependency, the Functional Dependency (FD), that describes mappings of values from one attribute to another.

*Example 2.* An example structure based on the dataset for an eDBN is given in Figure 1. In the dataset every applicant has the same value for area and young\_farmer within a single trace. Therefore there exist a Functional Dependency from applicant to both area and young\_farmer in the current time step. And since these values do not change between events in a trace we also link each of these attributes from the previous time step to the current time step.

The attribute activity in the previous time step has a Conditional Dependency to activity in the current time step because the current activity depends on the previous. Branches in the workflow may occur so multiple activities can follow a certain activity, giving the explanation of why this is a CD and not a FD.

After finding the different dependencies within the model we have to learn the parameters in order to be able to score events. These parameters consist of all the Conditional Probability Tables, the Functional Dependency Mappings and the probability of encountering new values or new relations as explained in [3]. We add these *new values* and *new relations* measures to the model to be able to handle new (unseen) values and relations. DBNs return a probability of 0 when encountering a value that was not present in the training dataset, but our log files do allow the occurrence of new values (a new employee entering the company, ...) in the logs. Using these measures we can take this into account without severely penalizing the score for an event in case of a justified new value.

Parent	Value	Probability
mail income	mail valid	1.0
initialize	initialize	0.6
	begin editing	0.3
	performed	0.1
begin editing	finish editing	0.7
	calculate	0.3

Table 2: CPT for relation between activity\_previous and activity

*Example 3.* In order to learn all the parameters for the model in Figure 1 we start by learning the CPT for the relation  $\text{activity\_previous} \rightarrow \text{activity}$ . An example CPT is shown in Table 2. For the different FDs we store the mappings that are present between values of the attributes in the training data. The FDT of the relation  $\text{applicant} \rightarrow \text{young\_farmer}$  will link all applicants to the fact whether they are young farmers or not. In order to determine the new values and new relations rate, we calculate for every attribute in the current time step the probability of encountering a new value or relation.

The score for an event  $e$  given the model can be calculated by:

$$\text{Score}(e) = \prod_{A \in \{\text{attributes of } e\}} P(e.A | \text{Parents}(A))$$

with

$$\begin{aligned} P(e.A | \text{Parents}(A)) = & \text{value}_A(e.A) \\ & \times CPT(e.A | \text{Parents}(A)) \\ & \times FDT(e.A) \end{aligned}$$

Where  $\text{value}_A$  indicates the score for the value itself (does it occur in the training set or is it a new value),  $CPT$  returns the probability for a particular value given the values of the attributes it depends on, and  $FDT$  returns a score based on whether or not the Functional Dependencies were broken or not.

The score for an event is built from the contributions of every attribute equally. This helps us to split the score into partial scores were we get the contribution of every attribute by itself. Furthermore the partial score itself can also be decomposed into its components giving an even better and more detailed way of determining the root cause of changes and differences.

Thanks to the decomposability of the score assigned to an event we created an expressive model that is able to easily give an explanation about how the score was computed. Furthermore thanks to how we composed the score, equal events and traces get approximately the same score making it easy to group traces together that show the same behavior and making it possible to detect drifts and differences.




Fig. 2: Workflow used to score the traces in the log.

## 4.2 Detecting Drifts

First we select the traces in the log we want to use as the reference set, we call these traces the *training dataset*. This training dataset forms the baseline where we will compare all other traces too, we refer to these traces as the *test dataset*. When detecting concept drift we always select the first  $n$  traces as our training dataset and compare them with the other traces. Traces that occur after a drift are supposed to receive a different score than the ones before the drift. The basic workflow of building the model and calculating the scores is given in Figure 2. We can also use this workflow with different kind of training data (E.g. only selecting events that satisfy a certain condition) to find different types of drifts or differences.

In order to be able to compare traces we assign a score to every trace based on the model we created using the training data. This score is calculated as the mean score for every event in the trace. We use this score in contrast to the  $n$ -th root of the product of all event scores (where  $n$  is the number of events in the trace) used in [3] as we are only looking to characterize a trace using this score. When we would use the original formula to compute the score for a trace, a single violation of a FD can result in the total score being 0. When looking for anomalies this is an important property as these FDs are supposed not to be broken. Also it would be impossible to find out if only one attribute had a partial score of 0 or multiple. By using the mean score we ensure that we can see this difference.

**Definition 1.** *The score for a trace is computed as:  $Score_{trace}(t) = \frac{\sum_{e \in t} Score(e)}{|t|}$ . With  $|t|$  equal to the number of events in a trace. When there is no confusion possible between the score of an event and a trace we omit the subscript.*

After we have calculated the scores for every trace we can sort them according to timestamp. Traces that represent the same Business Process (without changes in the workflow or attributes) are supposed to have scores closely related to each other. After a drift occurred the scores for the traces change due to the changes that were made to the process itself. The distribution of scores after such a change point is different than before the drift.

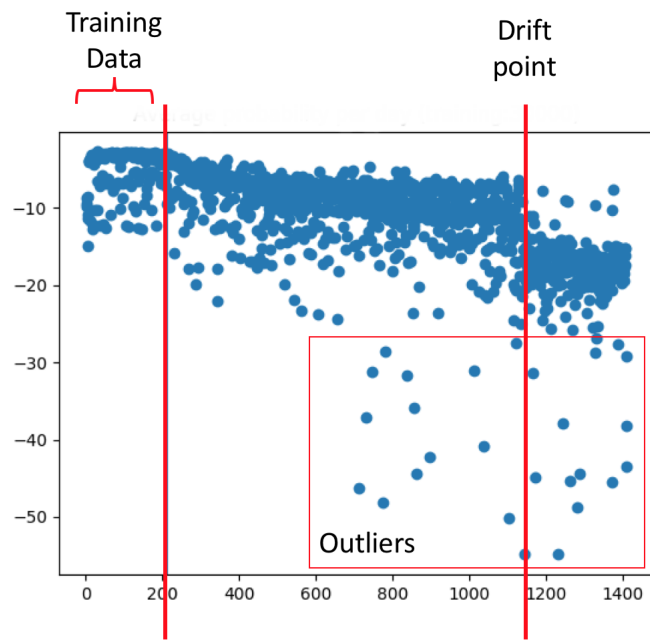


Fig. 3: Example of scores for different traces picked from a random dataset. The first vertical red line indicates the difference between the traces used for training and testing, the second red line indicates the drift.

**Trace-score plot** The first visual aid we introduce is the *trace-score plot* which can be seen in Figure 3. The score of every trace is plotted on this graph, giving us a first idea of the different variations present in the log. We use a logarithmic scale for plotting the different scores because the scores of the traces are very small as they are based on probabilities. The logarithmic scale helps us better visualize these differences. The x axis in this plot is time, mostly indicated by the incremental number of the trace in the log. This plot helps us to look at how the scores are distributed over time. A first phenomenon that we can observe is that the overall distribution of scores changes after a certain point in time, these changes represent drifts in the business process. A second interesting observation that can be made is that of outlying traces that deviate from the main distribution of scores. We can then decide to investigate these outlying traces in more detail, as checking all traces in detail is often not feasible.




Fig. 4: Sliding window of size 10 with the *reference window* (yellow) and *test window* (green).

**Drift plot** To better visualize the changes in the scores we use the two sample Kolmogorov-Smirnov statistical test for checking the hypotheses that the distributions in both samples are similar. The test returns a p-value between 0 and 1, where 1 stands for accepting the hypothesis that both distributions are equal and 0 rejecting it. To find the drift points in the entire log we use a sliding window technique. The first half of the window is used as the *reference window* and the second half as the *testing window*, as can be seen in Figure 4. We apply the two sample Kolmogorov-Smirnov test on both halves of the window to find possible drift points. The size of the window has to be manually set. A smaller window is more sensitive for finding the smaller drifts but is also more likely to return false positives and a larger window is less sensitive but might return more false negatives. Since we have all the scores calculated and they do not depend on the window size, it is doable to check some different window sizes and compare the plots they return. We can plot all the resulting p-values from the statistical test in order to visually see drift points as shown in Figure 5. Again we make use of a logarithmic scale since the differences are otherwise not visible. We will refer to this type of plots as the *drift plot*.




Fig. 5: Plotting all the p-values obtained by the Kolmogorov-Smirnov test. We can see two dips in this plot. The first dip indicates the change point between the training data and test data. The second dip indicates the real drift point.



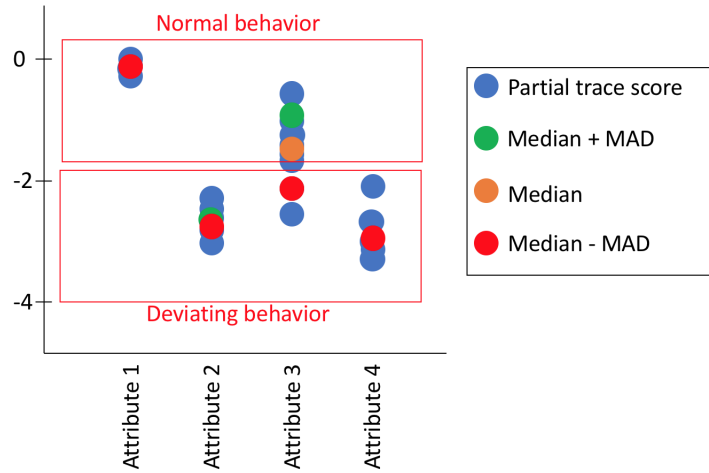


Fig. 6: Example attribute-density plot for 4 attributes.

**Attribute-density plot** To find the root cause for the drifts we found, we take a look at the partial scores of the events, instead of the total event scores themselves. We accumulate the scores according to the attributes. We can plot all the partial trace scores in an *attribute-density plot* as can be seen in Figure 6, as before we use a logarithmic scale. We add the median and median absolute deviation [13] in order to get a better insight in the density and distribution of scores for a single attribute. In contrast to the other types of plots these plots characterizes the data in terms of the different attributes. To be able to use these plots for comparing different segments of a log file we first select a reference segment which is used to learn the model used for calculating the partial event scores of all segments. Every segment can be visualized in a different attribute-density plot for comparison. Any technique can be used to help users find differences between these plots. In the remainder of this paper we simply plotted all the median values on the same plot in order to determine differences.

**Towards a general workflow** Figure 7 shows the relations of the three different plots we introduced. The trace-score plot and drift plot are used to divide the log file into different segments where each segment contains traces with similar behavior. This is done using the temporal aspect of the log file and our sliding window technique. To better understand the differences and equalities among the segments we first take a reference segment (or subset of a segment) that is used to train the model. This model represents the expected traces in the log file based on the reference segment. It can then be applied to every other segment (including the reference segment) to get the partial scores for the traces in that segments. These partial scores can be plotted on an attribute-density plot, mak-

ing it possible to find the root cause for the difference between the segments. This general workflow for determining the root cause is shown in Figure 8. All of the analysis performed in Section 5 are based on this general workflow.

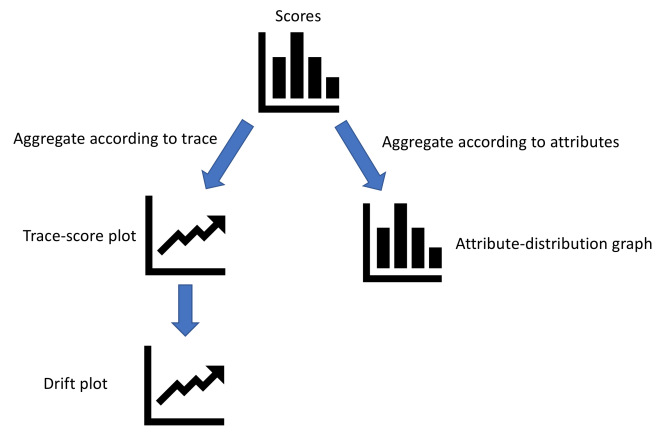


Fig. 7: Workflow of building the different types of plots and showing how they relate to each other.




Fig. 8: General workflow for describing changes between segments of the log file.

## 5 Analyzing the data

In this section we analyze the BPI Challenge 2018 data to show the effectiveness of our proposed workflow. And show that it is able to find interesting insights in the given data.

### 5.1 Preprocessing of the data

Although our method does not require any preprocessing we perform a little bit of preprocessing to avoid overly complex models without any added value over more simple models. We perform the following steps:

1. Since our implementation works with CSV-files we first use the Prom Tool [12] to convert the given XES-files.
2. We filter out all the *derived attributes* as we are only interested in the attributes that really depend on the activity and not the outcome of the process.
3. Next we filter out the attributes *eventid*, *event\_identify\_id* and *identity\_id*, since their respective values are unique for every event and would thus form Functional Dependencies with every other attribute in the event.
4. Because we are primarily looking to test our model for concept drift detection, we also omit the *year* attribute as the process changes every year and we do not want our model to be biased towards this attribute.

All of these preprocessing steps were repeated for every experiment conducted on the data, unless explicitly stated otherwise. Important to note is that our preprocessing steps did not alter the nature of the log or any of the processes in the log.

### 5.2 Finding drift points

To look for any drift points in the data we first learn a model according to the log and use this to test all the traces in the following way:

1. As a first step we read the data and perform the preprocessing step.
2. To learn our model we use the first 30,000 events from the log as the training dataset. We use the same dataset for training the model structure as for training all the parameters.
3. For testing we use all the preprocessed data, including the training set.
4. We calculate the score for every event in the log given the model.
5. We calculate the mean score for every trace. This is the *trace score* that we use to further visualize the drifts in the data.
6. As the first visual step we plot all the scores from the different traces in a trace-score plot, according to the timestamp of the first event of every trace. This is shown in Figure 9.
7. Looking at this plot we can already get some insights about the data:





Fig. 9: Annotated trace-score plot.

- The first cluster of traces indicates the traces used for training the model (we can ignore this cluster).
- We can clearly see a difference in scores around trace 14,000. Which is a clear indication for a drift point.
- A second possible drift point can be seen around trace 29,000, using the drift log in the next step will make it possible to really determine if this is a drift point.
- The traces cluster together in two distinct groups. Indicating that the data has mainly two types of traces, each with their own characteristics.
- We can see 4 outlying traces before the first drift point.


These insights are visually highlighted in Figure 9 and further explored in the next subsections.

- Next we perform the Kolmogorov-Smirnov test with a sliding window over the scores. The resulting drift plots for different window sizes can be found in Figure 10. These plots clearly show the 2 (3 including the training data) drifts occurring in the data, confirming our initial thoughts.


We show in the next subsections that our model is able to explore our findings even further in order to give an explanation for the observed behavior.



(a) Windows size of 400



(b) Windows size of 800



(c) Windows size of 1600

Fig. 10: Drift plots containing the p-values for different window sizes.

### 5.3 Inspecting the different attributes

In order to find the root cause of the different drifts in the log we use the attribute-density plot to visually represent the characteristics of a process. In order to do so we perform the following steps:


1. Learn the model as before.
2. From every event get the partial scores of every attribute and combine these for the different attributes (in contrast to accumulating the total score for an event as we did before).
3. Using these partial scores we create multiple attribute-density plots, one plot for every period between drifts. These plots can be used to compare the different segments with each other to determine the root cause for the drifts.
4. Figure 11 shows the attribute-density plots for the three different years.
5. In order to better compare these plots we plotted all the median values in one plot. This plot can be found in Figure 12.

Using the plots in Figure 11 and 12 we can see three big changes between the first year, and and the other two years.


- The first difference is found for the attribute *area*, this can be explained by the fact that farmers acquire new land or sell parts of there lands. It is also possible that the used technique for discretizing influences this attribute, as the different bins can change over the years.
- The following two changes are probably closely related to each other. Starting in the second year, a new document type was used for the applications, this explains the large differences for this attribute between the different years. And it is normal that due to the changes to the documents, the subprocesses also change.
- Between year 2 and 3 we have the largest difference in the subprocess attribute but less in the other attributes, indicating only some changes in the different subprocesses have occurred.

We have thus found two drift points: one at the beginning of each new year. The difference between year 1 and 2 was large as there is a significant change in the area but most importantly a new document type was used that had its consequences on the used subprocesses. It is also possible that this new document type influenced the way the area had to be filled in. Between year 2 and 3 the difference is less large, this was already visible in the trace-score plot. The most important change between these years was the way the subprocesses where used.


In order to find a more detailed analysis of why this attribute changed that much, other tools and methods can be used to give more detailed explanations and insights.



(a) Year 1



(b) Year 2



(c) Year 3

Fig. 11: The attribute-density plots for the different years, where we indicated the most changing attributes.




Fig. 12: Plot showing all the median values for every year.

#### 5.4 Comparing Departments with each other


The method we used in the previous sections can be generalized in order to look for differences between groups of traces. We now take a look at possible differences between departments. As a first experiment we use the same training data as before and plot an attribute-density plot for every single department. These plots can be found in Figure 13. When comparing these departments with each other we cannot see a clear difference in the scores for the attributes.

In order to better test differences between departments we train our model with events from department 1 and use the other departments as separate test segments. This approach can be generalized for testing any kind of differences. We take a reference dataset for training and use the resulting model for testing other segments of traces. As before we show the attribute-density plots created for every departments in Figure 14. We were looking for differences in activities and subprocesses, but when comparing the plots we cannot see a substantial difference between any attribute.


#### 5.5 Difference between types of traces

In the trace-score plot we saw that within every year we have two types of traces. In this section we will take a closer look by decomposing the score even further. We only investigate traces that occur in the first year, as otherwise other drifts and differences might interfere. The first step is to divide the traces in two parts: we have the first type (score between -8 en -10) and the second type (score between -10 and -12). These two types are also highlighted in Figure 9. To find a possible explanation for this difference we create an attribute-density plots for both types of traces. These plots can be found in Figure 15.







(a) Department 1



(b) Department 2




(c) Department 3




(d) Department 4


Fig. 13: Attribute-density plots for the different departments. When using the first 30,000 events as training data.



(a) Department 2



(b) Department 3



(c) Department 4

Fig. 14: Attribute-density plots after training with department 1.

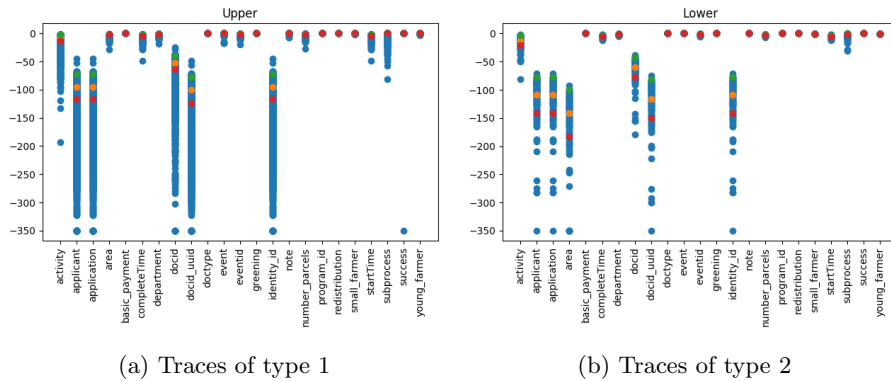


Fig. 15: Attribute-density plot for both types of traces

When comparing these two plots we can see that there is a big difference in the scores for the *area* attribute, which causes the split in two clusters. The score for the attribute *area* is further decomposable in different partial scores: one score is related to the value of the attribute itself. If it is already encountered in the data it returns 1, otherwise it returns the probability of encountering a new value for that attribute based on the training data. The next partial score we got from the Conditional Probability Table, but since the attribute has no Conditional Dependencies this returns 1. Last we get a score for every single Functional Dependency. The results of all these scores can be found in Figure 16 for both clusters.




Fig. 16: Detailed attribute-density plot for the attribute area for both types of traces.

The comparison between these two plots shows that the difference between the two types is caused by the fact that more unseen values for area are encountered in traces from type 2. Thanks to these plots we have a clear insight in which aspect of which attribute causes these deviations. Now we only have to analyze the area attribute into more detail without having to check all attributes present in the log. This effect is, again, due to the way the attribute area was discretized. In the first year we have 148 different values for the area attribute. Of these 148 values only 2 of these values did not appear in the training data, and causes the effect of the two types of traces. Since it is a result from how the data was initially preprocessed we cannot make any conclusions about the real process.

### 5.6 Investigating anomalous events in the first year

In the trace-score plot we saw that four dots were not part of any of the two clusters. In this section we use the attribute-density plot to quickly look for any significant differences between the median found in the attribute-density plot for the first year (in Figure 11). When comparing these traces to the expected scores we see that especially the *area* and *number\_parcels* attributes have scores that are unexpected when comparing them to the attribute-density plot for the first year. When we investigate these findings we see that these four traces are the only traces in the first year with both the area and number\_parcels equal to 0, thus indeed indicating outlying traces.




Fig. 17: Attribute-density plot for the outlier traces

## 6 Conclusion

In this report we show that the model we created in order to find anomalies in Business Process logs can easily be used to analyze Concept Drift and other differences in a log file. The way we score events and traces using this model is appropriate for generating more in-depth analysis of why certain traces do deviate. The analysis of the data proved the effectiveness of our method on a real-life dataset where we were able to split our data into the three different years present in the log by only using our drift points.

In this paper we were interested in answering two questions of the BPI Challenge 2018:

1. How can one characterize these differences as a particular instantiation of concept drift?
2. How can one characterize the differences between departments and is there indeed a relation?

To answer this first question we used our trace-score plot and drift plot in order to first detect the drift itself. As showed we did manage to accurately find the drift points in the data. In order to characterize these drifts we further investigated the segments in the log between these drifts. Using the attribute-density plot we were able to detect that, as expected, the document type and subprocesses caused the drifts between the years.

The approach of comparing different years with each other can be generalized in order to solve the second question. We used the first department for training our model and tested the other departments against this model. Besides from the normal deviations in attributes representing ids we did not note any difference between the departments.

As a last analysis we analyzed individual traces and clusters of traces. Again we used the attribute-density plots to find differences between either two clusters or between some events and the expected scores. This method proved to be useful as we were able to successfully explain certain differences in the log file.

The applications mentioned in this report are a starting point but give a clear picture of the possibilities of our eDBNs. The way our analysis was structured should give a insight in the logical steps one can follow to explore multivariate log files. In the future we plan to integrate all these possibilities within a single visual tool. In order to make a quick investigation of these log files more easy. We also plan on incorporating the temporal aspect within traces to be able to detect deviations in duration of a trace.

## References

1. van der Aalst, W.M.P.: Process Mining: Discovery, Conformance and Enhancement of Business Processes. 1st edn. Springer Publishing Company, Incorporated (2011)
2. Song, M., Günther, C.W., Van der Aalst, W.M.: Trace clustering in process mining. In: International Conference on Business Process Management, Springer (2008) 109–120

3. Pauwels, S., Calders, T.: Extending dynamic bayesian networks for anomaly detection in complex logs. arXiv preprint arXiv:1805.07107 (2018)
4. Keim, D.A., Mansmann, F., Schneidewind, J., Ziegler, H.: Challenges in visual data analysis. In: Information Visualization, 2006. IV 2006. Tenth International Conference on, IEEE (2006) 9–16
5. Van Dongen, B.F. (Boudewijn), Borchert, F. (Florian): Eindhoven university of technology. dataset. <https://doi.org/10.4121/uuid:3301445f-95e8-4ff0-98a4-901f1f204972> (2018)
6. Bose, R.J.C., van der Aalst, W.M., Žliobaitė, I., Pechenizkiy, M.: Handling concept drift in process mining. In: International Conference on Advanced Information Systems Engineering, Springer (2011) 391–405
7. Maaradji, A., Dumas, M., La Rosa, M., Ostovar, A.: Fast and accurate business process drift detection. In: International Conference on Business Process Management, Springer (2015) 406–422
8. Accorsi, R., Stocker, T.: Discovering workflow changes with time-based trace clustering. In: International Symposium on Data-Driven Process Discovery and Analysis, Springer (2011) 154–168
9. Carmona, J., Gavalda, R.: Online techniques for dealing with concept drift in process mining. In: International Symposium on Intelligent Data Analysis, Springer (2012) 90–102
10. Russell, S.J., Norvig, P.: Artificial intelligence: a modern approach (3rd edition) (2009)
11. Healy, M.: Understanding robust and exploratory data analysis. (1983)
12. Van Dongen, B.F., de Medeiros, A.K.A., Verbeek, H., Weijters, A., Van Der Aalst, W.M.: The prom framework: A new era in process mining tool support. In: International Conference on Application and Theory of Petri Nets, Springer (2005) 444–454
13. Leys, C., Ley, C., Klein, O., Bernard, P., Licata, L.: Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology* **49**(4) (2013) 764–766