

Decision Mining Revisited - Discovering Overlapping Rules

Felix Mannhardt^{1,2(✉)}, Massimiliano de Leoni¹, Hajo A. Reijers^{1,3},
and Wil M.P. van der Aalst¹

¹ Eindhoven University of Technology, Eindhoven, The Netherlands
{f.mannhardt,m.d.leoni,h.a.reijers,w.m.p.v.d.aalst}@tue.nl

² Lexmark Enterprise Software, Naarden, The Netherlands

³ Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

Abstract. Decision mining enriches process models with rules underlying decisions in processes using historical process execution data. Choices between multiple activities are specified through rules defined over process data. Existing decision mining methods focus on discovering mutually-exclusive rules, which only allow one out of multiple activities to be performed. These methods assume that decision making is fully deterministic, and all factors influencing decisions are recorded. In case the underlying decision rules are overlapping due to non-determinism or incomplete information, the rules returned by existing methods do not fit the recorded data well. This paper proposes a new technique to discover overlapping decision rules, which fit the recorded data better at the expense of precision, using decision tree learning techniques. An evaluation of the method on two real-life data sets confirms this trade off. Moreover, it shows that the method returns rules with better fitness and precision in under certain conditions.

Keywords: Decision mining · Process mining · Overlapping rules

1 Introduction

Organizations use process models representing their business processes for multiple reasons. Process models are used, for example, to document, specify, and analyze processes [1]. Generally, process models depict activities (i.e., units of work) and their dependencies in a graph representation, which specifies the order of activities in the process execution. During the execution of non-trivial processes, next to the ordering of activities, *decisions* between multiple alternative activities needs to be made. Those choices are explicitly modeled in process models as so-called *decision points*. A decision point specifies the alternatives available. An important challenge when using process models is to understand the decision that need to be made in a process, and the conditions under which certain alternative activities are performed. Awareness that modeling and analyzing decisions is key in process management is increasing. See for

The work of Dr. de Leoni has received funding from the European Community's Seventh Framework Program FP7 under grant agreement num. 603993 (CORE).

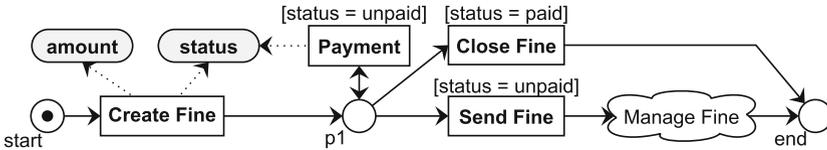


Fig. 1. Fragment of a process model taken from a road traffic fine management process with overlapping rules governing an exclusive choice based on the payment of the fine.

example the interest in the Decision Model and Notation (DMN) standard [2] supported by vendors such as Signavio and Camunda.

Process mining methods are able to discover process models with decision points by using event logs. Event logs contain information on performed activities (i.e., sequences of events) and are available in today's information systems [1]. *Decision mining* aims to discover the rules that are underlying those decisions. Those rules are determined using data recorded by information systems that support the process [3,4]. Events in event logs used for decision mining need to contain process data, which was available when the activity was performed (i.e., attributes). Take, for instance, the process model fragment depicted in Fig. 1, which shows a simplified fragment of a road-traffic fine management process [5]. After creating a fine notice (*Create Fine*) and recording the amount (*amount*), an exclusive choice between three alternatives has to be made. Either a payment is received (*Payment*), possibly in multiple installments, the fine is closed (*Close Fine*), or the police sends a fine notice (*Send Fine*) and the process continues with further management of the fine. Please note, that only one out of these three alternatives can be taken, which is different to an inclusive choice that allows the execution of multiple alternatives. The rules depicted in Fig. 1 drive this choice. Activity *Close Fine* can only be executed if the value of attribute *status* is *paid*. Thus, the process can only finish directly if the fine has been paid in a timely manner. Rules defined over the process data are an integral part of this process. Traditionally, decision mining methods [3,4,6] use decision tree learning techniques such as C4.5 [7] to determine rules governing the process execution based on event logs. For example, attribute values recorded for *status* and *amount* are used as *feature*, while the choice between *Payment*, *Send Fine*, *Close Fine* is used as *target class*. Then, mutually exclusive rules for each activity are built using the obtained decision tree, thus, the choice at the decision point is completely determined by the values of *status* and *amount*.

Existing decision mining techniques for exclusive choices rely on the strong assumption that the rules attached to the alternative activities of a exclusive choice need to be *mutually exclusive*. However, business rules are often non-deterministic and this “cannot be solved until the business rule is instantiated in a particular situation” [8]. This ambiguity can occur due to conflicting rules or missing contextual information [8]. For example, decisions taken by process workers may depend on contextual factors, which are not encoded in the system and, thus, not available in event logs. Moreover, even if those factors are encoded in the system event logs are often incomplete [9]. Without complete information in the event log the mutually exclusive rules underlying

decision-making cannot be discovered. Hence, this assumption is typically not met in reality. For example, the process model shown in Fig. 1, which is taken from a real-life process, contains overlapping rules for an exclusive choice. This means that on a decision point more than one of multiple activities may be executed under the same condition, i.e., when the same attribute values have been observed beforehand. In case the payment status is *unpaid*, the choice between Payment and Send Fine is deferred, i.e., an unpaid fine can be either paid directly or a notification is sent to the offender. The actual decision between either of those activities is not specified. It might depend on an unavailable contextual factor, e.g., some fines can be paid on-the-spot depending on the context. Please note that an exclusive choice with overlapping rules is different to an inclusive choice, which, for certain process executions, prescribes multiple alternatives activities to be executed. For example, in case of an inclusive choice, if status is *unpaid*, both payment and send fine are prescribed to occur. State of the art techniques [3,4,6] only use decision trees, and, hence, cannot discover overlapping rules on exclusive choices. For instance, current techniques fail to discover the rule in Fig. 1.

This paper proposes a technique that discovers *overlapping rules* in those cases that the underlying observations are characterized better by such rules. The technique is able to deliberately trade the precision of mutually-exclusive rules, i.e., only one alternative is possible, against fitness, i.e., the overlapping rules that are less often violated. In short, as in [3,4,6], our technique builds an initial decision tree based on observations from the event log. Then, for each decision tree leaf, the wrongly classified instances are used to learn a new decision tree leading to new rules. These new rules are used in disjunction with the initial rules yielding overlapping rules of the form $rule_1 \vee rule_2$. We evaluate our technique on two real-life data sets: an event log taken from a road traffic fine management process and an event log with pathways of patients in a hospital. The evaluation shows that our technique discovers overlapping rules in real-life data, and that those rules provide a better balance in terms of fitness and precision. For example, our technique discovers overlapping rules similar to the ones depicted in Fig. 1, whereas traditional method fail, e.g., to discover a rule for activity Payment.

As to the structure of this paper, we introduce necessary formalisms for process models and event logs (Sect. 2). Then, we present our discovery technique for process models with overlapping rules (Sect. 3). We evaluate our technique based on real-life event logs (Sect. 4). We discuss related work (Sect. 5), and conclude with a summary and sketched future work (Sect. 6).

2 Background

We present necessary preliminaries such as the formalism we use to represent process models and event logs, and criteria used to determine the quality of decision rules.

2.1 Process Model

Generally, our decision mining technique is independent of the formalism used to model the process, e.g., BPMN, UML-activity diagrams, EPCs or YAWL. We choose Data Petri Nets (DPN) [4] as modeling language because it has simple and clear semantics.

A **DPN** is a Petri net [10] that is extended with variables (i.e., data attributes). A Petri net consists of a set of **places**, a set of **transitions** and the flow relations that describe the bipartite graph between places and transitions. Transitions correspond to activities in the process. The state of a Petri net is defined by its marking. The marking of a Petri net assigns a number of tokens to each place. Executing a transition consumes one token from each of its input places and produces one token on each of its output places. We denote with output transitions of a place the set of transitions for which there is a directed edge from the place to those transitions. A transition can only be executed (fired) when there is at least one token in every input place.

A DPN is a Petri net with additional components: a set of **variables** defined over a universe of possible values; a set of **write operations** for each transition; a **guard expression** (guard) for each transition. Transitions update the values of variables through write operations. Furthermore, guards defined over the variables of the DPN further constrain when transitions may be executed. A transition in a DPN can be executed only if all its input places contain at least one token and the guard is satisfied by the current variable assignment. Unless specified otherwise we assume the guard `true` for each transition, i.e., the transition can fire regardless of the current variable assignment. The state of a DPN is defined by both the marking and the current values of all its variables. The behavior of a DPN corresponds to all sequences of transition firings starting from an initial state to any final state. The initial state is made of the initial marking, i.e., the initial number of tokens in each place and an empty set of variable values. Final states are all final markings. For sake of space we refer to [5] for a comprehensive introduction to DPNs.

Example 1. Figure 1 shows a simplified process in the DPN notation. The process starts with executing transition `Create Fine`, which writes attribute `status`. When executing transition `Create Fine` a token is removed from the place `source` and a new token is put in place p_1 . Now, a choice between the output transitions of p_1 , `Payment`, `Send Fine`, and `Close Fine`, has to be made. Therefore, place p_1 is called *decision point*. As there are guards placed on all three transitions their enablement depends on the current assignment of attribute `status`. For example, both transitions `Send Fine` and `Payment` can only be executed when `status` is *unpaid*. As the guards of `Send Fine` and `Payment` overlap, the choice between both transitions is non-deterministic. Assuming transition `Create Fine` assigned the value *paid* to `status`, i.e., the fine has been paid directly, then, only transition `Close Fine` can be executed and the process ends by reaching the final marking of the DPN.

2.2 Event Log

An event log stores information about the executed activities in a process [1]. Given a set of transitions, variables, and variable values, we define an **event log** \mathcal{E} as a collection of unique **events** [1, 11]. Each event $e \in \mathcal{E}$ is associated with a set $values(e)$ containing the latest values of all attributes recorded before the event occurred starting with an initial value. Moreover, each event refers to the execution of a transition $trans(e)$ and the set of variables that are written by event e is obtained by $writes(e)$. Table 1 shows an event log for the process model introduced in Fig. 1. Each row represents a unique

Table 1. Event Log \mathcal{E} with data attributes **status** and **amount**

Id	Case	Activity	Status	Amount				
e_1	1	Create Fine	unpaid	30	e_6	3	Create Fine	unpaid 30
e_2	1	Payment	unpaid	-	e_7	3	Payment	paid -
e_3	1	Send Fine	-	-	e_8	3	Close Fine	- -
e_4	2	Create Fine	unpaid	30
e_5	2	Send Fine	-	-				

recorded activity execution (i.e., event) together with the produced data (i.e., attributes). Special attributes like an **id**, the case identifier **case**, and the activity name **activity** are recorded for each event. The location of an event in the case, i.e., the order in which the events occurred, is uniquely identified through the **id** attribute.

Example 2. For the example event log in Table 1 we can determine the transition corresponding to event e_2 as $trans(e_2) = \text{Payment}$. Moreover, we can obtain the value of all attributes at the moment when e_2 occurred as $values(e_2) = ((\text{status} = \text{unpaid}), (\text{amount} = 30))$. Finally, the variables written by e_2 are $writes(e_2) = \{\text{status}\}$.

2.3 Quality Criteria - Fitness and Precision

We use two criteria to determine the quality of the guards defined on the output transitions of place p of a DPN given an event log \mathcal{E} : place fitness and place precision. We denote the set of events for transitions in the output transition of a place p with \mathcal{E}_p , i.e., $e \in \mathcal{E}$ iff $trans(e)$ is an output transition of p . We define the **place fitness** of a place based on the number of events in \mathcal{E}_p for which the guard is violated:

$$fitness_{\mathcal{E},p} = 1 - \frac{|\{e \in \mathcal{E}_p \mid \text{Guard of } trans(e) \text{ is violated}\}|}{|\mathcal{E}_p|}$$

The place fitness linearly decreases with an increase of the fraction of output transitions that fire violating the respective guard. The **place precision** of a place p is defined as is based on the possible behavior at p defined by the DPN, and actual behavior observed in \mathcal{E}_p . We use functions $pos_p : \mathcal{E} \rightarrow \mathbb{N}$ and $obs_p : \mathcal{E} \rightarrow \mathbb{N}$ returning the possible and observed executions of output transitions of place p before event e occurred. Work [12] describes how to obtain values for pos_p and obs_p given an event log and a DPN. Using pos_p and obs_p , we define the place precision as:

$$precision_{\mathcal{E},p} = \frac{\sum_{e \in \mathcal{E}_p} |obs_p(e)|}{\sum_{e \in \mathcal{E}_p} |pos_p(e)|}$$

Given a data set that allows for precise disjunctive guards, the place precision gets lower when guards on the output transitions are more overlapping.

3 Discovery of DPN with Overlapping Decision Rules

Given an event log containing information about process executions and a process model, the problem of decision mining can be regarded as that of discovering a DPN that characterizes the process: given a Petri net, we aim to discover the variables, write operations and guards of a DPN. Without loss of generality, we assume that event log \mathcal{E} is defined over the same transitions, variables and universe of values as the DPN.

Furthermore, we make four assumptions on the input event log \mathcal{E} . First, we assume that ignoring the variables, all recorded process instances are *compliant* with regard to the process represented by the Petri net. Second, for each event $e \in \mathcal{E}$ the executed transition can be uniquely determined, i.e., there are no unobservable transitions and each event can be uniquely mapped onto a single transition of the model. Third, we assume that events write attributes consistently, i.e., if an event writes an attribute v then all other events corresponding to the same transition also write attribute v . Fourth, we assume an initial value for each attribute. We show later in Sect. 3.3 that any event log can be transformed to an event log fulfilling these assumptions.

We discover the write operations of the DPN as follows. For each transition t of the Petri net, variable v is written by transition t if there exists an event $e \in \mathcal{E} : \text{trans}(e) = t$ that assigns a value to variable v (i.e., $v \in \text{writes}(e)$). We discover the guards for each transition t as described the next Sect. 3.1.

3.1 Overall Discovery Procedure

For each decision point p (i.e., places with more than two output transitions) we construct a set of observation instances related to p to be used to discover the guards. Given a set X , we denote the set of all multi-sets over a set X with $\mathbb{B}(X)$. Moreover, we use notation $X = [a^2, b]$ as short-hand notation to denote the multi-set $X = [a, a, b]$. Finally, we use \uplus to denote the sum of two multi-sets, i.e., $X \uplus [b, c] = [a^2, b^2, c]$.

Definition 1 (Observation Instances). *Let V be a set of variables and let U be a universe of possible values. Let P be the set of places of a Petri net and let $p \in P$ be a decision point. We denote with $p^\bullet = \{t_1, \dots, t_m\}$ the output transitions of p . Function $I \in P \rightarrow \mathbb{B}((V \rightarrow U) \times T)$ returns the multi-set of observation instances for p :*

$$I(p) = \uplus_{e \in \mathcal{E}, \text{trans}(e) \in p^\bullet} [(values(e), \text{trans}(e))]$$

For each event $e \in \mathcal{E}$ that refers to an output transition of p , i.e., $\text{trans}(e) \in p^\bullet$, the set of observation instances of p contains an instance $(\mathbf{x}, t) \in I(p)$, with $x \in (V \rightarrow U)$ being the observed values of the attributes, and t the observed transition. The values of \mathbf{x} are obtained by taking the latest observed value for the attributes in preceding events.

Example 3. Given the process model introduced in Example 1 and the event log \mathcal{E} introduced in Example 2 the multiset of observation instances for place p_1 is $I(p_1) = [((\text{status} = \text{unpaid}, \text{amount} = 30), \text{Payment})^2, ((\text{status} = \text{unpaid}, \text{amount} = 30), \text{Send Fine})^2, ((\text{status} = \text{paid}, \text{amount} = 30), \text{Close Fine})]$.

Algorithm 1. discoverGuards

Input: Petri net (Places P , Transitions T), Compliant Event Log (\mathcal{E}), Minimum Instances (v), Merge Ratio (ϵ)
Result: Guards of the DPN G

```

1 foreach  $p \in P$  s.t.  $|p^*| > 1$  do  $\psi_p \leftarrow \text{buildEstimator}(p, I, v \cdot |I(p)|, \epsilon)$ 
2 foreach  $t \in T$  do
3    $G(t) \leftarrow \text{true}$ 
4   foreach  $p \in {}^*t$  do  $G(t) \leftarrow G(t) \wedge \psi_p(t)$ 
5 end
6 return  $G$ 

```

Algorithm 1 describes the overall discovery method for the entire Petri net. Using the observation instances I_p , we build the guard function ψ_p for each decision point $p \in P$ through function `buildEstimator`. Having obtained the guard function, we assign each transition the conjunction of all rules obtained from their input places [4]. Function `buildEstimator` is described in the next Sect. 3.2.

3.2 Discovering Overlapping Rules

Whereas the construction of the observation instances $I(p)$ and the overall discovery procedure share similarities with the previous work [4], our technique considerably differs in how the actual rules are obtained through function `buildEstimator`. Our contribution is a new algorithm that discovers guards that may be partially overlapping, i.e., two or more transitions may be enabled for some state reachable in the DPN. We now describe how to compute function `buildEstimator`, which discovers overlapping guards for place p given the observation instances $I(p)$ and two user-defined parameters, the minimum number of instances n and the merge ratio ϵ . As our approach makes use of decision trees, we introduce the concept of decision tree builder.

Definition 2 (C4.5 Decision Tree Builder). *Let T be a set of transitions. Let $O = \mathbb{B}((V \rightarrow U) \times T)$ be a multi-set of observation instances over a set V of variables with values U . Let $n \in \mathbb{N}$ be the minimum number of instances on a leaf for the splitting criterion in the decision tree induction. Let $\mathcal{F}(V)$ be the universe of formulas over variable V . Function `buildTreen(O)` $\in 2^{\mathcal{F}(V) \times T}$ returns the leaves of a C4.5 decision tree built using the supplied set of instances. A leaf $(\text{expr}, t) \in \text{buildTree}_n(O)$ predicts transition $t \in T$ under condition $\text{expr} \in \mathcal{F}(V)$.*

The rule for a leaf of the decision tree is obtained by taking the conjunction of all conditions represented by those nodes that are encountered on a path from the leaf up to the root node [4]. As in [4], a **base decision tree** $\text{baseTree} = \text{buildTree}_n(I(p))$ is first built. We obtain the initial guard of transition t by taking the disjunction of all rules predicting t :

$$\psi(t) \leftarrow \begin{cases} \bigvee_{(\text{expr}, t) \in \text{baseTree}} (\text{expr}) & \text{if } \exists \text{expr}' \text{ s.t. } (\text{expr}', t) \in \text{baseTree} \\ \text{false} & \text{otherwise} \end{cases}$$

For each leaf $l = (\text{expr}, t)$ of the base decision tree, we extract those instances \bar{I}_l that have been wrongly classified by the base classifier: \bar{I}_l contains all those instances

$(x, t') \in I(p)$ such that the predicted transition t' is different from the transition t in leaf l , i.e., $t' \neq t$. Next, we build a **new decision tree** $subTree$ based on \bar{I}_l for each leaf, i.e., $subTree = buildTree_{n'}(\bar{I}_l)$. Since the size of \bar{I}_l can be significantly smaller than that of I , we scale down the parameter n to $n' = n \cdot |\bar{I}_l| / |I|$. The idea is that the second decision tree $subTree_l$ can *further discriminate* between the observed transitions among the wrongly classified instances, thus, possibly introducing partial overlap with the existing rule. There are two possible cases:

1. a decision tree with more than one leaf is found: $|subTree_l| > 1$;
2. a single-node decision tree with a single leaf $(true, t')$: $subTree_l = (true, t')$

In the **first case**, we build rules for each leaf $l' = (subExpr, t') \in subTree_l$ by taking the conjunction of the rule $expr$ from leaf l of tree $baseTree$ and the newly discovered rule $subExpr$. We obtain the new guard of transitions t' by adding newly discovered rules in disjunction to the existing ones:

$$\psi(t') \leftarrow \psi(t') \vee \bigvee_{(subExpr, t') \in subTree_l} (expr \wedge subExpr).$$

In the **second case** the decision tree represents a majority vote, i.e., the transition that was most often observed within the wrong instances is predicted. In this case we add rule $expr$ to the existing guard $\psi(t')$, i.e., $\psi(t') \leftarrow \psi(t') \vee (expr)$, but only if two conditions are met that avoid overfitting the data. First, $|\bar{I}_l| > n$, and, second, the fraction of observation instance in \bar{I}_l referring to t' is larger than the user-specified merge ratio ϵ . Finally, we assume that all transitions for which no rule could be found are always enabled. For a formal discussion of our algorithm refer to the technical report [13].



Fig. 2. Two decision trees that are discovered on the example data set. The number of instances is written in the root node (rectangle). The number of wrongly and correctly predicted instances is written next to leaf nodes (circle)

Example 4. Given the multi-set of observations for place p_1 obtained from the events given in Table 1, we build the guard estimation functions for each output transition of p_1 . First an initial decision tree is built. Assume that this initial decision tree consists of two leafs $l_1 = (status = paid, Close Fine)$ and $l_2 = (status = unpaid, Send Fine)$ as shown in Fig. 2. Ten of the 30 instances classified as `Close Fine` are wrongly classified. In all those instances transition `Payment` was observed. Those wrongly classified instances give evidence that also transition `Payment` is performed when `status` is `unpaid`. As depicted on the right-hand side of Fig. 2 an additional decision tree is built for those 10 instances. As the set of wrong observation instances \bar{I}_l only contains instances for transition `Payment`, the additional decision tree only consists of one leaf

that always predicts transition `Payment` by majority vote. If the user-defined threshold n is below the number of wrongly classified instances $|\bar{I}_2| = 10$, then, the condition `status = unpaid` will be added to the guard function of `Payment`. The same rule as shown in Fig. 1 is then discovered. Please note that this example is deliberately simplified. In a real-life setting event logs contain more than two attribute and, therefore, the additional decision would consist of multiple leafs.

3.3 Dealing with Real-Life Event Logs

In Sect. 3 we have made several assumptions to explain the key idea: the event log fits the Petri net perfectly, the set of attributes written by an event is consistent throughout the log, and every attribute value is initialized in the event corresponding to the first transition. Generally, real-life event logs do not satisfy this requirement. However, we can deal with these issues as shown next.

Non-compliant Event Log and Duplicate Transitions. An event log might contain events that cannot be matched reliably to a single transition in the DPN (e.g., in presence of noise or duplicate transitions). Similarly, for some transitions that are required according to the model no event has been recorded (e.g., when the recording is incomplete or for invisible routing transitions). Therefore, it might not be possible to determine $trans(e)$ for every event $e \in \mathcal{E}$. Using *alignment-based* techniques such as [5, 11] we can determine a closest corresponding **process trace**, i.e., the sequence of transition executions leading to a final state of the DPN, for each *log trace*. Work [4] uses the same technique to deal with non-compliant event logs and shows that for reasonably compliant event logs the error introduced by such an alignment is negligible.

Inconsistent Attributes. In Sect. 3, we restricted the set of write operations for a transition to those variables that are consistently given a value by every event e in $trans(e)$. In real-life logs attribute values can be missing due to temporary recording errors leading to an inconsistent recording of attributes for some transitions. Moreover, the *alignment-based* techniques might need to introduce artificial events with missing attribute values. Therefore we introduce a user-defined threshold K like in [4], but add a way to deal with missing values to it. A variable v is added to the set of write operations of a transition t when the variable is observed to be given a value by $K\%$ of the events e of t (i.e., $trans(e) = t$). As a result, attributes might be missing from the set of attributes written for an event $e \in \mathcal{E}$. Every time an event e does not assign a value to variable v even though it should, we assume its value $values(e)(v)$ to be \diamond . Symbol \diamond indicates that the value is *missing*. The C4.5 algorithm can deal with such missing values.

Unassigned Attributes. Decision trees cannot deal with uninitialized attributes (similar to NULL values in databases). In real-life event logs, attributes might be uninitialized if some of the first events of the log's traces do not assign a value to all attributes. This issue can be mitigated by defining default values that are used when attribute have not taken on values yet. As an example, let us assume an attribute `APPROVAL` with two literal values assigned by events: `No` and `Yes`. As default value, we can introduce a third possible value `Unknown`, make it possible to discover rules `APPROVAL = Unknown`.

4 Evaluation

We evaluate our technique using two real-life data sets, and compare the obtained results to standard methods like decision tree induction algorithms. An implementation of our technique is available in the *MultiPerspectiveExplorer* [14] package of the open-source process mining framework ProM.

4.1 Evaluation Setup

Approaches. We compared the performance of our approach expressed in terms of place fitness and place precision with three other methods. We choose two methods at the extreme ends of the respective measure, and one method that naïvely introduces overlap. In total, we compared the following four approaches:

- WO.** The model without rules, i.e., the guard `true` is used for all transitions. This results in a perfect place fitness, no guard is violated.
- DTF.** The model with rules discovered by a decision tree as in work [4] using `false` as guard for transitions that are not observed in the tree. This method will always result in a perfect place precision as there is only one enabled transition.
- DTT.** The model with rules discovered by a decision tree as in work [4] using `true` as guard for transitions not observed in the tree. This method naïvely introduces overlap by enabling all those transitions.
- DTO.** The model with rules discovered by the approach reported on in this paper as described in Sect. 3.1.

The DTF and WO methods are at the extreme ends of the respective measures. Our approach aims at providing better place fitness (i.e., less violated guards) at the expense of some place precision (i.e., multiple enabled transitions). Therefore, our approach should provide better place fitness than the DTF method together with better place precision than the a model without rules (WO). Method DTT is included to investigate whether our approach improves over a naïve method to introduce overlap.

Event Logs and Process Models. We used two anonymized real-life data sets: *road fines* and *sepsis*. The road fines event log was taken from an information system handling road-traffic fines by an Italian local police force [15]. The road fines log contains more than 150,000 cases with approximately 500,000 events. There are 9 data attributes recorded including the fine amount and the payment amount. The sepsis event log contains events about the pathways of patients within a hospital that have a suspicion for sepsis, a life threatening condition typically caused by an infection. This event log contains data recorded during a full year resulting in 1056 cases with about 15,000 events. There are 39 data attributes recorded, e.g., the results of blood tests and information about the patient from checklists filled in by nurses. For both event logs, we obtained a normative process model of the control-flow, thus, without any guards. We used the control-flow of the process model presented in [5] for the road fines data set. For the sepsis data set, we created a model with help of domain experts from the hospital. Both models allow for most the behavior observed in the logs, but are lacking precision.

We checked this using the fitness measure defined for DPNs in work [5] (road fines: 99.7%, sepsis: 98.6%) and the precision measure for DPNs proposed by [12] (road fines: 63.9%, sepsis: 16.5%). Therefore, both models are good candidates for adding precision through discovered rules.

Experimental Design. We performed experiments for every decision point for the road-fines and sepsis models, with the exception of four decision points of the sepsis model for which no technique was able to discover rules. We use the C4.5 implementation of the WEKA toolkit with the pruning feature activated; the merge ratio ϵ was set to 0.5 in all experiments. For each technique, we used 10 different values of *minimum number of instances* (`minInstances`) parameter that were equally distributed within a certain interval, which is determined as follows. The smallest value of the interval was chosen such that the discovered guards were not composed by more than 7 atoms. This choice was based on the assumption that guards with more than 7 atoms are too complex and humanly unreadable and, hence, of no business value. The upper bound of the interval was the smallest value that could still return a rule, i.e. larger values would return no rules. In fact, a too large value of the `minInstances` parameter would constrain the decision tree to be representative of so many instances that no reliable rule can be returned. It is worth observing that the interval potentially changed with varying the decision point and the technique (DTO, DTF and DTF) being considered.

4.2 Results and Discussion

We conducted the experiments and recorded the obtained place fitness and place precision for 15 places and 10 parameter settings¹. The boxplot in Fig. 3 shows the results. In the following paragraphs, we compare our method to the three other approaches.

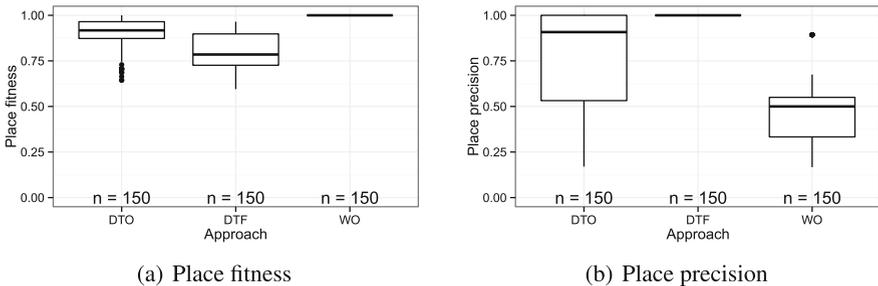


Fig. 3. Place fitness and local precision achieved by the proposed method (DTO) compared to the standard decision tree classifier (DTF), and the model without guards (WO)

¹ The data used for the evaluation is available under <http://purl.tue.nl/844997340832257>. For confidentiality reasons we cannot share the sepsis event log.

DTO vs. WO. Compared to WO, the results from our experiment (Fig. 3(b)) show clearly that DTO provides rules that increase the place precision against the process model without guards. The large spread of the obtained place precision indicates that, for some decision points and some parameter settings, our approach deliberately trades precision to obtain better fitting guards. This result is in line with the expectation that our approach returns overlapping rules that lose some precision for a better fitness.

DTO vs. DTF. The experimental results show that DTO discovers decision rules that lead to a better place fitness than the rules discovered by DTF (Fig. 3(a)) with, on average, a limited trade-off for lower precision. The outliers for DTO in Fig. 3(a) deserve some discussion. We inspected them and found that for some combinations of parameter settings and places, our approach failed to discover overlapping guards. It discovers the exact same rules as returned by DTT. Mostly, this happens for decision points with only two outgoing transitions and high settings of the `minInstances` parameter. This can be expected, as for decision trees with instances from two classes $\{A, B\}$ the wrong instances on a leaf $l = (expr, A)$ predicting transition A can only belong to the other transition B . Therefore, our approach will not discover a second decision tree with leaves predicting B , but rather use rule $expr$ from leaf l for the majority vote transition B . Our approach only allows this if the number of instances for B is above the setting of the `minInstances` parameter. Therefore, for high settings of `minInstances` and decision points with two outgoing transitions our approach is unlikely to improve over the normal decision tree classifier approach.

DTO vs. DTT. We also compare DTO against DTT, which naively assumes the guard true for transitions that are not predicted by the decision tree. For this comparison, we compare the results for decision points with more than two outgoing transitions, $|p^\bullet| > 2$, as the results obtained through DTT differ from the results of DTF only for those decision points. Figure 4(b) shows the fitness and precision for those places averaged over 10 parameter settings as discussed in Sect. 4.1. Each place is given a name for reference. The results in Fig. 4(b) show that DTO is able to discover overlapping guards that fit the observations better: for all of the considered decision points, the decision

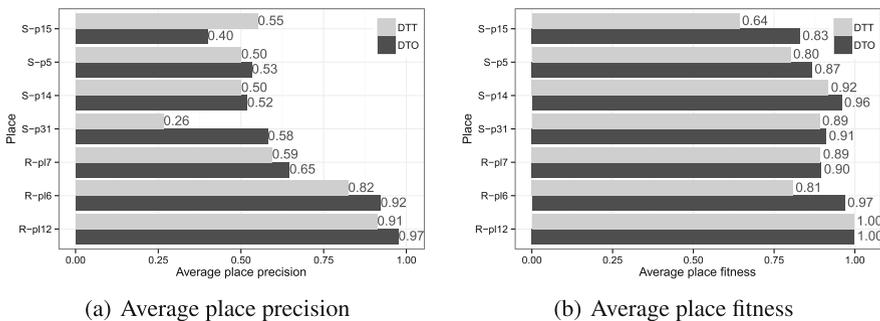


Fig. 4. Average place fitness and place precision achieved by the DTO method compared to the DTT method. Only decision points with more than two choices are shown.

rules returned by our approach increase the place fitness against those rules returned by DTT. Furthermore, the results show that for all except one decision point our approach discovers rules with higher place precision. In other words, it discovers more precise guards without losing fitness. In fact, for decision point S-p31 our approach obtained an average place precision of 0.58 whereas the rule returned by the DTT approach scores only 0.26. Our approach discovers guards for all six outgoing transitions whereas DTT only discovers guards for three transitions. On the remaining three DTT uses `true` as a guard, i.e., always enabled. The only decision point for which DTO obtains a worse precision score than DTT is S-p15. Our approach discovers guards that correspond to `true` for all three alternatives. However, this is not necessarily a bad representation of the observed data. In fact, the guards discovered by DTT cause the lowest place fitness in our experiment 0.65, i.e., the discovered guards are wrong in one third of the cases.

Example. Figure 5 shows a part of the DPN that we used for the sepsis data set. Table 2 shows the guards discovered by DTF, DTT and DTO, the approach presented in this paper, for the three alternative activities on decision point S-p5. All rules are based on two attributes: Lactate (L) and Hypotensie (H). DTF discovers the rule that patients with a lactate measurement (i.e., $L > 0$) are generally admitted to normal care and patients without lactate measurement ($L \leq 0$) leave the hospital. The guard for the admission to intensive care is returned as `false`. This leads to the situation where patients are never admitted to intensive care even if it is part of the model and observed. Obviously, this cannot be correct. DTF is unable to find a mutually-exclusive rule that includes this alternative activity given the information recorded in the event log. DTT discovers the same rules but naïvely assumes the guard `true` for admission to intensive care. Clearly, the DTT results are not satisfying as DTT would convey no rules about the admission of patients to intensive care. Our approach - DTO - discovers that patients with a lactate measurement ($L > 0$) can always be admitted to normal care. As an alternative to normal care, if attribute $H = \text{true}$ then patients can also be admitted to intensive care; otherwise, if $H = \text{false}$ patients leave the hospital. The guards for the activities overlap and the final decision is likely to be made on the basis of contextual factors, which are not encoded in the event log.

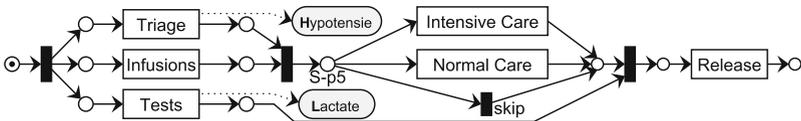


Fig. 5. Fragment of the process model used for the sepsis data set. After a triage form is filled, infusions are given and blood tests are taken patients are admitted to normal care, intensive care, or not admitted (skip). Two relevant attributes are recorded: Lactat (L) and Hypotensie (H).

Limitations and Threats to Validity. The results show that our technique is successful in uncovering overlapping rules in processes from event logs, and that these rules provide in some cases a much better characterization of the observed behavior. Still, the

Table 2. Guards discovered by the compared approaches at decision point S-p5

Approach	Normale care	Intensive care	Not admitted
DTF	$L > 0$	<code>false</code>	$L \leq 0$
DTT	$L > 0$	<code>true</code>	$L \leq 0$
DTO	$L > 0$	$L > 0 \wedge H = \text{true}$	$(L > 0 \wedge H = \text{false}) \vee L \leq 0$

proposed technique has some limitations and we evaluated our technique using only two real-life event logs. More experimental validation using event logs from different settings is required. An inherent limitation of our approach is that it can only use the majority vote to introduce overlapping guards for a decision point with two output transitions. This might cause the guard of one transition to be turned into the rule `true`, e.g. when the initial guards were based on a single condition. Our approach tends to discover guards that are more complex: Guard may become unreadable if algorithm's parameters are not carefully chosen.

5 Related Work

There are several approaches for decision mining given an event log with historical data about the process [1, 3, 4, 6, 16–21]. In all of these approaches, the decision-mining problem is translated into a classification problem, and solved using classification analysis techniques such as C4.5 [7]. However, every approach only discovers mutually-exclusive rules. Most related from the traditional classification field to our approach is work about multi-label classification [22, 23]. In a multi-label classification problem classes are not mutually exclusive, instances can be labeled with multiple classes, and the goal is to find the correct set of classes for unseen instances. Our setting is still different as we deal with instances that are only associated with one class, i.e., the executed transition. So, there is no work about discovering overlapping rules in the context of process models. Classifier chains methods are the closest to our work. They decompose the problem into multiple binary classification problems, one for each label [24]. This method assumes that instances are labeled with multiple classes. Also related to our work are methods for association rule mining [25]. The main problem of association-rule mining is that a potentially large set of rules is usually returned, failing to provide insights that are easy to interpret.

6 Conclusion

We propose a new technique for the discovery of overlapping rules in process models using event data. Existing techniques only return rules that assume completely deterministic decisions. This assumption rarely holds in reality. Our technique is the *first proposal* of a discovery technique that introduces overlapping rules. The technique aims to create process models that trade the precision of mutually-exclusive rules for the fitness of overlapping rules when the observed behavior gives evidence to such rules.

The evaluation using several real data sets shows that our technique is able to produce models with overlapping rules that fit the observed behavior better without losing too much precision. For some decision points, with more than 2 alternative activities, our technique returns rules that are both more fitting and more precise than the existing method [4]. As future work, we aim to investigate the application of other machine-learning techniques to decision mining. Moreover, we want to address limitations of decision mining techniques for data sets with imbalanced distributions of classes. We found that our technique helps to reveal rules when one transition is only observed for a small fraction of the cases, but a more thorough investigation of this phenomenon is needed.

References

1. van der Aalst, W.M.P.: *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, Heidelberg (2011)
2. Object Management Group (OMG): *Decision Model And Notation (DMN) Version 1.0, formal/2015-09-01* (2015)
3. Rozinat, A., van der Aalst, W.M.P.: Decision mining in ProM. In: Dustdar, S., Fiadeiro, J.L., Sheth, A.P. (eds.) *BPM 2006*. LNCS, vol. 4102, pp. 420–425. Springer, Heidelberg (2006)
4. de Leoni, M., van der Aalst, W.M.P.: Data-aware process mining: discovering decisions in processes using alignments. In: *SAC 2013*, pp. 1454–1461. ACM (2013)
5. Mannhardt, F., de Leoni, M., Reijers, H.A., van der Aalst, W.M.P.: Balanced multi-perspective checking of process conformance. *Computing* **98**(4), 407–437 (2016)
6. Grigori, D., Casati, F., Castellanos, M., Dayal, U., Sayal, M., Shan, M.C.: Business process intelligence. *Comput. Ind.* **53**(3), 321–343 (2004)
7. Quinlan, J.R.: *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco (1993)
8. Rosca, D., Wild, C.: Towards a flexible deployment of business rules. *Expert Syst. Appl.* **23**(4), 385–394 (2002)
9. Bose, R.P.J.C., Mans, R.S., van der Aalst, W.M.P.: Wanna improve process mining results? In: *IEEE Symposium on Computational Intelligence and Data Mining*, pp. 127–134 (2013)
10. Desel, J., Esparza, J.: *Free Choice Petri Nets*. Cambridge University Press, New York (1995)
11. van der Aalst, W.M.P., Adriansyah, A., van Dongen, B.F.: Replaying history on process models for conformance checking and performance analysis. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **2**(2), 182–192 (2012)
12. Mannhardt, F., de Leoni, M., Reijers, H.A., van der Aalst, W.M.P.: Measuring the precision of multi-perspective process models. In: *Business Process Management Workshops - BPM 2015* (2015, to appear)
13. Mannhardt, F., de Leoni, M., Reijers, H.A., van der Aalst, W.M.P.: Decision mining revisited - discovering overlapping rules. Technical report, BPMcenter.org, BPM Center Report BPM-01-06 (2016)
14. Mannhardt, F., de Leoni, M., Reijers, H.A.: The multi-perspective process explorer. In: *BPM Demo Session 2015*, vol. 1418, pp. 130–134. CEUR-WS.org (2015)
15. de Leoni, M., Mannhardt, F.: Road traffic fine management process. Eindhoven University of Technology. Dataset (2015). doi:[10.4121/uuid:270fd440-1057-4fb9-89a9-b699b47990f5](https://doi.org/10.4121/uuid:270fd440-1057-4fb9-89a9-b699b47990f5)
16. Jareevongpiboon, W., Janecek, P.: Ontological approach to enhance results of business process mining and analysis. *Bus. Process. Manag. J.* **19**(3), 459–476 (2013)

17. Catalkaya, S., Knuplesch, D., Chiao, C., Reichert, M.: Enriching business process models with decision rules. In: Lohmann, N., Song, M., Wohed, P. (eds.) *BPM 2013 Workshops*. LNBIP, vol. 171, pp. 198–211. Springer, Heidelberg (2014)
18. Ghattas, J., Soffer, P., Peleg, M.: Improving business process decision making based on past experience. *Decis. Support Syst.* **59**, 93–107 (2014)
19. Bazhenova, E., Weske, M.: Deriving decision models from process models by enhanced decision mining. In: *Business Process Management Workshops - BPM 2015* (2015, to appear)
20. Dunkl, R., Rinderle-Ma, S., Grossmann, W., Anton Fröschl, K.: A method for analyzing time series data in process mining: application and extension of decision point analysis. In: Nurcan, S., Pimenidis, E. (eds.) *CAiSE Forum 2014*. LNBIP, vol. 204, pp. 68–84. Springer, Heidelberg (2015)
21. de Leoni, M., Dumas, M., García-Bañuelos, L.: Discovering branching conditions from business process execution logs. In: Cortellessa, V., Varró, D. (eds.) *FASE 2013 (ETAPS 2013)*. LNCS, vol. 7793, pp. 114–129. Springer, Heidelberg (2013)
22. Boutell, M.R., Luo, J., Shen, X., Brown, C.M.: Learning multi-label scene classification. *Pattern Recogn.* **37**(9), 1757–1771 (2004)
23. Tsoumakas, G., Katakis, I.: Multi-label classification: an overview. *Int. J. Data Warehouse. Min.* **2007**, 1–13 (2007)
24. Read, J., Pfahringer, B., Holmes, G., Frank, E.: Classifier chains for multi-label classification. *Mach. Learn.* **85**(3), 333–359 (2011)
25. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. *SIGMOD Rec.* **22**(2), 207–216 (1993)