

Revisiting the Alpha Algorithm To Enable Real-Life Process Discovery Applications – Extended Report

Aaron Küsters^[0009–0006–9195–5380], Wil M.P. van der Aalst^[0000–0002–0955–6940]

Process and Data Science (PADS), RWTH Aachen University, Germany
kuesters@pads.rwth-aachen.de
wvdaalst@pads.rwth-aachen.de

Abstract. The Alpha algorithm was the first process discovery algorithm that was able to discover process models with concurrency based on incomplete event data while still providing formal guarantees. However, as was stated in the original paper, practical applicability is limited when dealing with exceptional behavior and processes that cannot be described as a structured workflow net without short loops. This paper presents the Alpha+++ algorithm that overcomes many of these limitations, making the algorithm competitive with more recent process mining approaches. The different steps provide insights into the practical challenges of learning process models with concurrency, choices, sequences, loops, and skipping from event data. The approach was implemented in ProM and tested on various publicly available, real-life event logs.

Keywords: Process Discovery · Process Mining · Process Models · Petri Nets

1 Introduction

The original *Alpha algorithm* was developed over twenty years ago [1,3]. The goal of the algorithm was to show the challenges related to discovering process models with concurrency from example traces. It was formally proven that, a process modeled as a structured workflow net without short loops, can be rediscovered from an event log that is directly-follows complete [3]. Despite this remarkable theoretical result, the Alpha algorithm has limited practical relevance for two main reasons:

- The original algorithm did not attempt to filter out infrequent behavior. Since exceptional behavior is not separated from frequent behavior, it is generally impossible to uncover structure from real-life event logs.
- The original algorithm assumed that the process can be modeled as a free-choice Petri net with unique visible activity labels. Most real-life processes can *not* be modeled as a structured workflow net without short loops and unique visible labels.

These limitations were already acknowledged in the papers proposing the algorithm, e.g., the focus of [3] was on showing the theoretical limits of process discovery based on directly-follows complete event logs. Many of the later process discovery approaches use these insights. Various extensions of the Alpha algorithm have been proposed, e.g., [18] extends the core algorithm to deal with long-term dependencies, and [19] extends the core algorithm to deal with invisible activities (e.g., skipping). Region-based process-discovery approaches provide formal guarantees. State-based regions were introduced by Ehrenfeucht and Rozenberg [9] in 1989 and generalized by Cortadella et al. [7]. In [2], it is shown how these state-based regions can be applied to process mining by first creating a log-based automaton using different abstractions. In [6,14], refinements are proposed to tailor state-based regions toward process discovery. Language-based regions work directly on traces without creating an automaton first; see, for example, the approaches presented in [5,20,21].

Variants of the Alpha algorithm and the region-based approaches have problems dealing with infrequent behavior and are rarely used in practice. The region-based approaches are also infeasible for larger models and logs. Approaches such as the eST-Miner [12] and the different variants of the inductive miner [10,11] aim to provide formal guarantees but can also handle infrequent behavior. Variants of the inductive miner have also been implemented in various commercial systems (e.g., Celonis). The so-called split-miner uses a combination of approaches to balance recall and precision [4].

The goal of this paper is to go back to the original ideas used by the Alpha algorithm and make the algorithm work in practical settings. The result is the *Alpha+++ algorithm*, which, not only extends the core algorithm, but also removes problematic noisy activities, adds invisible activities, repairs loops, and post-processes the resulting Petri net. The approach uses a broad combination of novel ideas, making the Alpha algorithm competitive when compared with the state-of-the-art. The ideas incorporated in the Alpha+++ algorithm may also be used in combination with other approaches (e.g., identifying problematic activities and introducing artificially created invisible activities).

The remainder of this paper is organized as follows. Section 2 introduces event logs, directly-follows graphs, and the original Alpha algorithm. Section 3 describes the Alpha+++ algorithm. The algorithm has been implemented in ProM (cf. Section 4) and evaluated using various event logs (cf. Section 5). Section 6 concludes the paper.

2 Preliminaries

2.1 Event Logs

Process mining starts from event data. An event may have many different attributes. However, here we focus on discovering the control flow and assume that each *event* has a *case* attribute, an *activity* attribute, and a *timestamp* attribute. We only use the timestamps to order events related to the same case. Therefore,

each case can be described as a sequence of activities, also called *trace*. An *event log* is a multiset of traces, as different cases can exhibit the same trace.

Definition 1 (Event Log). \mathcal{U}_{act} is the universe of activity names. A trace $\sigma = \langle a_1, a_2, \dots, a_n \rangle \in \mathcal{U}_{act}^*$ is a sequence of activities. An event log $L \in \mathcal{B}(\mathcal{U}_{act}^*)$ is a multiset of traces.

For example, $L_1 = [\langle a, b, c, d \rangle^{400}, \langle a, b, d \rangle^{250}, \langle d, a, b, c \rangle^4, \langle d, a, b \rangle^2]$ is an event log containing 656 cases with 4 different variants. Variant $\langle a, b, c, d \rangle$ is the most frequent one, i.e., $L_1(\langle a, b, c, d \rangle) = 400$.

We write $actMult(L) = [\sigma(i) \mid \sigma \in L \wedge 1 \leq i \leq |\sigma|]$ for the multiset of activities in an event log L and $act(L) = \{a \mid a \in actMult(L)\}$ for the set of activities.

2.2 Directly-Follows Graphs

A *Directly-Follows Graph* (DFG) is a graph showing how often one activity is followed by another. A DFG consists of the activities as nodes and has an arc from an activity $a \in \mathcal{U}_{act}$ to an activity $b \in \mathcal{U}_{act}$ if a is directly followed by b . Two special nodes, corresponding to a start and an end node, are added additionally.

Definition 2 (Directly-Follows Graph). A *Directly-Follows Graph* (DFG) is a pair $G = (A, \xrightarrow{c})$, where $A \subseteq \mathcal{U}_{act}$ is a set of activities and $\xrightarrow{c} \in \mathcal{B}((A \times A) \cup (\{\blacktriangleright\} \times A) \cup (A \times \{\blacksquare\}) \cup (\{\blacktriangleright\} \times \{\blacksquare\}))$ is a multiset of arcs. \blacktriangleright is the start node and \blacksquare is the end node.

Note that a DFG has arc weights. Hence, \xrightarrow{c} is a multiset, where $\xrightarrow{c}(a, b)$ denotes how often a is followed by b . We write $a \xrightarrow{c} b$ if and only if $\xrightarrow{c}(a, b) > 0$ holds. Similarly, we say that $a \xrightarrow{c}_{\geq t} b$ holds if and only if $\xrightarrow{c}(a, b) \geq t$.

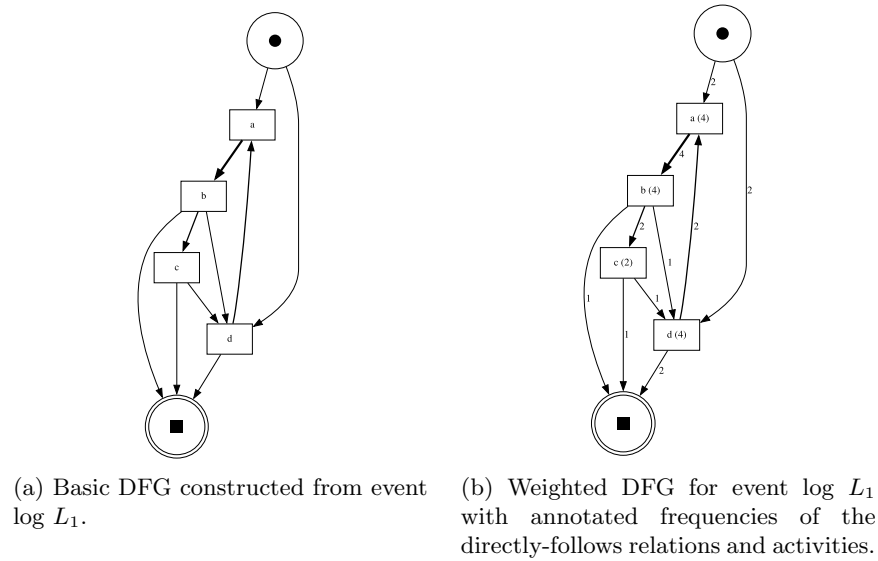
The construction of a DFG from an event log is straightforward.

Definition 3 (Constructing DFGs from Event Logs). Let $L \in \mathcal{B}(\mathcal{U}_{act}^*)$ be an event log. We can construct a DFG $disc_{dfg}(L) = (A, \xrightarrow{c})$ based on the directly-follows relations of event log L , with the set of activities $A = \{a \in \sigma \mid \sigma \in L\}$ and the multiset of arcs $\xrightarrow{c} = [(\sigma_i, \sigma_{i+1}) \mid \sigma \in L' \wedge 1 \leq i < |\sigma|]$, where $L' = [\langle \blacktriangleright \rangle \cdot \sigma \cdot \langle \blacksquare \rangle \mid \sigma \in L]$ denotes the event log where artificial start and end activities have been added.

Given an event log L , we can construct a DFG $disc_{dfg}(L) = (A, \xrightarrow{c})$, and in the context of L refer to the directly-follows relations in L represented by \xrightarrow{c} directly.

2.3 Petri Nets

We would like to discover process models which can represent more complex control-flow structures, like choices, loops, and concurrency. Therefore, we use *labeled* Petri nets as a target format for process discovery. The reader is assumed to be familiar with the Petri net basics.

Fig. 1: Example DFGs for event log L_1 .

Definition 4 (Labeled Petri Net). A labeled Petri net is a tuple $N = (P, T, F, l)$ with a set of places P , a set of transitions T (where $T \cap P = \emptyset$), a flow relation, $F \subseteq (P \times T) \cup (T \times P)$, and a labeling function $l \in T \rightarrow \mathcal{U}_{act}$. We write $l(t) = \tau$ if $t \in T \setminus \text{dom}(l)$ (i.e., t is a silent transition that cannot be observed).

A marking is represented by a multiset of places $M \in \mathcal{B}(P)$. For a node $x \in P \cup T$, we define the preset of x as $\bullet x = \{y \in P \cup T \mid (y, x) \in F\}$ and the postset of x as $x \bullet = \{y \in P \cup T \mid (x, y) \in F\}$. We focus on so-called *accepting* Petri nets, i.e., Petri nets with a defined initial and final state.

Definition 5 (Accepting Petri Net). An accepting Petri net is a triplet $AN = (N, M_{init}, M_{final})$ where $N = (P, T, F, l)$ is a labeled Petri net, $M_{init} \in \mathcal{B}(P)$ is the initial marking, and $M_{final} \in \mathcal{B}(P)$ is the final marking. \mathcal{U}_{AN} is the set of accepting Petri nets.

The language defined by an accepting Petri net is then simply given by the set of traces corresponding to all firing sequences that start in the initial marking M_{init} and end in the final marking M_{final} . A firing sequence leading from M_{init} to M_{final} is converted into a trace, i.e., a sequence of activities. Note that transitions that fire are mapped onto the corresponding activities. If a transition t is silent (i.e., $l(t) = \tau$), no corresponding activity is created when firing t . Hence, the language of an accepting Petri net is a set of traces.

2.4 Alpha Algorithm

A *process discovery algorithm* aims to discover a model from event data such that the language of the model best characterizes the example behavior seen in the event log.

Definition 6 (Process Discovery Algorithm). *A process discovery algorithm is a function $disc \in \mathcal{B}(\mathcal{U}_{act}^*) \rightarrow \mathcal{U}_{AN}$, i.e., based on a multiset of traces, an accepting Petri net is discovered.*

The classical Alpha process discovery algorithm was introduced in [3]. To be able to better explain the extensions presented in this paper, we split the description into three main parts. From an input event log L , place candidates are constructed based on the directly-follows relations of the log. The resulting set of place candidates is pruned to remove dominated candidates. Finally, the discovered Petri net is constructed.

Candidate Building

$$Cnd = \{(A, B) \mid \emptyset \subsetneq A, B \subseteq act(L) \wedge \forall_{a \in A} \forall_{b \in B} (a \xrightarrow{L} b) \\ \wedge \forall_{a, a' \in A} (a \not\xrightarrow{L} a') \wedge \forall_{b, b' \in B} (b \not\xrightarrow{L} b')\}$$

Candidate Pruning

$$Sel = \{(A_1, A_2) \in Cnd \mid \forall_{(A'_1, A'_2) \in Cnd} ((A_1 \subseteq A'_1 \wedge A_2 \subseteq A'_2) \\ \Rightarrow (A_1, A_2) = (A'_1, A'_2))\}$$

Petri Net Construction Let $PN = ((P, T, F, l), M_{init}, M_{final})$, where:

- $P = \{p_{(A,B)} \mid (A, B) \in Sel\} \cup \{i_W, o_W\}$
- $T = \{t_a \mid a \in act(L)\}$
- $F = \{(t_a, p_{(A,B)}) \mid (A, B) \in Sel \wedge a \in A\} \cup \{(p_{(A,B)}, t_b) \mid (A, B) \in Sel \\ \wedge b \in B\} \cup \{(i_W, t_s) \mid \exists_\sigma \langle s \rangle \cdot \sigma \in L\} \cup \{(t_e, o_W) \mid \exists_\sigma \sigma \cdot \langle e \rangle \in L\}$
- $l = \{(t_a, a) \mid a \in act(L)\}$
- $M_{init} = [i_W]$
- $M_{final} = [o_W]$

3 Alpha+++

In this section, we introduce the Alpha+++ process discovery algorithm based on the classical Alpha algorithm. Through certain pre-processing steps on the event log and a corresponding DFG, as well as fitness-based place filtering, this algorithm is especially well suited for real-life event logs.

The input for this process discovery algorithm is an event log L . In particular, only ordered traces of activities with corresponding frequencies are required. For the main steps of the algorithm, a DFG based on the event log L is used exclusively. Traces of the event log are only used for replay to remove unfitting place candidates. For simplicity, we assume that the traces of L already include artificial start and end activities, in particular, we assume $\{\blacktriangleright, \blacksquare\} \subseteq act(L)$.

We introduce the steps of the algorithm in the following order:

1. *Determine Activities*, where the set of activities used throughout the algorithm is determined. Problematic activities are removed from the event log and artificial activities are added, resulting in a repaired event log \hat{L} .
2. *Create an Advising DFG*, where an *advising DFG* is constructed based on the DFG corresponding to the repaired log \hat{L} , retaining only some of the original DFG edges.
3. *Candidate Building*, where a set of place candidates is built based on the directly-follows relation of the activities.
4. *Candidate Pruning*, where through efficient multistep filtering unfit or undesirable place candidates are discarded.
5. *Petri Net Construction*, where a Petri net is constructed based on the activities of the event log, the added artificial activities and the remaining place candidates.
6. *Post-Processing Petri Net*, where the repaired event log is replayed on the Petri net to remove problematic places.

3.1 Determine Activities

First, we determine the set of activities used in the later steps. Initially, starting with the set of activities occurring in the log, we first remove problematic activities that can cause issues with discovering place candidates later on. Next, we also add artificial activities to allow discovery of place candidates for certain loop and skip constructs.

Removing Problematic Activities Problematic activities can significantly alter the directly-follows relations of an event log, which are used in the later steps to identify place candidates. In the most extreme case, if a problematic activity randomly occurs between any other two activities in all traces, all the directly-follows information between two other activities would be lost.

We select a subset $\mathcal{A}_L \subseteq act(L)$ of activities to keep and remove the other problematic activities $act(L) \setminus \mathcal{A}_L$. There are many possible approaches to identifying problematic activities, such as calculating a problem-score per activity and considering all values above a certain threshold as problematic. For instance, for a simple problem-score, the fraction of directly-follows relation involving an activity which are parallel, i.e., also occur in the opposite direction, could be considered. This would, for example, allow to correctly identify the problem in the aforementioned extreme case.

Adding Artificial Activities Discovering Petri net constructs involving silent transitions is a non-trivial task for a DFG-based algorithm. Additionally, in later steps, we want to use traces from the log to assess the fitness of place candidates. Silent activities make calculating fitness scores significantly harder, as then token-based replay is no longer sufficient and computationally expensive alignments have to be computed. As a solution, we propose adding *artificial activities* to traces. They are not part of the activity set of the event log and

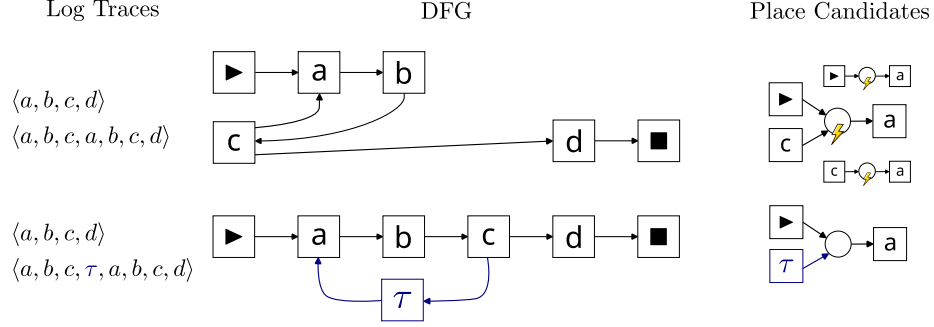


Fig. 2: Two event logs (traces shown on the left) and their DFGs. In the first event log (L_{\cup}) the directly-follows relation between \blacktriangleright and a is the same as between c and a . This causes issues, as the corresponding place candidates all have very low fitness. The added artificial activity τ inserted between the looped sequence $\langle a, b, c \rangle$ solves this problem, as the problematic directly-follows relation between c and a is replaced.

are only used to find and evaluate place candidates. In the final discovered Petri net, these artificial activities are then translated as silent transitions. This allows discovering Petri nets with silent transitions, while still retaining the advantages of token-based replay fitness evaluation during the algorithm steps. We add artificial activities for two types of constructs: *Loops* and *Skips*.

Adding artificial activities for loops is necessary, as the directly-follows relation between an end activity and a start activity of a loop can cause the discovery of problematic places. For example, consider the event log $L_{\cup} = [\langle a, b, c, d \rangle, \langle a, b, c, a, b, c, d \rangle]$. Clearly, this event log can be nicely expressed by a Petri net containing a loop construct, which allows repeating the activities a, b, c . However, the directly-follows relation $c \xrightarrow{L_{\cup}} a$ prevents discovering this loop accurately, as shown in Figure 2.

We detect loop constructs based on the directly-follows relations of the input event log L . For a given threshold $d \in \mathbb{R}^+$, we can define the set of detected loops:

Definition 7 (Detected Loops). Let $loops$ be the function that maps an event log to the set of detected loop start and end activities.

$$\begin{aligned}
 loops(L) = \{ (b, a) \in act(L) \times act(L) \mid & \exists_{(x_1, \dots, x_k) \in act(L)^*, i \in \{1, \dots, k\}} (x_i = a \\
 & \wedge \forall_{i \in \{1, \dots, k-1\}} (x_i \xrightarrow{L}_{\geq d} x_{i+1}) \\
 & \wedge x_k \xrightarrow{L}_{\geq d} b \wedge b \xrightarrow{L}_{\geq d} a) \}
 \end{aligned}$$

The parameter d determines the minimal DFG edge weight to consider when looking for loops. For example, with a threshold $d=1$ and the event log L_{\cup} , we can calculate $loops(L_{\cup}) = \{(c, a)\}$. As loop constructs can make a process model

very imprecise, we do not want to falsely detect loop behavior from rather infrequent directly-follows relations. For convenience, we can also consider threshold values d relative to the mean directly-follows weight.

For each detected loop endpoint pair $(b, a) \in \text{loops}(L)$, we want to add an artificial activity loop $\text{loop}_{b,a} \notin \text{act}(L)$. We write $\mathcal{A}_{loop} = \{\text{loop}_{b,a} \mid (b, a) \in \text{loops}(L)\}$ to denote the set of added artificial loop activities. Additionally, we define a transformation function which transforms a trace $\sigma \in L$ to a trace $\sigma' \in (\mathcal{A}_L \cup \mathcal{A}_{loop})^*$.

Definition 8 (Loop Repair Function). Let $\text{repair}_\circlearrowleft$ be the function that transforms a trace σ into a repaired trace with added artificial loop activities.

$$\text{repair}_\circlearrowleft(\sigma, L) = \begin{cases} \langle b, \text{loop}_{b,a}, a \rangle \cdot \text{repair}_\circlearrowleft(\sigma', L) & \text{if } \exists (b,a) \in \text{loops}(L) \sigma = \langle b, a \rangle \cdot \sigma' \\ \langle \rangle & \text{if } \sigma = \langle \rangle \\ \langle x \rangle \cdot \text{repair}_\circlearrowleft(\sigma', L) & \text{otherwise, with } \sigma = \langle x \rangle \cdot \sigma' \end{cases}$$

This function will be later used to transform the input event log L into a repaired event log, in which artificial activities have been added to relevant traces.

Next, we describe how artificial activities can assist in correctly discovering activity *Skips*, as shown in Figure 3.

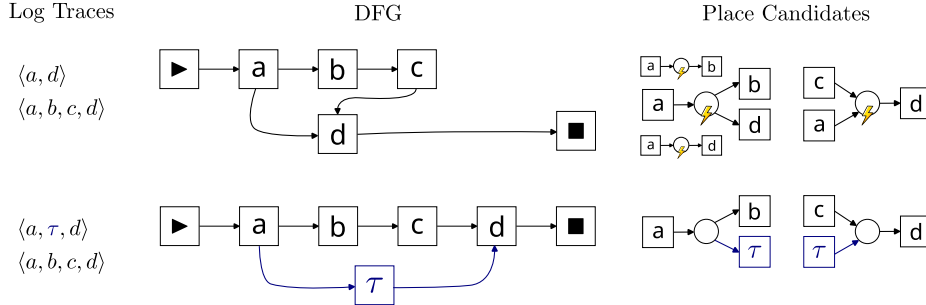


Fig. 3: Two event logs and their DFGs showcasing the motivation for repairing implicit skips. The directly-follows relation between a and d would suggest considering place candidates with poor fitness. The second log, where an artificial activity τ is inserted where b and c are skipped, mitigates this problem by replacing the directly-follows relation between a and d .

For a directly-follows-weight threshold $d \in \mathbb{R}^+$, the detected skips for event log L are defined by the following function, which provides the set of activities that have been detected as being “skippable” after an activity $a \in (\mathcal{A}_L \cup \mathcal{A}_{loop})$.

$$\begin{aligned} \text{skips}(a, L) = & \{b \in \text{act}(L) \mid a \xrightarrow{L} b \wedge a \xrightarrow{L} a \wedge b \xrightarrow{L} a \wedge b \xrightarrow{L} b \wedge a, b \notin \{\blacktriangleright, \blacksquare\} \\ & \wedge \emptyset \subsetneq \{x \in \text{act}(L) \mid b \xrightarrow{L} x\} \subseteq \{x \in \text{act}(L) \mid a \xrightarrow{L} x\}\} \end{aligned}$$

If $B \in \text{skips}(a, L)$ we assume that all activities $b \in B$ are optional steps after a . To allow appropriate model discovery in the rest of the algorithm, the log is repaired using a new artificial activity $\text{skip}_{a,B} \notin \text{act}(L)$. The set of all artificial skip activities is denoted by $\mathcal{A}_{\text{skip}}$. This artificial skip activity is inserted everywhere in a trace σ of L , where activity a is not directly followed by an activity $b \in B$ in σ (i.e., b was skipped). For that, we define the following transformation function:

$$\text{repair}_\tau(\sigma, L) = \begin{cases} \langle \rangle & \text{if } \sigma = \langle \rangle \\ \langle x \rangle \cdot \text{repair}_\tau(\sigma', L) & \text{if } \sigma = \langle x \rangle \cdot \sigma' \wedge \text{skips}(x, L) = \emptyset \\ \langle a, \text{skip}_{a,B} \rangle \cdot \text{repair}_\tau(\langle x \rangle \cdot \sigma', L) & \text{if } \sigma = \langle a, x \rangle \cdot \sigma' \wedge x \notin \text{skips}(a, L) \\ \langle a, x \rangle \cdot \text{repair}_\tau(\sigma', L) & \text{if } \sigma = \langle a, x \rangle \cdot \sigma' \wedge x \in \text{skips}(a, L) \end{cases}$$

We can now construct a *repaired* event log \hat{L} from the input event log L based on the previously identified set of detected loops $\text{loops}(L)$ and skips $\text{skips}(L)$. For that, we use their corresponding artificial activity set $\mathcal{A}_{\text{loop}}$ and $\mathcal{A}_{\text{skip}}$ as well as their corresponding trace transformation functions repair_\circ and repair_τ to transform the input event log L into a repaired event log \hat{L} . Note that $\text{act}(\hat{L}) = (\mathcal{A}_L \dot{\cup} \mathcal{A}_{\text{loop}} \dot{\cup} \mathcal{A}_{\text{skip}})$.

$$\hat{L} = [\text{repair}_\tau(\text{repair}_\circ(\sigma, L \setminus \mathcal{A}_L), L \setminus \mathcal{A}_L) \mid \sigma \in L \setminus \mathcal{A}_L]$$

3.2 Create an Advising DFG

Next, we extract a pruned DFG from the repaired event log \hat{L} , which ignores infrequent directly-follows relations. This DFG is used as guidance using the following algorithm steps. Note that this step does not modify the repaired event log: The output of this step is a pruned DFG containing the activities $\text{act}(\hat{L})$ as nodes. Edges between activities a and b are retained if their weight corresponds to at least 1% of the sum of the weights of all incoming edges to b or 1% of the sum of all outgoing edges from a . The value of 1% was determined as a good cutoff through experimentation. In addition, edges with weights below an absolute threshold value $n \in \mathbb{N}_0$ are also removed.

For the repaired event log \hat{L} and a given DFG-weight threshold $n \in \mathbb{N}_0$, we define the *advising DFG* (abbreviated as aDFG) as follows:

$$\begin{aligned} \min W(a, b) &= 0.01 \cdot \min \left\{ \sum_{c \in \text{act}(\hat{L})} \overset{\hat{L}}{\Rightarrow}(c, b), \sum_{c \in \text{act}(\hat{L})} \overset{\hat{L}}{\Rightarrow}(a, c) \right\} \\ \text{aDFG} &= \left(\text{act}(\hat{L}), \left[(a, b) \in \text{act}(\hat{L})^2 \mid \overset{\hat{L}}{\Rightarrow}(a, b) \geq \max \{n, \min W(a, b)\} \right] \right) \end{aligned}$$

3.3 Candidate Building

With the repaired event log and the aDFG, we can continue with building place candidates. Place candidates are composed of two sets of activities: The first set

corresponds to the transitions that should add a token to this place in a Petri net. The second set corresponds to transitions that should remove a token from this place.

The set of all place candidates is given by:

$$\begin{aligned} Cnd_0 = \{ & (A_1, A_2) \mid A_1, A_2 \subseteq act(\hat{L}) \wedge \forall_{a_1 \in A_1} \forall_{a_2 \in A_2} (a_1 \xrightarrow{\text{aDFG}} a_2) \\ & \wedge \forall_{a_1 \in A_1} \forall_{a_2 \in A_1 \setminus A_2} (a_1 \not\xrightarrow{\text{aDFG}} a_2) \\ & \wedge \forall_{a_1 \in A_2 \setminus A_1} \forall_{a_2 \in A_2} (a_1 \not\xrightarrow{\text{aDFG}} a_2) \\ & \wedge \exists_{a_1 \in A_1 \setminus A_2} \exists_{a_2 \in A_2 \setminus A_1} (a_2 \xrightarrow{\text{aDFG}} a_1) \} \end{aligned}$$

3.4 Candidate Pruning

The set of place candidates Cnd_0 includes many unfit places, which would produce process models with very low fitness. Furthermore, some place candidates might be dominated by others (e.g., the place candidate $(\{a\}, \{f\})$ is dominated by the candidate $(\{a, b\}, \{e, f\})$). Pruning the set of place candidates requires an efficient approach, as the number of place candidates can easily grow huge. We propose a three-step pruning approach. First, place candidates are filtered purely based on activity counts. If the difference in frequency of the input and output activity set is relatively large, the place candidate is rather unfit. This condition can be checked very efficiently. Next, the local fitness of the place candidate is calculated based on local trace replay. Local trace replay takes the order of the activities in the traces into account, and thus can detect even more unfit place candidates. Finally, to remove dominated place candidates, we retain only maximal place candidates.

Balance-based Pruning: For the balance-based pruning, we consider the number of activity occurrences in the log \hat{L} using $actMult(\hat{L})$. For a set of activities, $A \subseteq act(\hat{L})$ we can then sum the frequencies together as $count(\hat{L}, A) = \sum_{a \in A} actMult(\hat{L})(a)$. Based on that, we define the *balance* of a candidate (A_1, A_2) :

$$balance(\hat{L}, A_1, A_2) = \frac{|count(\hat{L}, A_1) - count(\hat{L}, A_2)|}{\max\{count(\hat{L}, A_1), count(\hat{L}, A_2)\}}$$

The balance of a candidate is between 0 and 1. Higher values are an indication that the place candidate is unfit. Based on a balance threshold $b \in [0, 1]$, candidates with a higher balance value than b can be filtered out:

$$Cnd_1 = \{(A_1, A_2) \in Cnd_0 \mid balance(\hat{L}, A_1, A_2) \leq b\}$$

Fitness-based Pruning: Let $fit(\sigma, (A_1, A_2), k)$ be defined as follows:

$$fit(\sigma, (A_1, A_2), k) = \begin{cases} 1 & \text{if } \sigma = \langle \rangle, k = 0 \\ 0 & \text{if } \sigma = \langle \rangle, k \neq 0 \\ 0 & \text{if } \sigma = \langle a \rangle \cdot \sigma', k = 0, a \notin A_1, a \in A_2 \\ fit(\sigma', (A_1, A_2), k + 1) & \text{if } \sigma = \langle a \rangle \cdot \sigma', a \in A_1, a \notin A_2 \\ fit(\sigma', (A_1, A_2), k - 1) & \text{if } \sigma = \langle a \rangle \cdot \sigma', k \geq 1, a \notin A_1, a \in A_2 \\ fit(\sigma', (A_1, A_2), k) & \text{if } \sigma = \langle a \rangle \cdot \sigma', (a \in A_1 \cap A_2 \vee a \notin A_1 \cup A_2) \end{cases}$$

Note that $fit(\sigma, (A_1, A_2), 0) = 1$ if the place candidate (A_1, A_2) fits the trace; otherwise it takes the value 0.

The traces relevant for a place candidate (A_1, A_2) are defined by the following function:

$$rel(A_1, A_2) = \left[\sigma = \langle a_1, \dots, a_n \rangle \in \hat{L} \mid \exists_{i \in \{1, \dots, n\}} (a_i \in A_1 \vee a_i \in A_2) \right]$$

We consider traces relevant for a place candidate, if they contain at least one activity that is in the set of outgoing or ingoing activities of that place candidate. For a single activity, we use the notation $rel(a) := rel(\{a\}, \emptyset)$ to denote the traces containing that activity.

We write $fit(\sigma, (A_1, A_2)) := fit(\sigma, (A_1, A_2), 0)$ and

$$fit(\hat{L}, (A_1, A_2)) := \sum_{\sigma \in rel(A_1, A_2)} fit(\sigma, (A_1, A_2))$$

for ease of notation.

For a given local candidate fitness threshold $t \in [0, 1]$, the candidates remaining after the local fitness replay pruning are then given as:

$$mfit(A_1, A_2) = \min \left\{ \frac{\sum_{\sigma \in rel(a)} fit(\sigma, (A_1, A_2))}{|rel(a)|} \mid a \in A_1 \cup A_2 \right\}$$

$$Cnd_2 = \left\{ (A_1, A_2) \in Cnd_1 \mid \frac{fit(\hat{L}, (A_1, A_2))}{|rel(A_1, A_2)|} \geq t \wedge mfit(A_1, A_2) \geq t \right\}$$

Maximal Candidate Selection: Finally, as the last candidate pruning step, all dominated place candidates are removed, just like in the original Alpha algorithm.

$$Sel = \{(A_1, A_2) \in Cnd_2 \mid \forall_{(A'_1, A'_2) \in Cnd_2} ((A_1 \subseteq A'_1 \wedge A_2 \subseteq A'_2) \Rightarrow (A_1, A_2) = (A'_1, A'_2))\}$$

3.5 Petri Net Construction

Based on the remaining place candidates, an accepting Petri net is constructed as the tuple $((P, T, F, l), M_{init}, M_{final})$, where

- $P = \{p_{(A_1, A_2)} \mid (A_1, A_2) \in Sel\}$
- $T = \{t_a \mid a \in act(\hat{L}) \setminus \{\blacktriangleright, \blacksquare\}\}$
- $F = \{(t_a, p_{(A_1, A_2)}) \mid (A_1, A_2) \in Sel \wedge a \in A_1 \setminus \{\blacktriangleright, \blacksquare\}\} \cup \{(p_{(A_1, A_2)}, t_a) \mid (A_1, A_2) \in Sel \wedge a \in A_2 \setminus \{\blacktriangleright, \blacksquare\}\}$
- $l = \{(t_a, a) \mid a \in \mathcal{A}_L\} \cup \{(t_a, \tau) \mid a \in (\mathcal{A}_{loop} \cup \mathcal{A}_{skip})\}$
- $M_{init} = [p_{(A_1, A_2)} \in P \mid \blacktriangleright \in A_1]$
- $M_{final} = [p_{(A_1, A_2)} \in P \mid \blacksquare \in A_2]$

are the components defined using the results of the previous steps.

3.6 Post-Processing Petri Net

Let $replay(p, PN, \sigma)$ be the replay function, which takes the value 1 exactly when the place p of the Petri net PN can replay trace σ (i.e., there is no missing or remaining token in p at any time when replaying σ on PN).

For a given local place replay fitness threshold $r \in [0, 1]$, we can then define the result of the post-process replay as $((P', T, F', l), M'_{init}, M'_{final})$, where the set of updated places P' is given by:

$$P' = \left\{ p_{(A_1, A_2)} \mid (A_1, A_2) \in Sel \wedge \frac{\sum_{\sigma \in rel(A_1, A_2)} replay(p, PN, \sigma)}{|rel(A_1, A_2)|} \geq r \right\}$$

The flow relation and initial and final markings are also updated correspondingly:

- $F' = \{(i, o) \in F \mid i \in P' \wedge o \in P'\}$
- $M'_{init} = [p \in M_{init} \mid p \in P']$
- $M'_{final} = [p \in M_{final} \mid p \in P']$

Note, that this post-processing uses a similar local fitness measure as the fitness-based pruning presented in section 3.4. One key difference is the handling of self-loops (where the post-processing replay is more restrictive). Another aspect to consider is that the candidate fitness pruning occurs before deleting dominated place candidates, while the post-processing fitness replay occurs after this pruning. In particular, choosing more aggressive filtering in the post-processing can lead to a decrease in places and an increase in disconnected transitions, which is often undesirable.

The final accepting Petri net discovered is $((P', T, F', l), M'_{init}, M'_{final})$.

4 Implementation

We implemented the Alpha+++ algorithm as a ProM¹ plugin (Java) and also created a Python implementation² for large-scale evaluation on a variety of real-life event logs. The ProM plugin (AlphaRevisitExperiments³) can be installed

¹ <https://promtools.org/>

² <https://github.com/aarkue/alpha-revisit-python>

³ <https://svn.win.tue.nl/repos/prom/Packages/AlphaRevisitExperiments/>

in ProM Nightly versions and can be used in standard mode to simply discover a Petri net or in interactive mode to experiment with different algorithm step options and view additional information (e.g., how many place candidates were pruned in which step). In both versions, the Alpha+++ preset can be selected out of the preset list on the top. The parameters used throughout the algorithm steps can then be changed. Additionally, the different algorithm steps can be swapped with alternatives or skipped, allowing for further experimentation.

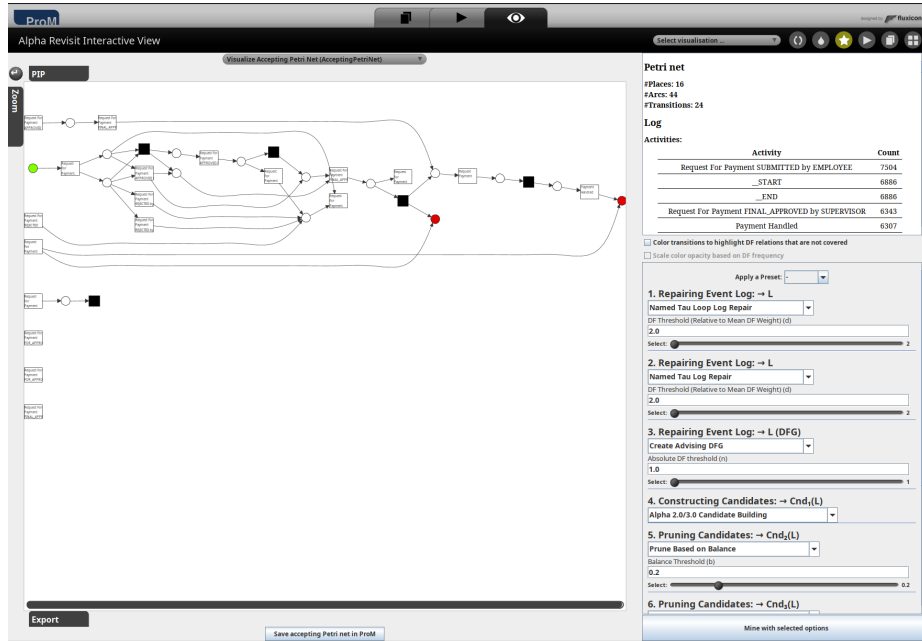


Fig. 4: A screenshot of the interactive mode of the developed ProM plugin *Alpha Revisit Experiments*. On the right, the steps and different step parameters can be configured. The main section on the left shows the discovered Petri net.

5 Evaluation

To evaluate the proposed Alpha+++ algorithm ($\alpha+++$), we discovered Petri nets for five real-life event logs, shown in Table 1. For comparison, we also discovered models using the Inductive Miner Infrequent (IMf) and the standard Alpha algorithm (α). We subsequently calculated alignment-based fitness, precision and F1-scores using PM4Py⁴.

⁴ <https://pm4py.fit.fraunhofer.de/> (Version 2.6.1)

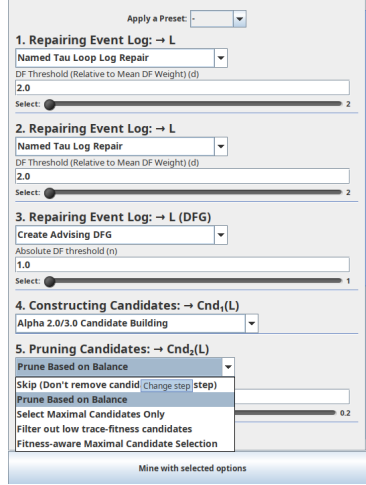


Fig. 5: A closeup of the configuration panel of the ProM plugin. For each step, there are multiple possible algorithm implementations available, which can be selected from a dropdown-menu. Different presets that form a complete process discovery algorithm, e.g., the Alpha+++ algorithm presented here, can be applied at the top.

Table 1: Overview of the event logs used for evaluation. We used a random sample of 3000 cases from the BPI Challenge 2019 log for computational reasons, as it allowed for alignment-based evaluation of the discovered models.

Event Log	#Events	#Activities	#Traces	#Variants	Reference
RTFM	561,470	11	150,370	231	[8]
Sepsis	15,214	16	1,050	846	[13]
BPI Challenge 2019 (Sample of 3000 Cases)	18,972	34	3,000	470	[15]
BPI Challenge 2020 (Request for Payment)	36,796	19	6,886	89	[17]
BPI Challenge 2020 (Domestic Declaration)	56,437	17	10,500	99	[16]

For IMf, we evaluated four models per event log using noise thresholds of 0.1, 0.2, 0.3 and 0.4. For α , we used four variant filtering approaches upfront: Either only selecting the 10 most common variants or the n most common variants to cover at least 10%, 50% or 80% of traces. For $\alpha+++$, we chose artificial activity thresholds of 2 and 4 (relative to the mean directly-follows weight) for the log repair steps. Here, a lower threshold value causes more artificial activities to be added. For each artificial activity threshold, we selected five combinations of the balance b , local candidate fitness t and local place replay fitness r thresholds. Note, that for t and r a value closer to 1 and for b a value closer to 0 is more restrictive. We did not apply problematic activities filtering.

The evaluation results are shown in Table 2. Overall, the fitness and F1-scores of $\alpha+++$ are competitive compared to the IMf. 8 of the 20 models discovered with α are not easy sound (i.e., no final marking is reachable), and thus no alignment scores could be computed. The remaining 12 models exhibit rather low fitness for some logs but very high precision across the board, significantly boosting the corresponding F1-values. Although our approach does not guarantee easy soundness, all 50 Petri nets discovered with $\alpha+++$ are easy sound and allow computation of alignments. There are notable differences across the different event logs: $\alpha+++$ performs significantly worse compared to the IMf on the Sepsis log in terms of F1-score, caused by lower precision scores, as the models discovered with $\alpha+++$ seem to be underfitting. On the two BPI Challenge 2020 logs, $\alpha+++$ outperforms the IMf in most configurations, often also exhibiting better fitness and precision scores simultaneously.

The influence of the parameters of $\alpha+++$ is mostly as expected: More restrictive b, t, r values improve the fitness of the models while decreasing the precision.

Manual inspection of the discovered models reveals that the models discovered with $\alpha+++$ are mostly rather simple and often consist of several disconnected model fragments. Furthermore, multiple models exhibit redundant structures involving silent transitions (e.g., a place with one labeled transition as preset and one silent transition as postset). Such constructs could be removed by further post-processing of the Petri net. For more details and a comprehensive list of the discovered models, see section B and section C of the appendix.

6 Conclusion

In this paper, we revisited the Alpha algorithm to overcome its limitations, focusing on real-life event logs. For that, we presented the Alpha+++ algorithm which, like the Alpha algorithm, primarily uses directly-follows relations to discover Petri nets. Alpha+++ pre-processes event logs by adding artificial activities for potential loop or skip constructs. This allows discovering silent transitions while still assessing the fitness of places by easily computable token-based replay instead of expensive alignment computations. Subsequently, place candidates are generated based on a pruned DFG. A multistep candidate filtering approach efficiently removes place candidates with low fitness, configurable

through parameters. We implemented the Alpha+++ algorithm both as a ProM plugin and in Python. The ProM plugin is available in ProM nightly builds and also features an interactive mode to allow experimenting with different algorithm steps and parameters. We evaluated the Alpha+++ on five real-life event logs and compared the results to the classical Alpha algorithm and the widely adapted Inductive Miner Infrequent. Overall, the results indicate that the Alpha+++ algorithm is competitive in terms of fitness and precision. In general, the different step parameter configurations tested reliably determine the trade-off between fitness and precision.

Further research should include further evaluation of the algorithm. For that, additional performance metrics like simplicity or generality could be included and also compared to other process discovery algorithms. It is particularly interesting to see if there are any patterns regarding algorithm parameters, event log properties, and model performance. Such observations could enable automatic parameter selection based on the log, and thus simplify Alpha+++ to a well-performing one-in-all algorithm. Additionally, a more comprehensive qualitative analysis of the discovered models is needed. Further research could also explore if any theoretical guarantees, such as easy-soundness, are attainable, e.g., using more sophisticated post-processing of the discovered Petri net.

References

1. W.M.P. van der Aalst and B.F. van Dongen. Discovering Workflow Performance Models from Timed Logs. In Y. Han, S. Tai, and D. Wikarski, editors, *International Conference on Engineering and Deployment of Cooperative Information Systems (EDCIS 2002)*, volume 2480 of *Lecture Notes in Computer Science*, pages 45–63. Springer-Verlag, Berlin, 2002.
2. W.M.P. van der Aalst, V. Rubin, H.M.W. Verbeek, B.F. van Dongen, E. Kindler, and C.W. Günther. Process Mining: A Two-Step Approach to Balance Between Underfitting and Overfitting. *Software and Systems Modeling*, 9(1):87–111, 2010.
3. W.M.P. van der Aalst, A.J.M.M. Weijters, and L. Maruster. Workflow Mining: Discovering Process Models from Event Logs. *IEEE Transactions on Knowledge and Data Engineering*, 16(9):1128–1142, 2004.
4. A. Augusto, R. Conforti, M. Marlon, M. La Rosa, and A. Polyvyanyy. Split Miner: Automated Discovery of Accurate and Simple Business Process Models from Event Logs. *Knowledge Information Systems*, 59(2):251–284, 2019.
5. R. Bergenthum, J. Desel, R. Lorenz, and S. Mauser. Process Mining Based on Regions of Languages. In G. Alonso, P. Dadam, and M. Rosemann, editors, *International Conference on Business Process Management (BPM 2007)*, volume 4714 of *Lecture Notes in Computer Science*, pages 375–383. Springer-Verlag, Berlin, 2007.
6. J. Carmona, J. Cortadella, and M. Kishinevsky. A Region-Based Algorithm for Discovering Petri Nets from Event Logs. In *Business Process Management (BPM 2008)*, pages 358–373, 2008.
7. J. Cortadella, M. Kishinevsky, L. Lavagno, and A. Yakovlev. Deriving Petri Nets from Finite Transition Systems. *IEEE Transactions on Computers*, 47(8):859–882, August 1998.
8. M. de Leoni and F. Mannhardt. Road Traffic Fine Management Process, 2015. https://data.4tu.nl/articles/dataset/Road_Traffic_Fine_Management_Process/12683249.
9. A. Ehrenfeucht and G. Rozenberg. Partial (Set) 2-Structures - Part 1 and Part 2. *Acta Informatica*, 27(4):315–368, 1989.
10. S.J.J. Leemans, D. Fahland, and W.M.P. van der Aalst. Discovering Block-Structured Process Models from Event Logs Containing Infrequent Behaviour. In N. Lohmann, M. Song, and P. Wohed, editors, *Business Process Management Workshops, International Workshop on Business Process Intelligence (BPI 2013)*, volume 171 of *Lecture Notes in Business Information Processing*, pages 66–78. Springer-Verlag, Berlin, 2014.
11. S.J.J. Leemans, D. Fahland, and W.M.P. van der Aalst. Scalable Process Discovery and Conformance Checking. *Software and Systems Modeling*, 17(2):599–631, 2018.
12. L.L. Mannel and W.M.P. van der Aalst. Discovering Process Models with Long-Term Dependencies While Providing Guarantees and Handling Infrequent Behavior. In L. Bernardinello and L. Petrucci, editors, *Application and Theory of Petri Nets and Concurrency (Petri Nets 2022)*, volume 13288 of *Lecture Notes in Computer Science*, pages 303–324, 2022.
13. F. Mannhardt. Sepsis Cases - Event Log, 2016. https://data.4tu.nl/articles/dataset/Sepsis_Cases_-_Event_Log/12707639.
14. M. Solé and J. Carmona. Process Mining from a Basis of State Regions. In J. Lilius and W. Penczek, editors, *Applications and Theory of Petri Nets 2010*, volume 6128 of *Lecture Notes in Computer Science*, pages 226–245. Springer-Verlag, Berlin, 2010.

15. B. van Dongen. BPI Challenge 2019, 2019. https://data.4tu.nl/articles/dataset/BPI_Challenge_2019/12715853.
16. B. van Dongen. BPI Challenge 2020: Domestic Declarations, 2020. https://data.4tu.nl/articles/dataset/BPI_Challenge_2020_Domestic_Declarations/12692543.
17. B. van Dongen. BPI Challenge 2020: Request For Payment, 2020. https://data.4tu.nl/articles/dataset/BPI_Challenge_2020_Request_For_Payment/12706886.
18. L. Wen, W.M.P. van der Aalst, J. Wang, and J. Sun. Mining Process Models with Non-Free-Choice Constructs. *Data Mining and Knowledge Discovery*, 15(2):145–180, 2007.
19. L. Wen, J. Wang, W.M.P. van der Aalst, B. Huang, and J. Sun. Mining Process Models with Prime Invisible Tasks. *Data and Knowledge Engineering*, 69(10):999–1021, 2010.
20. J.M.E.M. van der Werf, B.F. van Dongen, C.A.J. Hurkens, and A. Serebrenik. Process Discovery using Integer Linear Programming. *Fundamenta Informaticae*, 94:387–412, 2010.
21. S.J. van Zelst, B.F. van Dongen, W.M.P. van der Aalst, and H.M.W. Verbeek. Discovering Workflow Nets Using Integer Linear Programming. *Computing*, 100(5):529–556, 2018.

A Generality and Simplicity Evaluation

We additionally evaluated all the Petri nets discovered for the evaluation in section 5 in terms of simplicity and generalization. For that, we again utilized PM4Py using the corresponding evaluation functions. In those, simplicity is measured using the inverse arc degree of nodes in the Petri net, while generalization relies on how frequent model elements are revisited during replay⁵.

⁵ See also <https://pm4py.fit.fraunhofer.de/documentation#evaluation>.

B Disconnected Transitions

In many of the models discovered with Alpha+++ there are one or more transitions that are disconnected from the rest of the model. As expected, this phenomenon is most frequently observable for more restricting algorithm parameters, which aggressively filter out low-fitness place candidates and places. In Table 4, we counted the number of disconnected labeled transitions for each of the models and calculated what percentage of the activities of the log are represented by disconnected transitions on average.

Table 4: Disconnected labeled transitions in models discovered with Alpha+++. The average percentage refers to the average fraction of activities in the event log that are labels of disconnected transitions.

	Alpha+++ Algorithm										Average %
	Artificial Activity Threshold of 2.0					Artificial Activity Threshold of 4.0					
	<i>b=0.5</i>	<i>b=0.3</i>	<i>b=0.2</i>	<i>b=0.2</i>	<i>b=0.1</i>	<i>b=0.5</i>	<i>b=0.3</i>	<i>b=0.2</i>	<i>b=0.2</i>	<i>b=0.1</i>	
	<i>t=0.5</i>	<i>t=0.7</i>	<i>t=0.8</i>	<i>t=0.8</i>	<i>t=0.9</i>	<i>t=0.5</i>	<i>t=0.7</i>	<i>t=0.8</i>	<i>t=0.8</i>	<i>t=0.9</i>	
	<i>r=0.5</i>	<i>r=0.6</i>	<i>r=0.7</i>	<i>r=0.8</i>	<i>r=0.9</i>	<i>r=0.5</i>	<i>r=0.6</i>	<i>r=0.7</i>	<i>r=0.8</i>	<i>r=0.9</i>	
RTFM	1	1	2	2	2	1	1	2	2	2	14.55
Sepsis Cases	1	5	5	7	7	6	7	12	12	13	46.88
BPI Challenge 2019 (Sample of 3000 Cases)	8	10	14	14	22	7	10	14	14	24	40.29
BPI Challenge 2020 (Requests for Payment)	0	0	0	0	0	0	0	0	0	0	0
BPIC 2021	0	0	0	0	0	1	1	1	1	1	2.94

In Alpha+++ disconnected transitions are not handled separately and always included in the final Petri net. Note, that always including the transitions without restrictions is not the only option, and has non-negligible effects on the fitness and precision of the complete model. In particular, the decision whether an activity should be included as a disconnected transition, i.e., without restrictions, or not at all can be made separately for each transition based on frequency information of the corresponding activity.

In particular, certain disconnected transitions can be omitted with the goal to maximize certain performance metrics, like precision or F1 score. Note that the fitness of the model will consistently decrease whenever a disconnected (labeled) transition is removed, as the resulting Petri net is more restrictive, i.e., not allowing this activity to occur at all. Thus, it is in particular interesting to attempt to increase the precision of the end model while not decreasing the fitness too much. As a heuristic, we propose removing the labeled, disconnected transitions in the reverse order of their frequency in the log (i.e., transitions with activity labels that are infrequent in the log are removed first). We call this technique of determining the removal order *greedy*. Removing a transition with a rather infrequent activity label often leads to only a small decrease in fitness and a bigger gain in precision. In Figure 6 we applied this technique to a model discovered using the Sepsis event log. As such, we removed disconnected, labeled

transitions one by one in the given order, and measured the fitness, precision and F1 score in every step. The x-axis additionally shows the number of events that are no longer covered by a transition (i.e., the sum of events in the log with an activity label corresponding to a transition that was removed).

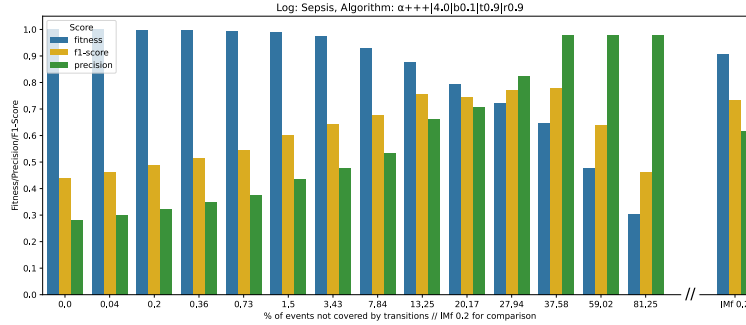


Fig. 6: Fitness, precision and F1 score values of the resulting models when removing transitions one by one greedily. Clearly, in most steps, the greedy ordering has the desired effect of only decreasing the fitness slightly while increasing the precision significantly. Performance values for the Inductive Miner Infrequent 0.2 are given as a reference point.

We ran experiments to compare how well the greedy activity removal performs compared to the brute-force approach of evaluating all possible removal orders. For that, we selected to maximize the F1 score and evaluated the Petri nets obtained by removing the transitions corresponding to the first k infrequent activities. In the brute force approach, we also selected a set of k disconnected labeled transitions to remove, but choose that set from all possibilities that maximizes the F1 score of the resulting net. This allows us to compare the F1 score for every number k of removed transitions between the removal set based on the greedy ordering or the best possibility for maximizing the F1 score. As such, the brute force approach provides an upper bound of achievable F1 score improvements and is always guaranteed to be at least as high as the F1 score achieved using greedy removal.

Figure 7 shows the observed gain in F1 score from the brute-force approach compared to the greedy ordering for one model discovered with the Sepsis event log. Our experiments indicate that for most use-cases, the greedy activity removal order should be preferred because of the significantly better running time. While the brute-force approach might find better removal sets in some configurations, it requires evaluating models for every possible removal set of a given size, which is computationally expensive. In particular, we observed that for many discovered models the greedy removal order was, in fact, already optimal regarding the F1 score.

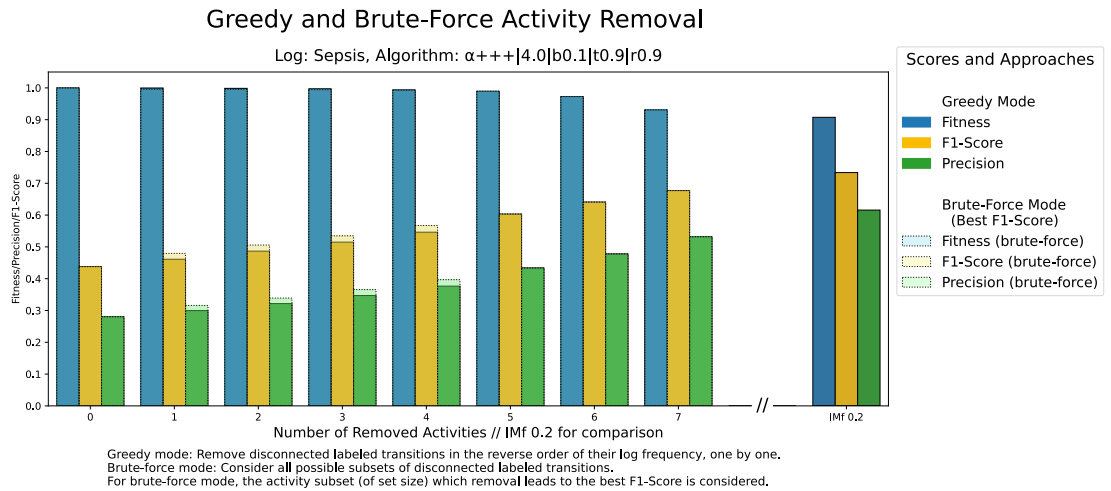


Fig. 7: Comparison of the greedy activity removal order versus the best removal set of that size (w.r.t. F1 score). The possible F1 score improvements of the brute-force approach (shown as light orange on top of bar) are noticeable for some removal counts.

The brute-force and greedy transition removal post-processing, as well as the corresponding evaluation and visualization functions, are also published open-source⁶.

⁶ <https://github.com/aarkue/alpha-revisit-python/blob/main/remove-disconnected-acts.ipynb>

C Discovered Process Models

In the next pages, we include all the process models used for evaluation (see section 5). The logs are enumerated in the following order:

1. RTFM
2. Sepsis Cases
3. BPI Challenge 2019 (Sample of 3000 Cases)
4. BPI Challenge 2020 (Request for Payment)
5. BPI Challenge 2020 (Domestic Declarations)

To stay concise, we abbreviate the algorithms as before and additionally introduce the following notation for the different configurations of the $\alpha+++$ algorithm:

$$\alpha+++; \underbrace{2.0}_{\text{DF Threshold}} ; b \underbrace{0.3}_{\text{Balance Thresh.}} ; t \underbrace{0.7}_{\text{Local Candidate Fitness Thresh.}} ; r \underbrace{0.6}_{\text{Replay Fitness Thresh.}}$$

For example, the configuration $\alpha+++;2.0;b0.3;t0.7;r0.6$ represents that the model was discovered with $\alpha+++$ using a relative DF threshold of 2.0 (used for adding artificial activities), a balance threshold of 0.3 (used to filter place candidates), a local candidate fitness threshold of 0.7 (used to prune the candidates further) and finally a replay place fitness threshold of 0.6 (used to post-process the Petri net by removing unfit places).

We (visually) post-processed the models discovered with $\alpha+++$, to allow for easier interpretation and comparison of the models, as the models often exhibit multiple disconnect parts. In particular, we added artificial start and end transitions, marked by \blacktriangleright or \blacksquare , and connected them to an added place, which also connects all disconnected process parts. Note, that this does not alter the underlying semantic of the Petri net, if interpreted correctly.

C.1 RTFM

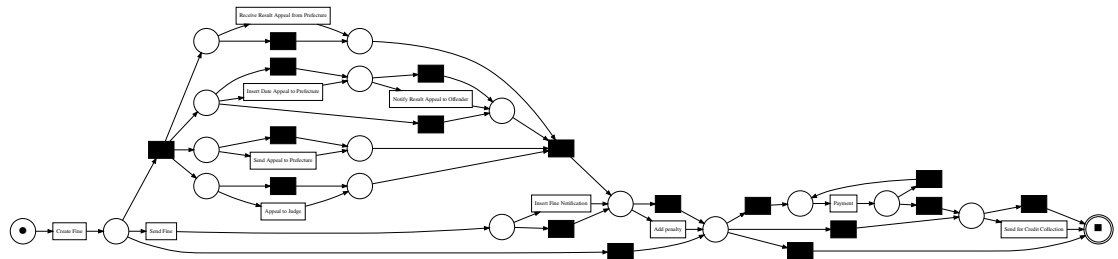


Fig. 8: Model discovered using IMf 0.1 on RTFM

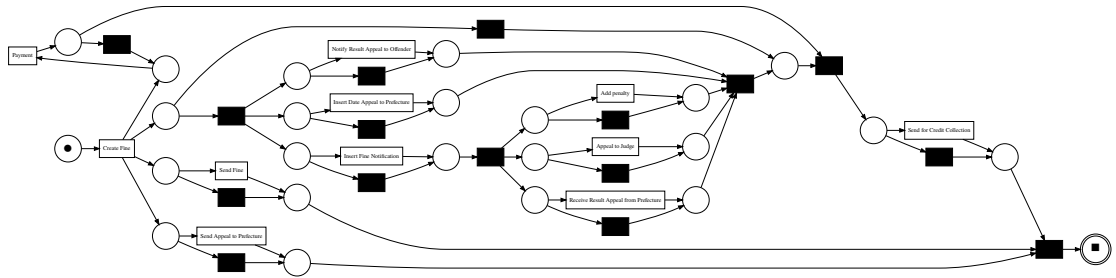


Fig. 9: Model discovered using IMf 0.2 on *RTFM*

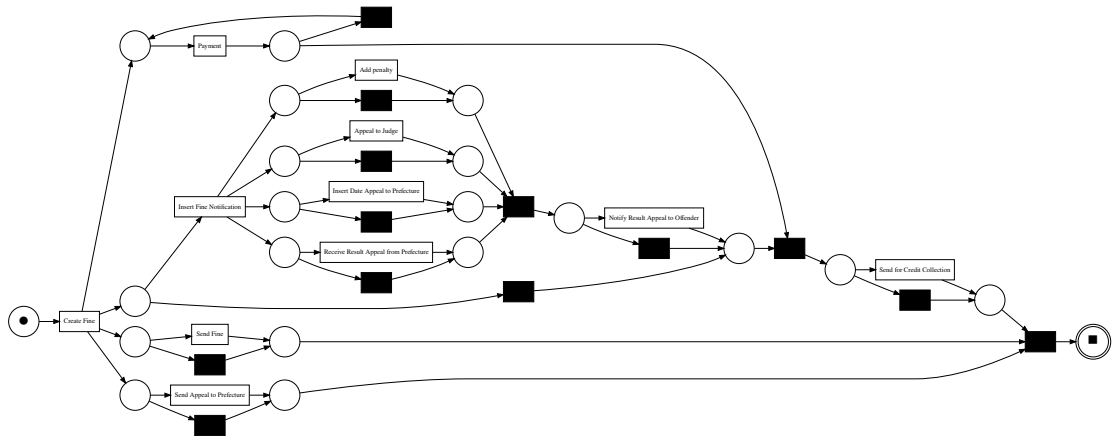


Fig. 10: Model discovered using IMf 0.3 on *RTFM*




Fig. 11: Model discovered using IMf 0.4 on *RTFM*




Fig. 12: Model discovered using α Top10 on *RTFM*




Fig. 13: Model discovered using α 10% Cov on *RTFM*




Fig. 14: Model discovered using α 50%Cov on *RTFM*




Fig. 15: Model discovered using α 80%Cov on *RTFM*

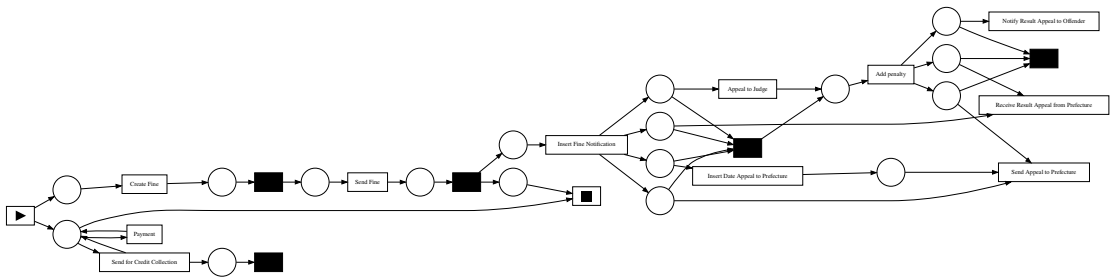


Fig. 16: Model discovered using $\alpha+++;2.0;b0.5;t0.5;r0.5$ on *RTFM*

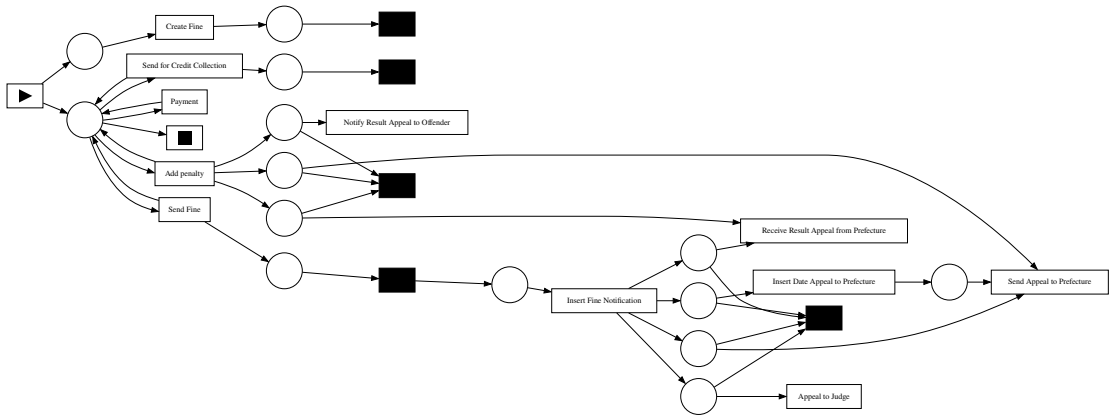


Fig. 17: Model discovered using $\alpha+++;2.0;b0.3;t0.7;r0.6$ on *RTFM*

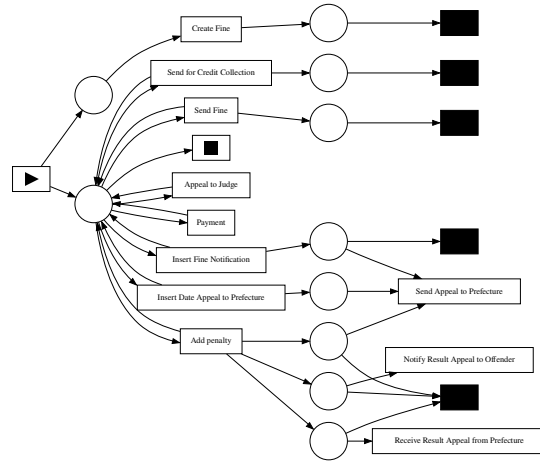


Fig. 18: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.7$ on *RTFM*

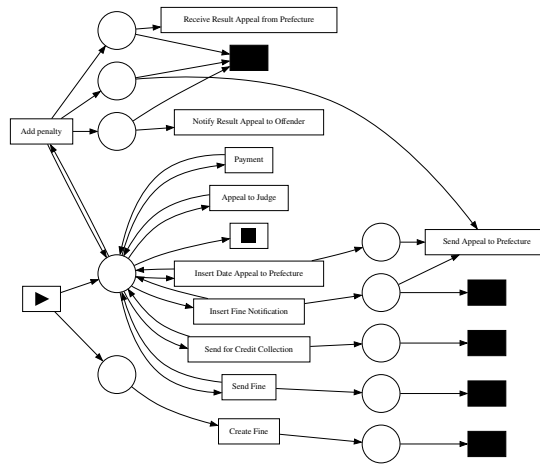


Fig. 19: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.8$ on *RTFM*

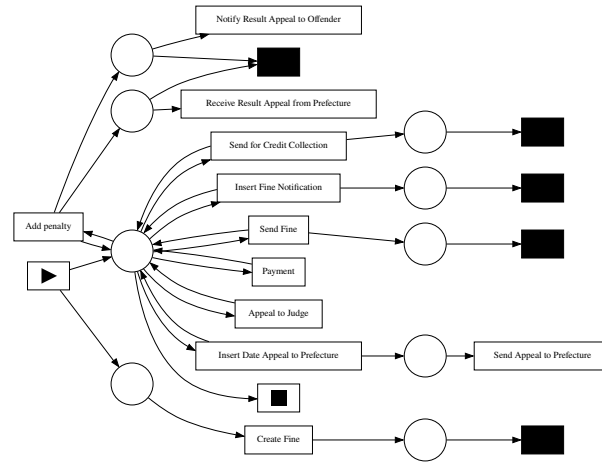


Fig. 20: Model discovered using $\alpha^{++}; 2.0; b0.1; t0.9; r0.9$ on *RTFM*

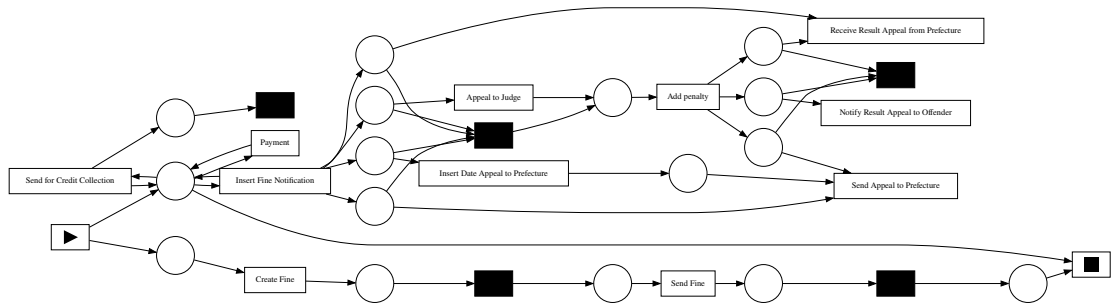


Fig. 21: Model discovered using $\alpha^{++}; 4.0; b0.5; t0.5; r0.5$ on *RTFM*

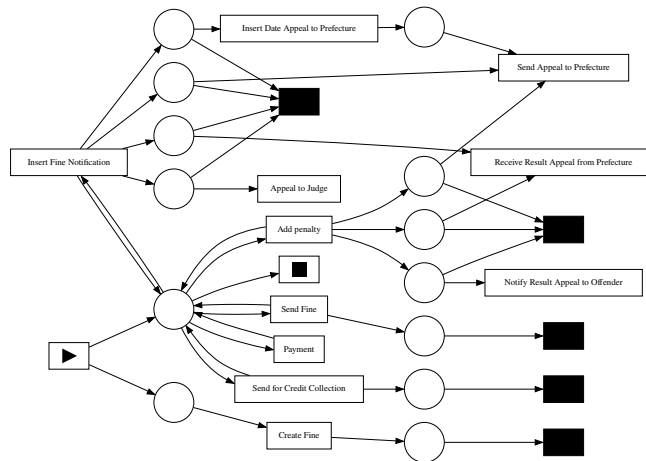


Fig. 22: Model discovered using $\alpha^{++}; 4.0; b0.3; t0.7; r0.6$ on *RTFM*

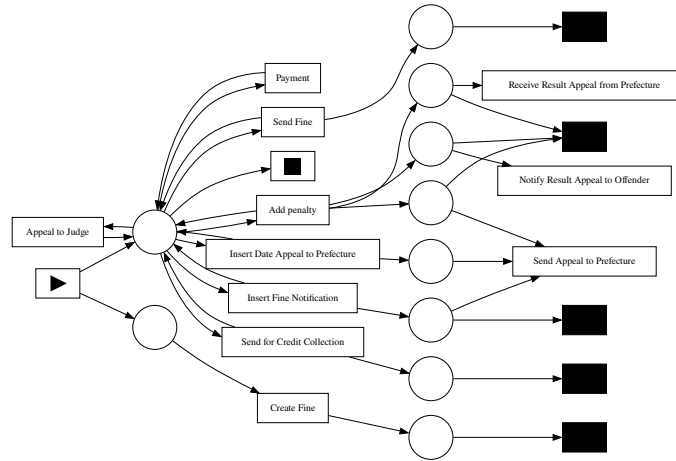


Fig. 23: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.7$ on *RTFM*

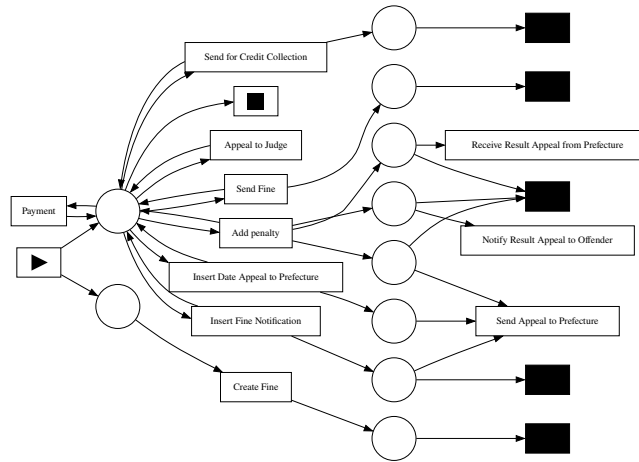


Fig. 24: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.8$ on *RTFM*

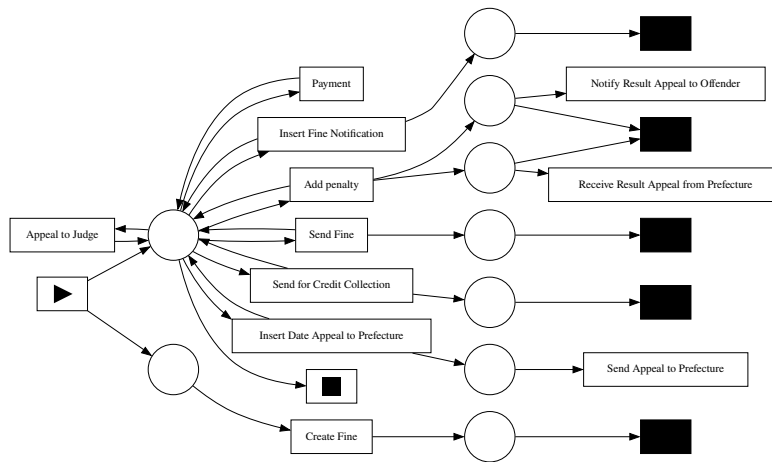


Fig. 25: Model discovered using $\alpha+++;4.0;b0.1;t0.9;r0.9$ on *RTFM*

C.2 Sepsis Cases




Fig. 26: Model discovered using IMf 0.1 on *Sepsis Cases*

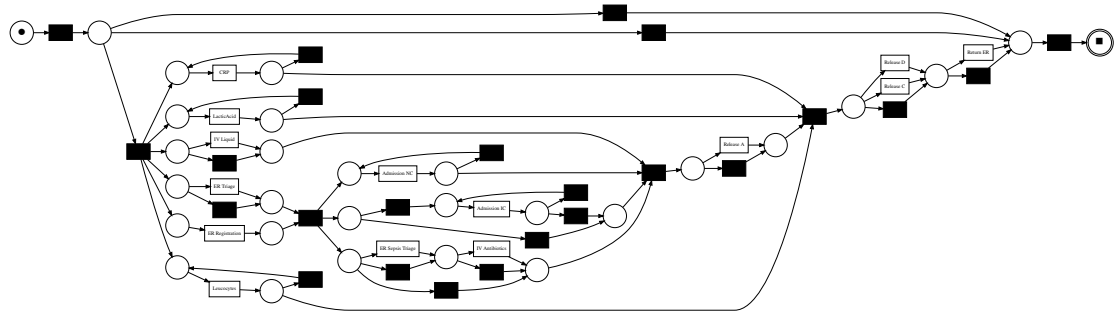


Fig. 27: Model discovered using IMf 0.2 on *Sepsis Cases*

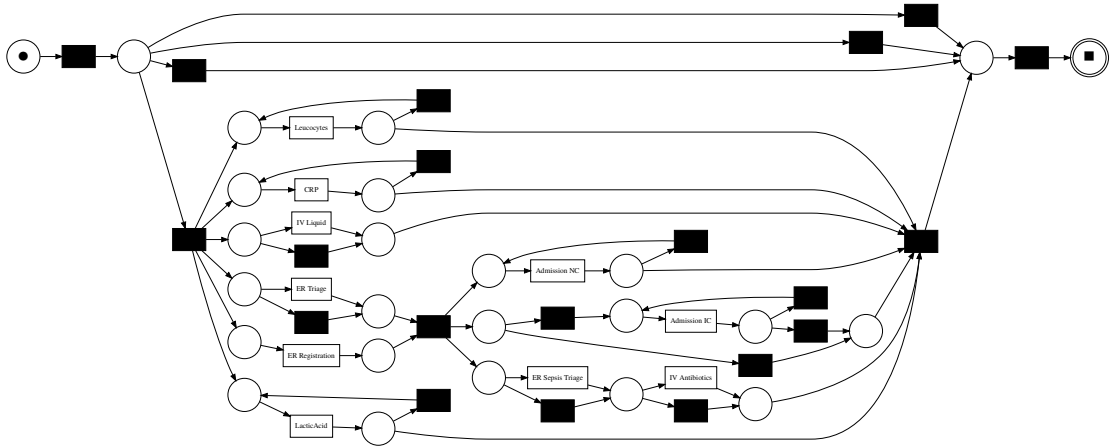


Fig. 28: Model discovered using IMf 0.3 on *Sepsis Cases*

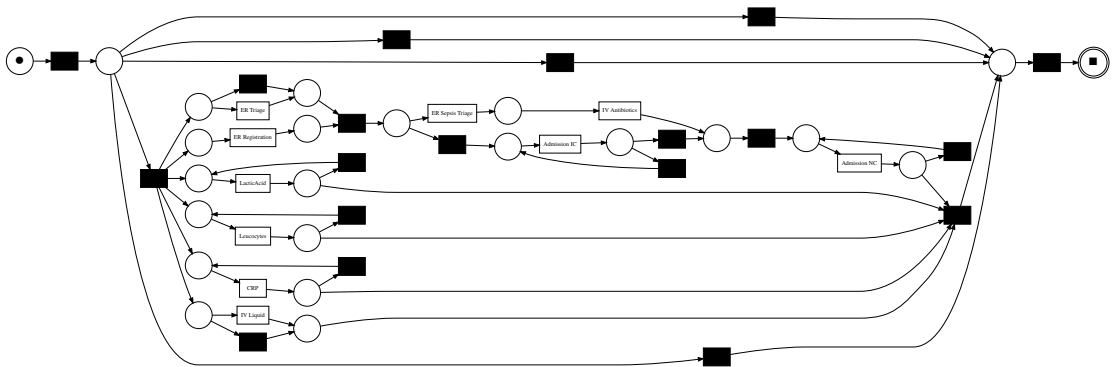


Fig. 29: Model discovered using IMf 0.4 on *Sepsis Cases*




Fig. 30: Model discovered using α Top10 on *Sepsis Cases*




Fig. 31: Model discovered using α 10%Cov on *Sepsis Cases*




Fig. 32: Model discovered using α 50%Cov on *Sepsis Cases*

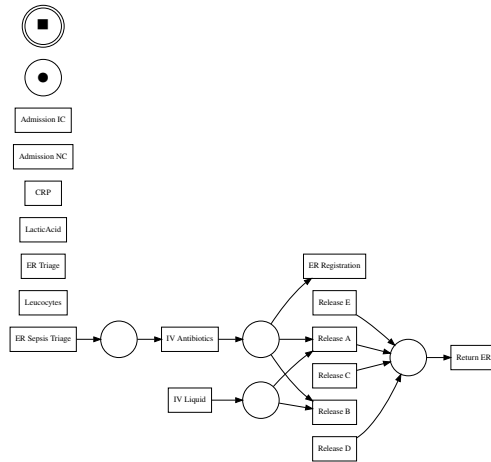


Fig. 33: Model discovered using α 80%Cov on *Sepsis Cases*

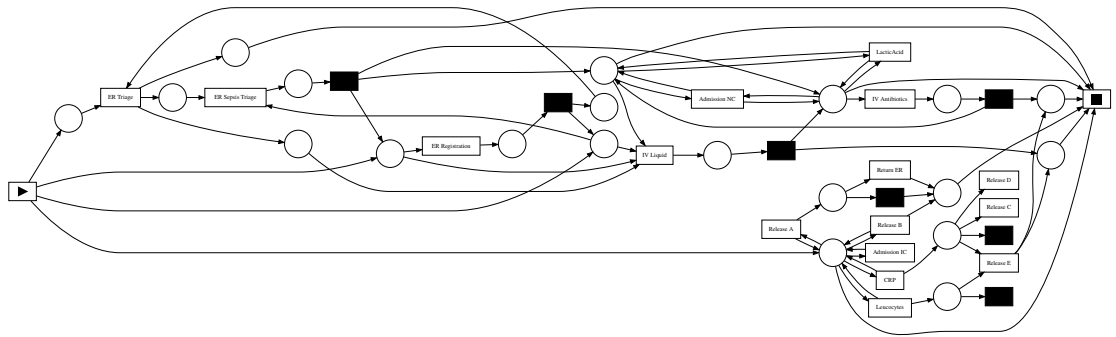


Fig. 34: Model discovered using $\alpha^{++}; 2.0; b0.5; t0.5; r0.5$ on *Sepsis Cases*

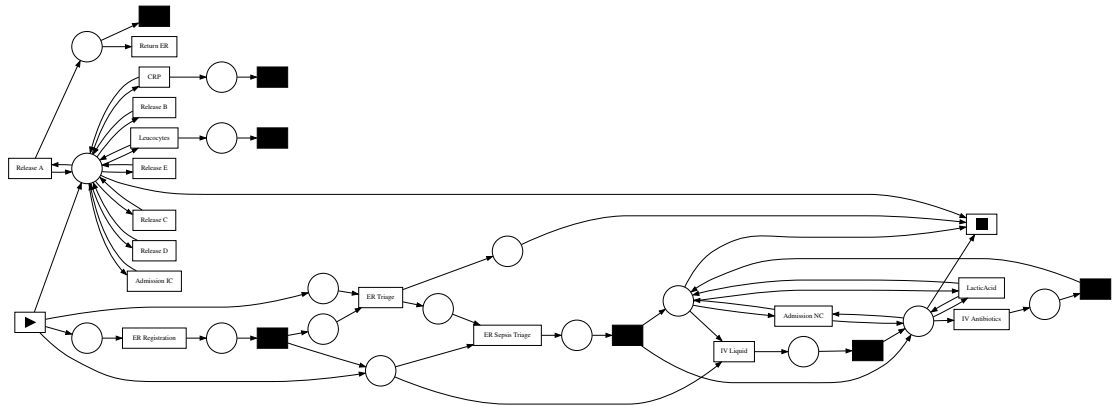


Fig. 35: Model discovered using $\alpha^{++}; 2.0; b0.3; t0.7; r0.6$ on *Sepsis Cases*

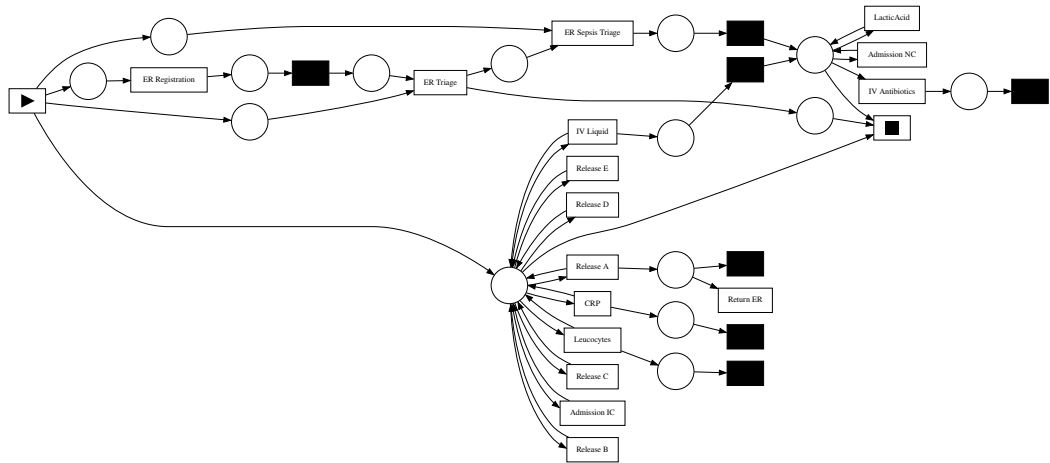


Fig. 36: Model discovered using $\alpha^{++}; 2.0; b0.2; t0.8; r0.7$ on *Sepsis Cases*

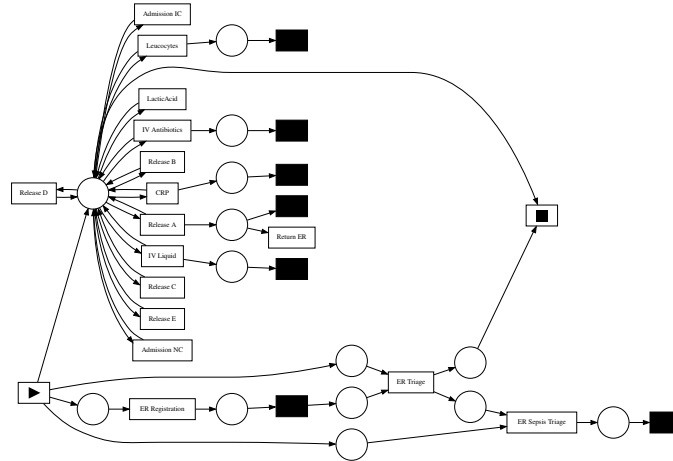


Fig. 37: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.8$ on *Sepsis Cases*

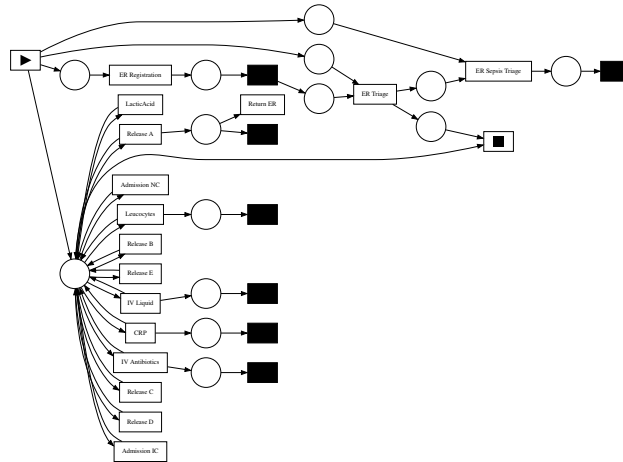


Fig. 38: Model discovered using $\alpha+++;2.0;b0.1;t0.9;r0.9$ on *Sepsis Cases*

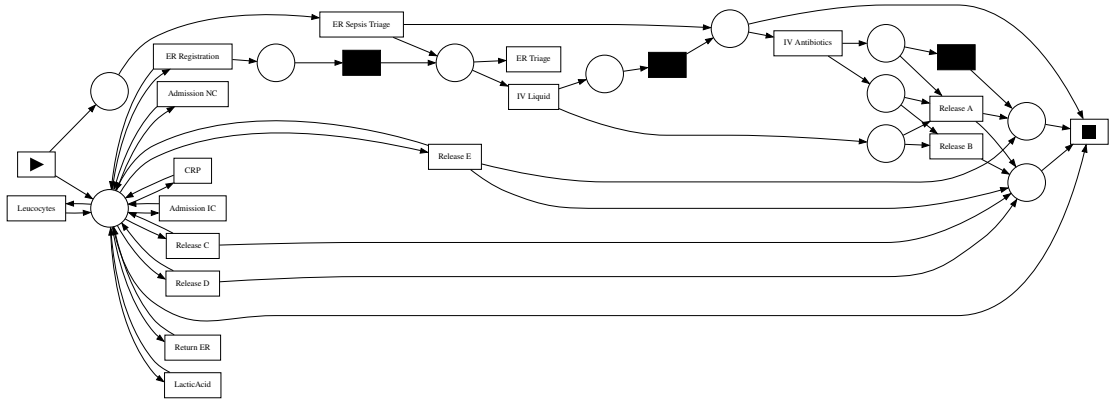


Fig. 39: Model discovered using $\alpha^{++};4.0;b0.5;t0.5;r0.5$ on *Sepsis Cases*




Fig. 40: Model discovered using $\alpha^{++};4.0;b0.3;t0.7;r0.6$ on *Sepsis Cases*




Fig. 41: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.7$ on *Sepsis Cases*




Fig. 42: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.8$ on *Sepsis Cases*

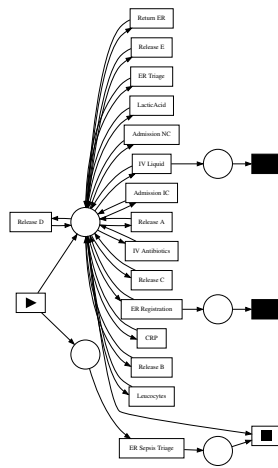


Fig. 43: Model discovered using $\alpha+++;4.0;b0.1;t0.9;r0.9$ on *Sepsis Cases*

C.3 BPI Challenge 2019 (Sample of 3000 Cases)




Fig. 44: Model discovered using IMf 0.1 on *BPI Challenge 2019 (Sample of 3000 Cases)*

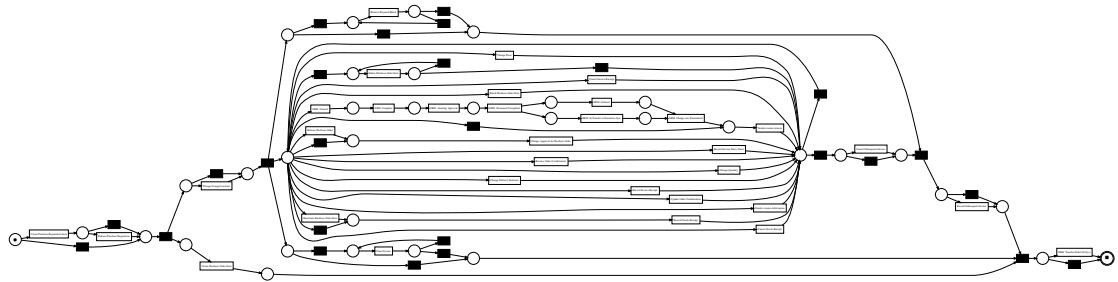


Fig. 45: Model discovered using IMf 0.2 on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 46: Model discovered using IMf 0.3 on *BPI Challenge 2019 (Sample of 3000 Cases)*

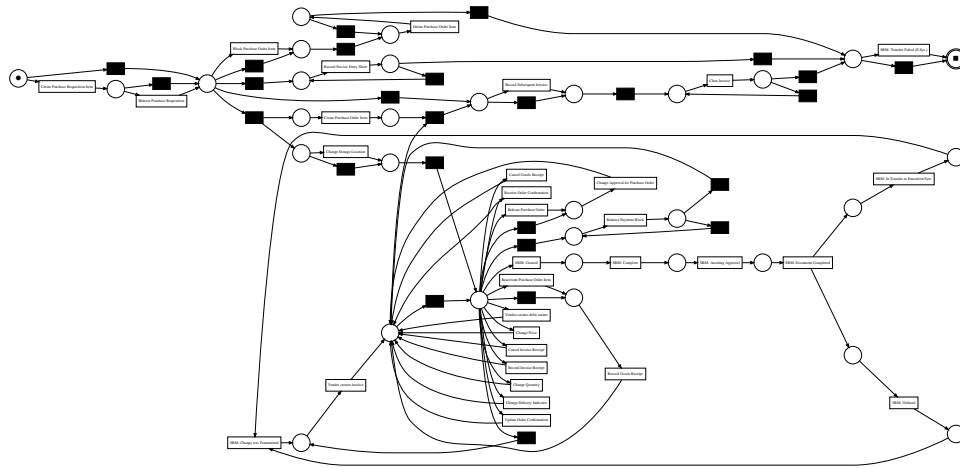


Fig. 47: Model discovered using IMf 0.4 on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 48: Model discovered using α Top10 on *BPI Challenge 2019 (Sample of 3000 Cases)*



Fig. 49: Model discovered using α 10%Cov on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 50: Model discovered using α 50%Cov on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 51: Model discovered using α 80%Cov on *BPI Challenge 2019 (Sample of 3000 Cases)*

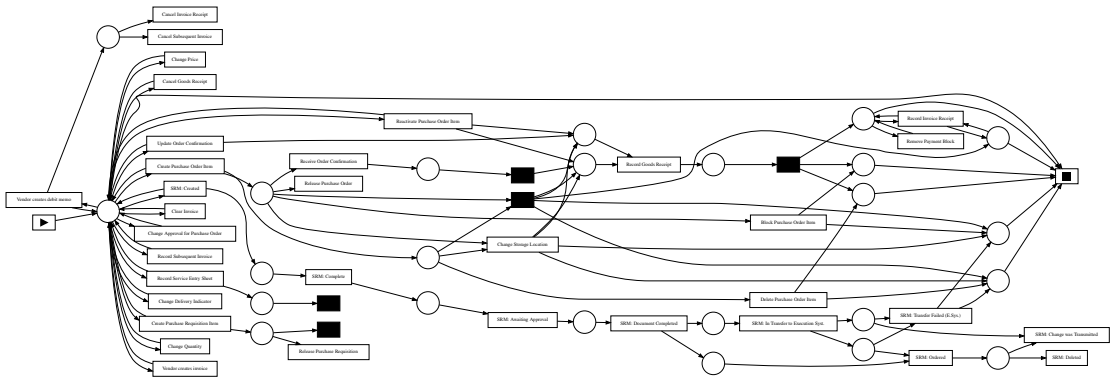


Fig. 52: Model discovered using $\alpha+++; 2.0; b0.5; t0.5; r0.5$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

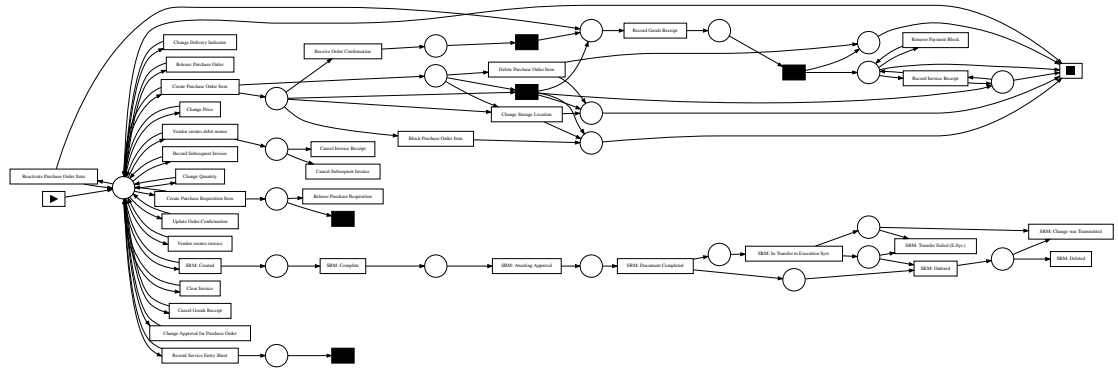


Fig. 53: Model discovered using $\alpha+++;2.0;b0.3;t0.7;r0.6$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

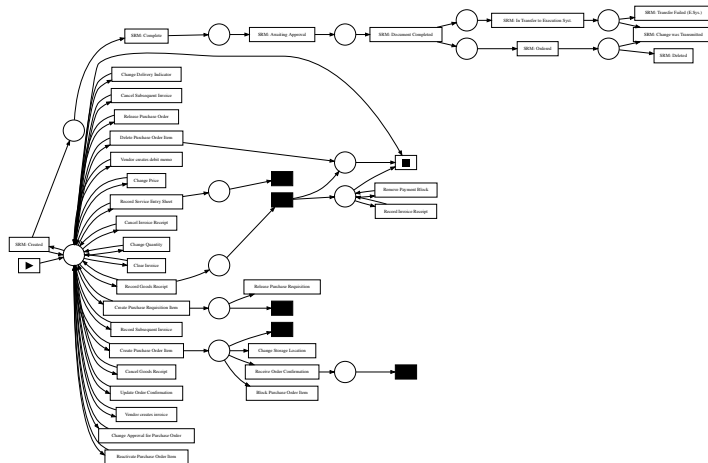


Fig. 54: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.7$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

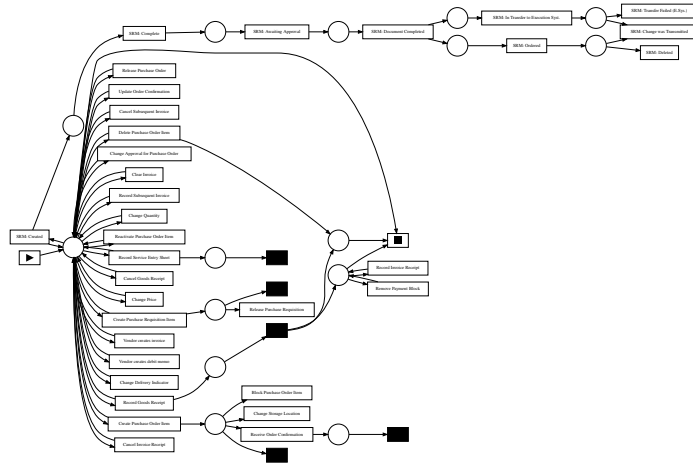


Fig. 55: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.8$ on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 56: Model discovered using $\alpha+++;2.0;b0.1;t0.9;r0.9$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

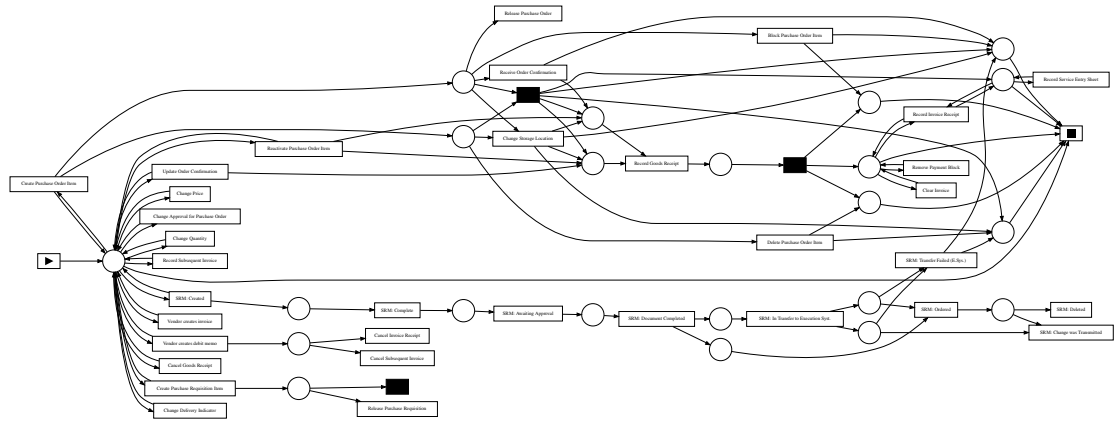


Fig. 57: Model discovered using $\alpha+++;4.0;b0.5;t0.5;r0.5$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

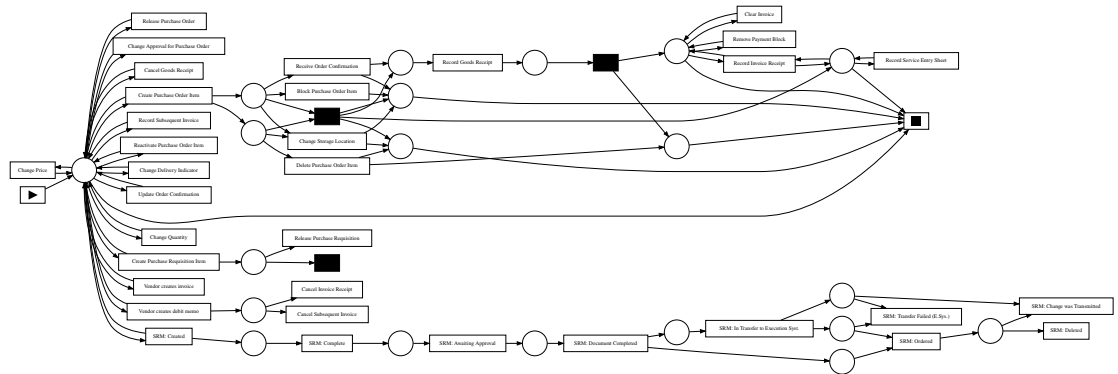


Fig. 58: Model discovered using $\alpha+++;4.0;b0.3;t0.7;r0.6$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

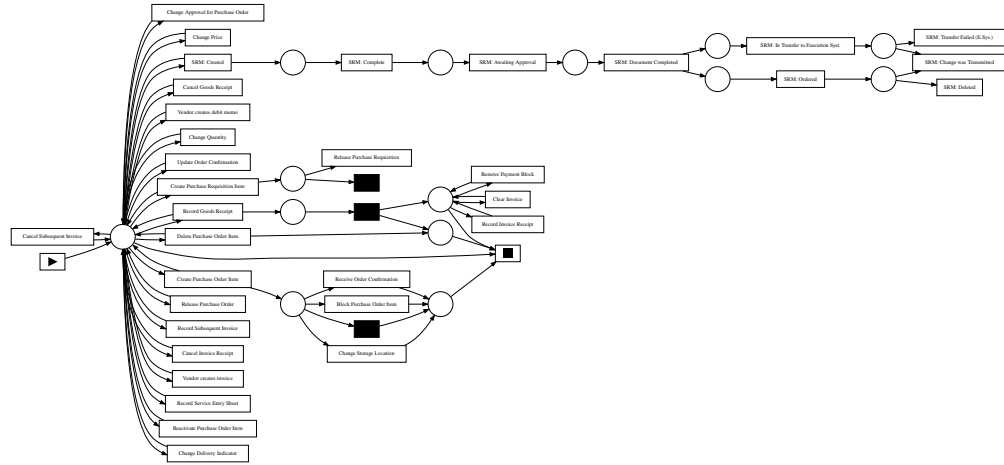


Fig. 59: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.7$ on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 60: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.8$ on *BPI Challenge 2019 (Sample of 3000 Cases)*




Fig. 61: Model discovered using $\alpha+++;4.0;b0.1;t0.9;r0.9$ on *BPI Challenge 2019 (Sample of 3000 Cases)*

C.4 BPI Challenge 2020 (Request for Payment)




Fig. 62: Model discovered using IMf 0.1 on *BPI Challenge 2020 (Request for Payment)*



Fig. 63: Model discovered using IMf 0.2 on *BPI Challenge 2020 (Request for Payment)*




Fig. 64: Model discovered using IMf 0.3 on *BPI Challenge 2020 (Request for Payment)*




Fig. 65: Model discovered using IMf 0.4 on *BPI Challenge 2020 (Request for Payment)*




Fig. 66: Model discovered using α Top10 on *BPI Challenge 2020 (Request for Payment)*




Fig. 67: Model discovered using α 10%Cov on *BPI Challenge 2020 (Request for Payment)*




Fig. 68: Model discovered using α 50%Cov on *BPI Challenge 2020 (Request for Payment)*




Fig. 69: Model discovered using α 80%Cov on *BPI Challenge 2020 (Request for Payment)*




Fig. 70: Model discovered using $\alpha+++;2.0;b0.5;t0.5;r0.5$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 71: Model discovered using $\alpha+++;2.0;b0.3;t0.7;r0.6$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 72: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.7$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 73: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.8$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 74: Model discovered using $\alpha+++;2.0;b0.1;t0.9;r0.9$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 75: Model discovered using $\alpha+++;4.0;b0.5;t0.5;r0.5$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 76: Model discovered using $\alpha+++;4.0;b0.3;t0.7;r0.6$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 77: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.7$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 78: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.8$ on *BPI Challenge 2020 (Request for Payment)*




Fig. 79: Model discovered using $\alpha+++;4.0;b0.1;t0.9;r0.9$ on *BPI Challenge 2020 (Request for Payment)*

C.5 BPI Challenge 2020 (Domestic Declarations)




Fig. 80: Model discovered using IMf 0.1 on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 81: Model discovered using IMf 0.2 on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 82: Model discovered using IMf 0.3 on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 83: Model discovered using IMf 0.4 on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 84: Model discovered using α Top10 on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 85: Model discovered using α 10%Cov on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 86: Model discovered using α 50%Cov on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 87: Model discovered using α 80%Cov on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 88: Model discovered using $\alpha+++;2.0;b0.5;t0.5;r0.5$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 89: Model discovered using $\alpha+++;2.0;b0.3;t0.7;r0.6$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 90: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.7$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 91: Model discovered using $\alpha+++;2.0;b0.2;t0.8;r0.8$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 92: Model discovered using $\alpha+++;2.0;b0.1;t0.9;r0.9$ on *BPI Challenge 2020 (Domestic Declarations)*

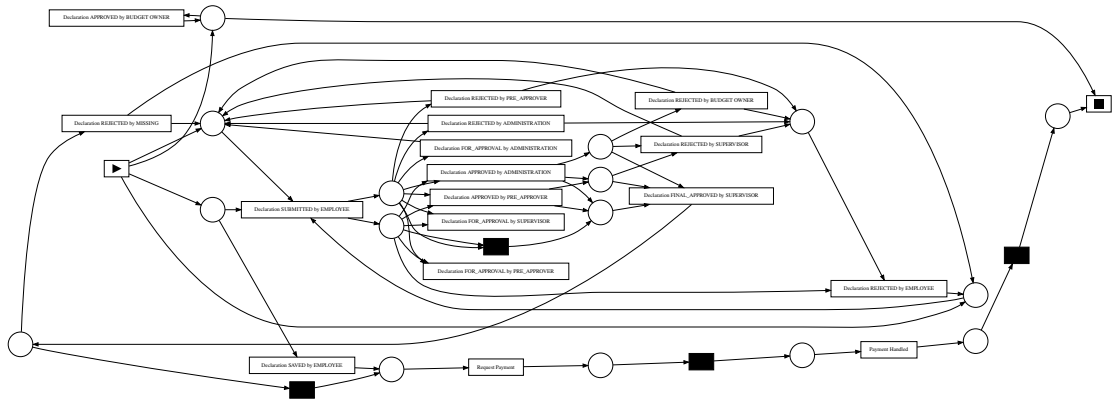


Fig. 93: Model discovered using $\alpha+++;4.0;b0.5;t0.5;r0.5$ on *BPI Challenge 2020 (Domestic Declarations)*

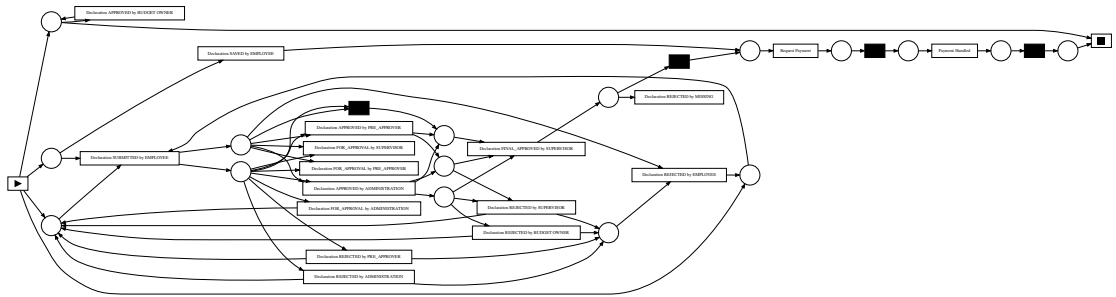


Fig. 94: Model discovered using $\alpha+++;4.0;b0.3;t0.7;r0.6$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 95: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.7$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 96: Model discovered using $\alpha+++;4.0;b0.2;t0.8;r0.8$ on *BPI Challenge 2020 (Domestic Declarations)*




Fig. 97: Model discovered using $\alpha+++;4.0;b0.1;t0.9;r0.9$ on *BPI Challenge 2020 (Domestic Declarations)*