

Mining Uncertain Event Data in Process Mining^{*}

Marco Pegoraro [✉][0000-0002-8997-7517] and Wil M.P. van der Aalst^[0000-0002-0955-6940]

Process and Data Science Group (PADS)
Department of Computer Science, RWTH Aachen University, Aachen, Germany
{pegoraro,wvdaalst}@pads.rwth-aachen.de
<http://www.pads.rwth-aachen.de/>

Abstract. Nowadays, more and more process data are automatically recorded by information systems, and made available in the form of *event logs*. Process mining techniques enable process-centric analysis of data, including automatically discovering process models and checking if event data conform to a certain model. In this paper we analyze the previously unexplored setting of uncertain event logs: logs where quantified uncertainty is recorded together with the corresponding data. We define a taxonomy of uncertain event logs and models, and we examine the challenges that uncertainty poses on process discovery and conformance checking. Finally, we show how upper and lower bounds for conformance can be obtained aligning an uncertain trace onto a regular process model.

1 Introduction

Over the last decades, the concept of *process* has become more and more central in formally describing the activities of businesses, companies and other similar entities, structured in specific steps and phases. A process is thus defined as a well-structured set of activities, possibly performed by multiple actors (*resources*), which contribute to the completion of a specific task or to the achievement of a specific goal.

The processes that govern the innards of business companies are increasingly supported by software tools. Performing specific activities is both aided and recorded by *process-aware information systems* (PAISs), which support the definition and management of processes. The information regarding the execution of processes, which includes time, case and activity information, can then be extracted from PAISs in the form of an *event log*, a database or file containing the digital trace of the operations carried out in the context of the execution of a process and recorded as *events*. The discipline of *process mining* concerns the automatic analysis of event logs, with the goal of extracting knowledge regarding

^{*} In *International Conference on Process Mining* (ICPM 2019). DOI: 10.1109/ICPM.2019.00023. ©IEEE. Please do not print this document unless strictly necessary.

e.g. the structure of the process, the conformity of events to a specific normative process model, the performances in executing the process, the relationships between groups of actors in the process.

In this paper we will consider the analysis of a specific class of event logs: the logs that contain *uncertain event data*. Uncertain events are recordings of executions of specific activities in a process which are enclosed with an indication of uncertainty in the event attributes. Specifically, we consider the case where the attributes of an event are not recorded as a precise value but as a range or a set of alternatives.

The recording of uncertain event data is a common occurrence in process management. The *Process Mining Manifesto* [2] describes a fundamental property of event data as *trustworthiness*, the assumption that the recorded data can be considered correct and accurate. In a general sense, uncertainty as defined here is an explicit absence of trustworthiness, with an indication of uncertainty recorded together with the event data. In the taxonomy of event data proposed in the Manifesto the logs at the two lower levels of quality frequently lack trustworthiness, and thus can be uncertain. This encompasses a wide range of processes, such as event logs of document and product management systems, error logs of embedded systems, worksheets of service engineers, and any process recorded totally or partially on paper. There are many possible causes behind the recording of uncertain event data, such as:

- *Incorrectness*. In some instances, the uncertainty is simply given by errors occurred while recording the data itself. Faults of the information system, or human mistakes in a data entry phase can all lead to missing or altered event data that can be subsequently modeled as uncertain event data.
- *Coarseness*. Some information systems have limitations in their way of recording data - often tied to factors like the precision of the data format - such that the event data can be considered uncertain. A typical example is an information system that only records the date, but not the time, of the occurrence of an event: if two events are recorded in the same day, the order of occurrence is lost. This is an especially common circumstance in the processes that are, partially or completely, recorded on paper and then digitalized. Another factor that can lead to uncertainty in the time of recording is the information system being overloaded and, thus, delaying memorization of data. This type of uncertainty can also be generated by the limited sensibility of a sensor.
- *Ambiguity*. In some cases, the data recorded is not an identifier of a certain event attribute; in these instances, the data needs to be interpreted, either automatically or manually, in order to obtain a value for the event attribute. Uncertainty can arise if the meaning of the data is ambiguous and cannot be interpreted with precision. Example are data in the form of images, text, or video.

Aside from the causes, we can individuate other types of uncertain event logs based on the frequency of uncertain data. Uncertainty can be *infrequent*, when a specific attribute is only seldomly recorded together with explicit uncertainty; the uncertainty is rare enough that uncertain events can be considered outliers.

Conversely, *frequent* uncertain behavior of the attribute is systematic, pervasive in a high number of traces, and thus not to be considered an outlier. The uncertainty can be considered part of the process itself. These concepts are not meant to be formal, and are laid out to distinguish between logs that are still processable regardless of the uncertainty, and logs where the uncertainty is too invasive to analyze them with existing process mining techniques.

In this paper we propose a taxonomy of the different types of explicit uncertainty in process mining, together with a formal, mathematical formulation. As an example of practical application, we will consider the case of conformance checking [5], and we will apply it to uncertain data by assessing what are the upper and lower bounds on the conformance score for possible values of the attributes in an uncertain trace.

The rest of this paper is organized as follows. Section II discusses previous and related work in the management of uncertain data. Section III proposes a taxonomy of the different possible types of uncertain process data. Section IV contains the formal definitions needed to manage uncertainty. Section V describes a practical application of process mining over uncertain event data, the case of conformance checking through alignments. Section VI shows experimental results on computing conformance checking scores for uncertain data. Finally, Section VII concludes the paper and discusses about future work.

2 Related Work

As mentioned, the occurrence of data containing uncertainty - in a broad sense - is common both in more classic disciplines like statistics and Data Mining [8] and in process mining [2]; and logs that show an explicit uncertainty in the control flow perspective can be classified in the lower levels of the quality ranking proposed in the process mining manifesto.

Within process mining there exist various techniques to deal with a kind of uncertainty different from the one that we analyze here: missing or incorrect data. This can be considered as a form of non-explicit uncertainty: no measure or indication on the nature of the uncertainty is given in the event log. The work of Suriadi et al. [11] provides a taxonomy of this type of issues in event logs, laying out a series of data patterns that model errors in process data. In these cases, and if this behavior is infrequent enough to allow the event log to remain meaningful, the most common way for existing process mining techniques to deal with missing data is by filtering out the affected traces and performing discovery and conformance checking on the resulting filtered event log. While filtering out missing values is straightforward, various methodologies of event log filtering have been proposed in the past to solve the problem of incorrect event attributes: the filtering can take place thanks to a reference model, which can be given as process specification [12], or from information discovered from the frequent and well-formed traces of the same event log; for example extracting an automaton from the frequent traces [7], computing conditional probabilities of

frequent sequences of activities [9], or discovering a probabilistic automaton [13]. In the latter cases, the noise is identified as infrequent behavior.

Some previous work attempt to repair the incorrect values in an event log. Conforti et al. [6] propose an approach for the restoration of incorrect timestamps based on a log automaton, that repairs the total ordering of events in a trace based on correct frequent behavior. Fani Sani et al. [10] define outlier behavior as the unexpected occurrence of an event, the absence of an event that is supposed to happen, and the incorrect order of events in the trace; then, they propose a repairing method based on probabilistic analysis of the context of an outlier (events preceding or following the anomalous event). Again, both of these methods define anomalous/incorrect behavior on the basis of the frequency of occurrence.

The main driving reasons behind this work is to provide the means to treat uncertainty as a relevant part of a process; thus, we aim not to filter it out but model it. In conclusion, there are two novel aspects regarding uncertain data that we intend to address in this work. The first is the *explicitness of uncertainty*: we work with the underlying assumption that the actual value of the uncertain attribute, while not directly provided, is described formally. This is the case when meta-information about the uncertainty in the attribute is available, either deduced from the features of the information system(s) that record the logs or included in the event log itself. Note that, as opposed to all previous work on the topic, the fact that uncertainty is explicit in the data means that the concept of uncertain behavior is completely separated from the concept of infrequent behavior. The second is the goal of *modeling uncertainty*: we consider uncertainty part of the process. Instead of filtering or cleaning the log we introduce the uncertainty perspective in process mining by extending the currently available techniques to incorporate it.

3 A Taxonomy of Uncertain Event Data

The goal of this section of the paper is to propose a categorization of the different types of uncertainty that can appear in process mining. In process management, a central concept is the distinction between the data perspective (the event log) and the behavioral perspective (the process model). The first one is a static representation of process instances, the second summarizes the behavior of a process. Both can be extended with a concept of explicit uncertainty: this concept also implies an extension of the process mining techniques that have currently been implemented.

In this paper we will focus on uncertainty in event data, while the concept of uncertainty applied to models will be examined in a future work. Specifically, as an example application we will consider computing the conformance score of uncertain process data on classical models.

We can individuate two different notions of uncertainty:

- *Strong uncertainty*: the possible values for the attributes are known, but the probability that the attribute will assume a certain instantiation is unknown or unobservable.
- *Weak uncertainty*: both the possible values of an attribute and their respective probabilities are known.

In the case of a discrete attribute, the strong notion of uncertainty consists on a set of possible values assumed by the attribute. In this case, the probability for each possible value is unknown. Vice-versa, in the weak uncertainty scenario we also have a discrete probability distribution defined on that set of values. In the case of a continuous attribute, the strong notion of uncertainty can be represented with an interval for the variable. Notice that an interval do not indicate a uniform distribution; there is no information on the likelihood of values in it. Vice-versa, in the weak uncertainty scenario we also have a probability density function defined on a certain interval. Figure 1 summarizes this concepts. This leads to very simple representations of explicit uncertainty.

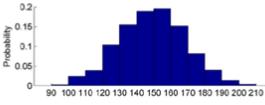
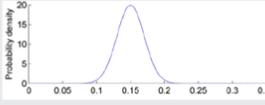
	Weak uncertainty	Strong uncertainty
Discrete data	Discrete probability distribution 	Set of possible values $\{x, y, z, \dots\}$
Continuous data	Probability density function 	Interval $\{x \in \mathbb{R} a \leq x \leq b\}$

Fig. 1. The four different types of uncertainty.

In this paper we consider only the control flow and time perspective of a process – namely, the attributes of the events that allow to discover a process model. These are the unique identifier of a process instance (case ID), the timestamp (often represented by the distance from a fixed origin point, e.g. the *Unix Epoch*), and the activity identifier of an event. Case IDs and activities are values chosen from a finite set of possible values; they are discrete variables. Timestamps, instead, are represented by numbers and thus are continuous variables.

We will also describe an additional type of uncertainty, which lays on the event level rather than the attribute level:

- *Indeterminate event*: there is a chance that the event did not take place even though it was recorded in the event log. Indeterminate events are indicated with a ? symbol, while determinate (regular) events are marked with a ! symbol.

Table 1. An example of strongly uncertain trace.

Case ID	Timestamp	Activity	Indet. event
{0, 1}	2011-12-05T00:00	A	!
0	2011-12-07T00:00	{B, C, D}	!
0	[2011-12-06T00:00, 2011-12-10T00:00]	D	?
0	2011-12-09T00:00	{A, C}	!
{0, 1, 2}	2011-12-11T00:00	E	?

Table 2. An example of weakly uncertain trace.

Case ID	Timestamp	Activity	Indet. event
{0:0.9, 1:0.1}	2011-12-05T00:00	A	!
0	2011-12-07T00:00	{B:0.7, C:0.3}	!
0	$\mathcal{N}(2011-12-08T00:00, 2)$	D	?:0.5
0	2011-12-09T00:00	{A:0.2, C:0.8}	!
{0:0.4, 1:0.6}	2011-12-11T00:00	E	?:0.7

Examples of strongly and weakly uncertain traces are shown in Tables 1 and 2 respectively.

4 Definitions

Let us now provide a formal definition of the concept of uncertainty applied to event data.

Definition 1 (Power Set). *The power set of a set A is the set of all possible subsets of A , and is denoted with $\mathcal{P}(A)$. $\mathcal{P}_{NE}(A)$ denotes the set of all the non-empty subsets of A : $\mathcal{P}_{NE}(A) = \mathcal{P}(A) \setminus \{\emptyset\}$.*

Definition 2 (Multiset). *A multiset is an extension of the concept of set that keeps track of the cardinality of each element. $\mathcal{B}(A)$ is the set of all multisets over some set A . Multisets are denoted with square brackets, e.g. $b = [x, x, y]$.*

Definition 3 (Sequence). *Given a set X , a finite sequence over X of length n is a function $s \in X^* : \{1, \dots, n\} \rightarrow X$, and it is written as $s = \langle s_1, s_2, \dots, s_n \rangle$. Over the sequence s we define $|s| = n$, $s[i] = s_i$ and $x \in s \iff x \in \text{set}(s)$.*

Definition 4 (Universes). Let \mathcal{U}_E be the set of all the event identifiers. Let \mathcal{U}_C be the set of all the case id identifiers. Let \mathcal{U}_A be the set of all the activity identifiers. Let \mathcal{U}_T be the totally ordered set of all the timestamp identifiers. Let $\mathcal{U}_O = \{!, ?\}$, where the “!” symbol denotes determined events, and the “?” symbol denotes indeterminate events.

Definition 5 (Events). Let us denote with $\mathcal{E}_C = \mathcal{U}_E \times \mathcal{U}_C \times \mathcal{U}_A \times \mathcal{U}_T$ the universe of certain events. $\mathcal{E}_{SU} = \mathcal{U}_E \times \mathcal{P}_{NE}(\mathcal{U}_C) \times \mathcal{P}_{NE}(\mathcal{U}_A) \times \mathcal{P}_{NE}(\mathcal{U}_T) \times \mathcal{U}_O$ is the universe of strongly uncertain events. $\mathcal{E}_{WU} = \{(e, f) \in \mathcal{U}_E \times (\mathcal{U}_C \times \mathcal{U}_A \times \mathcal{U}_T \not\rightarrow [0, 1]) \mid \sum_{(a,c,t) \in \text{dom}(f)} f(c, a, t) \leq 1\}$ is the universe of weakly uncertain events. Over a strongly uncertain event $(e, c_s, a_s, t_s, u) \in \mathcal{E}_{SU}$ we define the following projection functions: $\pi_c^{\mathcal{E}_{SU}}(e) = c_s$, $\pi_a^{\mathcal{E}_{SU}}(e) = a_s$, $\pi_t^{\mathcal{E}_{SU}}(e) = t_s$ and $\pi_o^{\mathcal{E}_{SU}}(e) = o$.

Definition 6 (Event logs). A certain event log is a set of events $L_C \subseteq \mathcal{E}_C$ such that every event identifier in L_C is unique. A strongly uncertain event log is a set of events $L_{SU} \subseteq \mathcal{E}_{SU}$ such that every event identifier in L_{SU} is unique. A weakly uncertain event log is a set of events $L_{WU} \subseteq \mathcal{E}_{WU}$ such that every event identifier in L_{WU} is unique.

A weakly uncertain event log $L_{WU} \subseteq \mathcal{E}_{WU}$ has a corresponding strongly uncertain event log $\overline{L_{WU}} = L_{SU} \subseteq \mathcal{E}_{SU}$ such that $L_{SU} = \{(e, c_s, a_s, t_s, o) \in \mathcal{E}_{SU} \mid \exists_{(e', f) \in L_{WU}}, e = e' \wedge c_s = \{c \in \mathcal{U}_C \mid \exists_{a,t}((c, a, t) \in \text{dom}(f) \wedge f(c, a, t) > 0)\} \wedge a_s = \{a \in \mathcal{U}_A \mid \exists_{c,t}((c, a, t) \in \text{dom}(f) \wedge f(c, a, t) > 0)\} \wedge t_s = \{t \in \mathcal{U}_T \mid \exists_{c,a}((c, a, t) \in \text{dom}(f) \wedge f(c, a, t) > 0)\} \wedge (o = ! \iff (\sum_{(c,a,t) \in \text{dom}(f)} f(c, a, t)) = 1) \wedge (o = ? \iff (\sum_{(c,a,t) \in \text{dom}(f)} f(c, a, t)) < 1)\}$.

Definition 7 (Realization of an event log). $L_C \subseteq \mathcal{E}_C$ is a realization of $L_{SU} \subseteq \mathcal{E}_{SU}$ if and only if:

- For all $(e, c, a, t) \in L_C$ there is a distinct $(e', c_s, a_s, t_s, o) \in L_{SU}$ such that $e' = e$, $a \in a_s$, $c \in c_s$ and $t \in t_s$;
- For all $(e, c_s, a_s, t_s, o) \in L_{SU}$ with $o = !$ there is a distinct $(e', c, a, t) \in L_C$ such that $e' = e$, $a \in a_s$, $c \in c_s$ and $t \in t_s$.

$\mathcal{R}_L(L_{SU})$ is the set of all such realizations of the log L_{SU} .

Note that these definition allow us to transform a weakly uncertain log into a strongly uncertain one, and a strongly uncertain one in a set of certain logs.

5 Conformance Checking on Uncertain Event Data

As a preliminary application of process mining over uncertain event data we now focus on conformance checking. Starting from an event log and a process model, conformance checking verifies if the event data in the log conforms to the model, providing a diagnostic of the deviations. Conformance checking serves

many purposes, such as checking if process instances follow a specific normative model, assessing if a certain execution log has been generated from a specific model, or verifying the quality of a process discovery technique.

The specific scenario we consider in this paper includes:

- Strong uncertainty on the activity;
- Strong uncertainty on the timestamp;
- Strong uncertainty on indeterminate events.

All three can happen concurrently. Table 3 shows such a trace, which we will use as running example. It is worth noticing that the specific case of uncertainty on the case ID causes a problem; since an event can have many possible case IDs, it can belong to different traces. In data format where the event are already aggregated into traces, such as the very common XES standard, this means that the information related to a trace can be *non local* to the trace itself, but can be stored in some other points of the log. We will focus on the problem of uncertainty on the case ID attribute in a future work.

Firstly, we will lay down some simplified notation in order to model in a more compact way the problem at hand.

Definition 8 (Simple traces and logs). $\sigma_C \in \mathcal{U}_A^*$ is a simple untimed trace. \mathcal{T}_C denotes the universe of simple untimed traces. $L_C^S \in \mathcal{B}(\mathcal{T}_C)$ is a simple untimed log.

$\sigma_{CT} \in (\mathcal{U}_A \times \mathcal{U}_T)^*$ is a simple timed trace if and only if $\sigma_{CT} = \langle (a_1, t_1), (a_2, t_2), \dots, (a_n, t_n) \rangle$ and $\forall_{1 \leq i < j \leq n}, t_i < t_j$. \mathcal{T}_{CT} denotes the universe of simple timed traces. $L_{CT}^S \in \mathcal{B}(\mathcal{T}_{CT})$ is a simple timed log.

$\sigma_U \in \mathcal{P}(\mathcal{U}_E \times \mathcal{P}_{NE}(\mathcal{U}_A) \times \mathcal{U}_T \times \mathcal{U}_T \times \mathcal{U}_O)$ is a simple uncertain trace if for all $(e, a_s, t_{min}, t_{max}, u) \in \sigma_U$, $t_{min} < t_{max}$ and all the event identifiers are unique. \mathcal{T}_U denotes the universe of simple uncertain traces. $L_U^S \in \mathcal{B}(\mathcal{T}_U)$ is a simple uncertain log if all the event identifiers in the log are unique. For $e_U^S = (e, a_s, t_{min}, t_{max}, o) \in \sigma_U$ we define the following projection functions: $\pi_a^{L_U^S}(e_U^S) = a_s$, $\pi_{t_{min}}^{L_U^S}(e_U^S) = t_{min}$, $\pi_{t_{max}}^{L_U^S}(e_U^S) = t_{max}$ and $\pi_o^{L_U^S}(e_U^S) = o$.

Definition 9 (Realization of a simple trace). $\sigma_{CT} \in \mathcal{T}_{CT}$ is a timed realization of $\sigma_U \in \mathcal{T}_U$ if and only if:

- For all $(a, t) \in \sigma_{CT}$ there is a distinct $(a_s, t_{min}, t_{max}, u) \in \sigma_U$ such that $a \in a_s$ and $t_{min} \leq t \leq t_{max}$;
- For all $(a_s, t_{min}, t_{max}, u) \in \sigma_U$ with $u = !$ there is a distinct $(a, t) \in \sigma_{CT}$ such that $a \in a_s$ and $t_{min} \leq t \leq t_{max}$.

We denote with $\mathcal{R}_T(\sigma_U)$ the set of all such timed realizations of the trace σ_U . For $\sigma_{CT} \in \mathcal{T}_{CT}$ we denote with $\pi_A(\sigma_{CT})$ the simple untimed trace $\sigma_C \in \mathcal{T}_C$ such that $|\sigma_{CT}| = |\sigma_C| = n$ and for all $1 \leq i \leq n$, $\sigma_{CT}(i) = (a, t): \sigma_C(i) = a$. For $\sigma_U \in \mathcal{T}_U$, $\mathcal{R}(\sigma_U) = \{\pi_A(\sigma_{CT}) \mid \sigma_{CT} \in \mathcal{R}_T(\sigma_U)\}$ is the set of all (untimed) realizations of σ_U .

These simplified traces and logs can be related to the more general framework described in the previous section through the following transformation: let $L_{SU} \subseteq \mathcal{E}_{SU}$ be a strongly uncertain log and let $g: \mathcal{U}_E \rightarrow \mathcal{U}_C$ be a function mapping events onto cases such that $dom(g) = \{e \mid (e, c_s, a_s, t_s, u) \in L_{SU}\}$ and for all $(e, c_s, a_s, t_s, u) \in L_{SU}$, $g(e) \in c_s$. Thus, for $c \in rng(g)$, $g^{-1}(c) = \{e \in \mathcal{U}_E \mid g(e) = c\}$. The simple uncertain event log defined by g on L_{SU} is $L_U^S = [\{(e, \pi_a^{\mathcal{E}_{SU}}(e), \min(\pi_t^{\mathcal{E}_{SU}}(e)), \max(\pi_t^{\mathcal{E}_{SU}}(e)), \pi_o^{\mathcal{E}_{SU}}(e)) \mid e \in g^{-1}(c)\} \mid c \in rng(g)]$.

The conformance checking algorithm that we are applying in this paper is based on *alignments*. Introduced by Adriansyah [3], conformance checking through alignments finds deviations between a trace and a Petri net model of a process by creating a correspondence between the sequence of activities executed in the trace and the firing of the transitions in the Petri net. An example of alignments is given in Figure 2.

Definition 10 (System Net). *A system net is a tuple $SN = (P, T, F, l, M_{init}, M_{final})$ with P the set of places, T the set of transitions, $P \cap T = \emptyset$, $F \subseteq (P \times T) \cup (T \times P)$ the flow relation and $l \in T \rightarrow \mathcal{U}_A$ a labeling function over transitions. A marking $M \in \mathcal{B}(P)$ is a multiset of places; $M_{init} \in \mathcal{B}(P)$ is the initial marking of the net, and $M_{final} \in \mathcal{B}(P)$ is the final marking of the net. \mathcal{U}_{SN} is the universe of system nets.*

A system net SN defines a directed graph with nodes $P \cup T$ and edges F . For any $x \in P \cup T$, $\bullet x = \{y \mid (y, x) \in F\}$ denotes the set of input nodes and $x \bullet = \{y \mid (x, y) \in F\}$ denotes the set of output nodes. A transition $t \in T$ is *enabled* in marking M of net SN , denoted as $(SN, M)[t]$, if each of its input places $\bullet t$ contains at least one token. An enabled transition t may *fire*, i.e., one token is removed from each of the input places $\bullet t$ and one token is produced for each of the output places $t \bullet$. If $t \notin dom(l)$, it is called *invisible*. To indicate invisible transitions we use the placeholder symbol τ ; by definition $\tau \notin dom(l)$. An occurrence of visible transition $t \in dom(l)$ corresponds to observable activity $l(t)$. Given a system net, $\phi(SN)$ is the set of all possible *visible* activity sequences, i.e. the labels of complete firing sequences starting in M_{init} and ending in M_{final} projected onto the set of observable activities. Given the set of activity sequences $\phi(SN)$ obtainable via complete firing sequences on a certain system net, we can define a perfectly fitting event log as a set of traces which activity projection is contained in $\phi(SN)$.

These definitions allow us to build *alignments* in order to compute the fitness of trace on a certain model. An alignment is a correspondence between a sequence of activities (extracted from the trace) and a sequence of transitions with the relative labels (fired in the model while replaying the trace). The first sequence indicates the “moves in the log” and the second indicates the “moves in the model”. If a move in the model cannot be mimicked by a move in the log, then a “ \gg ” (“no move”) appears in the top row; conversely, if a move in the log cannot be mimicked by a move in the model, then a “ \gg ” (“no move”) appears in the bottom row. “no moves” not corresponding to invisible transitions point to deviations between model and log. A *move* is a pair $(x, (y, t))$ where the first

element refers to the log and the second element to the model. A “ \gg ” in the first element of the pair indicates a move on model, in the second element it indicates a move on log.

An alignment is a sequence of moves such that after removing all “ \gg ” symbols, the top row corresponds to a trace in the log and the bottom row corresponds to a firing sequence starting in M_{init} and ending M_{final} . Notice that if $t \notin \text{dom}(l)$ is an invisible transition, the activation of t is indicated by a “ \gg ” on the log in correspondence of t and the placeholder label τ . Hence, the middle row corresponds to a visible path when ignoring the τ steps. Figure 2 shows a model with two examples of alignments, one of a fitting trace and the other of a non-fitting trace.

Definition 11 (Alignment). *Let $\sigma_C \in L_C$ be a trace and $t \in \phi_f(SN)$ a complete firing sequence of system net SN . An alignment of σ_C and t_* is a sequence of moves $\gamma \in A_M^*$ such that the projection on the first element (ignoring \gg) yields σ and the projection on the last element (ignoring \gg and transition labels) yields t .*

A trace and a model can have several possible alignments. In order to select the most appropriate one, we introduce a function that associate a *cost* to undesired moves - the ones associated with deviations.

Definition 12 (Cost of Alignment). *Cost function $\delta \in A_{LM} \rightarrow \mathbb{N}$ assigns costs to legal moves. The cost of an alignment $\gamma \in A_{LM}^*$ is the sum of all costs: $\delta(\gamma) = \sum_{(x,y) \in \gamma} \delta(x,y)$.*

In this paper we use a standard cost function δ_S that assigns cost zero to synchronous moves and moves on invisible transitions, and unit costs to moves on log or moves on model.

Definition 13 (Optimal Alignment). *Let $L_C \in \mathcal{B}(\mathcal{T}_C)$ be a simple untimed event log and let $SN \in \mathcal{U}_{SN}$ be a system net with $\phi(SN) \neq \emptyset$.*

- For $\sigma_C \in L_C$, we define: $\Gamma_{\sigma_C, SN} = \{\gamma \in A_{LM}^* \mid \exists t_* \in \phi_f(SN) \text{ } \gamma \text{ is an alignment of } \sigma_C \text{ and } t_*\}$.
- An alignment $\gamma \in \Gamma_{\sigma_C, SN}$ is optimal for trace $\sigma_C \in L_C$ and system net SN if for any $\gamma' \in \Gamma_{\sigma_C, SN}$: $\delta(\gamma') \geq \delta(\gamma)$.
- $\lambda_{SN} \in \mathcal{E}^* \rightarrow A_{LM}^*$ is a deterministic mapping that assigns any trace σ_C to an optimal alignment, i.e., $\lambda_{SN}(\sigma_C) \in \Gamma_{\sigma_C, SN}$ and $\lambda_{SN}(\sigma_C)$ is optimal.
- $\text{costs}(L_C, SN, \delta) = \sum_{\sigma_C \in L} \delta(\lambda_{SN}(\sigma_C))$ are the misalignment costs of the whole event log.

Depending on the possible values for a_s , t_{min} , t_{max} , and u there are multiple possible realizations of a trace. This means that, given a model, a simple uncertain trace could be fitting for certain realizations, but non-fitting for others. The question we are interested in answering is: given a simple uncertain trace and a Petri net process model, is it possible to find an *upper and lower bound* for the conformance score? More formally, when usually we are interested in the

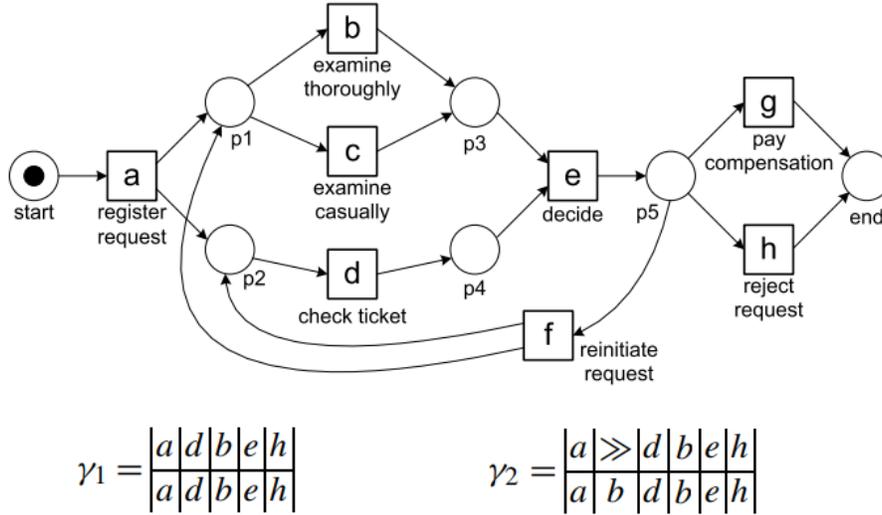


Fig. 2. Example of alignments on a model. The alignment γ_1 shows that the trace $\langle a, d, b, e, h \rangle$ is perfectly fitting the model. The alignment γ_2 shows that the trace $\langle a, b, d, b, e, h \rangle$ is misaligned with the model in one point.

optimal alignments (the ones with the minimal cost), we are now interested in the minimum and maximum cost of alignments in the realization set of a simple uncertain trace.

Table 3. The uncertain trace used as running example for the application of Conformance Checking on uncertainty.

Case ID	Timestamp	Activity	Indet. event
0	2011-12-05T00:00	A	!
0	2011-12-07T00:00	{B, C}	!
0	[2011-12-06T00:00 2011-12-10T00:00]	D	!
0	2011-12-09T00:00	{A, C}	!
0	2011-12-11T00:00	E	?

Definition 14 (Upper and Lower Bound on Alignment Cost for a Trace).

Let $\sigma_U \in \mathcal{T}_U$ be a simple uncertain trace, and let $SN \in \mathcal{U}_{SN}$ be a system net. The upper bound for the alignment cost is a function $\delta_{max}: \mathcal{T}_U \rightarrow \mathbb{N}$ such that $\delta_{max}(\sigma_U) = \max_{\sigma_C \in \mathcal{R}(\sigma_U)} \lambda_{SN}(\sigma_C)$. The lower bound for the alignment cost is a function $\delta_{min}: \mathcal{T}_U \rightarrow \mathbb{N}$ such that $\delta_{min}(\sigma_U) = \min_{\sigma_C \in \mathcal{R}(\sigma_U)} \lambda_{SN}(\sigma_C)$.

A simple way to compute the upper and lower bounds for the cost of an uncertain trace is using a bruteforce approach: enumerating the possible realizations of the trace, then searching for the costs of optimal alignments for all the realizations, and picking the minimum and maximum as bounds.

The technique to compute the optimal alignment [3] is as follows. Firstly, it creates an *event net*, a sequence-structured system net able to replay only the trace to align. The transitions in the event net have labels corresponding to the activities in the trace. Then, a *product net* should be computed; it is the union of the event net and the model together with synchronous transitions added. These additional transitions are paired with transitions in the event net and in the process model that have the same label; they are then connected with arcs from the input places and to the output places of those transitions. The product net is able to represent moves on log, moves on model and synchronous moves by means of firing transitions: the transitions of the event net correspond to moves on log, the transitions of the process model correspond to moves on model, the added synchronous transitions correspond to synchronous moves. The union of the initial and final markings of the event net and the process model constitute respectively the initial and final marking of the product net: every complete firing sequence on the product net corresponds to a possible alignment. Lastly, the product net is translated to a state space, and a state space exploration via the \mathbb{A}^* algorithm is performed in order to find the complete firing sequence that yields the lowest cost.

Let us define formally the construction of the event net and the product net:

Definition 15 (Event Net). *Let $\sigma_C \in \mathcal{T}_C$ be a simple untimed trace. The event net $en : \mathcal{T}_C \rightarrow \mathcal{U}_{SN}$ of σ_C is a system net $en(\sigma_C) = (P, T, F, l, M_{init}, M_{final})$ such that:*

- $P = \{p_i \mid 1 \leq i \leq |\sigma_C| + 1\}$,
- $T = \{t_i \mid 1 \leq i \leq |\sigma_C|\}$,
- $F = \bigcup_{1 \leq i \leq |\sigma_C|} \{(p_i, t_i), (t_i, p_{i+1})\}$
- $l: T \rightarrow \mathcal{U}_A$ such that for all $1 \leq i \leq |\sigma_C|$, $l(t_i) = \sigma_C[i]$,
- $M_{init} = \{p_1\}$,
- $M_{final} = \{p_{|P|}\}$.

Definition 16 (Product of two Petri Nets). *Let $S_1 = (P_1, T_1, F_1, l_1, M_{init_1}, M_{final_1})$ and $S_2 = (P_2, T_2, F_2, l_2, M_{init_2}, M_{final_2})$ be two system nets. The product net of S_1 and S_2 is the system net $S = S_1 \otimes S_2 = (P, T, F, l, M_{init}, M_{final})$ such that:*

- $P = P_1 \cup P_2$,
- $T \subseteq (T_1 \cup \{\gg\}) \times T_2 \cup \{\gg\}$ such that $T = \{(t_1, \gg) \mid t_1 \in T_1\} \cup \{(\gg, t_2) \mid t_2 \in T_2\} \cup \{(t_1, t_2) \in (T_1 \times T_2) \mid l_1(t_1) = l_2(t_2) \neq \tau\}$,
- $F \subseteq (P \times T) \cup (T \times P)$ such that
 - $F = \{(p, (t, \gg)) \mid p \in P_1 \wedge t \in T_1 \wedge (p, t) \in F_1\} \cup$
 - $\{(t, \gg), p \mid t \in T_1 \wedge p \in P_1 \wedge (t, p) \in F_1\} \cup$
 - $\{(p, (t, \gg)) \mid p \in P_2 \wedge t \in T_2 \wedge (p, t) \in F_2\} \cup$

- $$\begin{aligned} & \{((t, \gg), p) \mid t \in T_2 \wedge p \in P_2 \wedge (t, p) \in F_2\} \cup \\ & \{(p, (t_1, t_2)) \mid p \in P_1 \cup P_2 \wedge (t_1, t_2) \in T \cap (T_1 \times T_2)\} \cup \\ & \{((t_1, t_2), p) \mid p \in P_1 \cup P_2 \wedge (t_1, t_2) \in T \cap (T_1 \times T_2)\} \\ - & l: T \rightarrow \mathcal{U}_A \text{ such that for all } (t_1, t_2) \in T, l((t_1, t_2)) = l_1(t_1) \text{ if } t_2 = \gg, \\ & l((t_1, t_2)) = l_2(t_2) \text{ if } t_1 = \gg, \text{ and } l((t_1, t_2)) = l_1(t_1) \text{ otherwise,} \\ - & M_{init} = M_{init_1} \uplus M_{init_2}, \\ - & M_{final} = M_{final_1} \uplus M_{final_2}. \end{aligned}$$

We now present a technique which improves the performance of calculating the lower bound for conformance cost over using a bruteforce method. We will produce a version of the event net that embeds the possible behaviors of the uncertain trace. We define a *behavior net*, a Petri net that can replay all and only the realizations of an uncertain trace.

In order to obtain such a Petri net we first built a directed graph representing the uncertain trace as an intermediate step. We also need to present the concept of *transitive reduction*: given a directed graph G , its transitive reduction G' is a graph with the same set of vertices, the same reachability between vertices, and a minimal number of arcs, such that every pair of vertices is connected by at most one path. The transitive reduction of a directed acyclic graph always exists and is unique [4].

We can then define the *behavior graph*, which contains a vertex for each uncertain event in the trace and contains an edge between two vertices if the corresponding uncertain events may happen one directly after the other.

Definition 17 (Behavior Graph). Let $\sigma_U \in \mathcal{T}_U$ be a simple uncertain trace. A behavior graph $bg: \mathcal{T}_U \rightarrow \mathcal{U}_G$ is the transitive reduction of a directed graph (V, E) , where V is the set of vertices and $E \subseteq V \times V$ is the set of directed edges, such that:

- $V = \sigma_U$
- $E = \{(v_1, v_2) \mid v_1 \in V, v_2 \in V, \pi_{t_{max}}^{L_U^S}(v_1) < \pi_{t_{min}}^{L_U^S}(v_2)\}$

The behavior graph provides a structured representation of the uncertainty on the timestamp: when a specific vertex has two or more outbound edges, the events corresponding to the destination vertices can occur in any order, concurrently with each other. The property of \mathcal{U}_T of being totally ordered and the property $t_{min} < t_{max}$ of all simple uncertain traces ensure that the behavior graph is acyclic both before and after the transitive reduction. We can see the result on the example trace in Figures 3 and 4.

We then obtain a *behavior net* by replacing every vertex in the behavior graph with one or more transitions in a XOR configuration, each representing an activity contained in the π_A set of the corresponding uncertain event. The edges of the behavior graph become connection through places in the behavior net.

Definition 18 (Behavior Net). Let $\sigma_U \in \mathcal{T}_U$ be a simple uncertain trace, and let $bg(\sigma_U) = (V, E)$ be the corresponding behavior graph. A behavior net $bn: \mathcal{T}_U \rightarrow \mathcal{U}_{SN}$ is a system net $bn(\sigma_U) = (P, T, F, l, M_{init}, M_{final})$ such that:

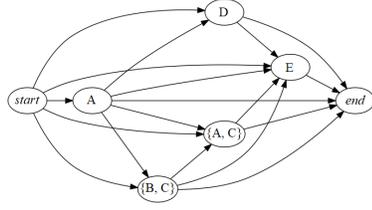


Fig. 3. The behavior graph of the trace in Table 3 before applying the transitive reduction.

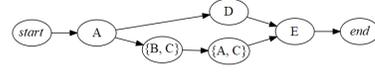


Fig. 4. The same behavior graph after the transitive reduction.

- $T = \{(v, a) \mid v \in V \wedge a \in \pi_a^{L_U^S}(v)\} \cup \{(v, \tau) \mid v \in V \wedge a \in \pi_a^{L_U^S}(v) \wedge \pi_o^{L_U^S}(v) = ?\}$
- $P = E \cup \{start, end\}$
- $F = \{((v_1, a), (v_1, v_2)) \mid (v_1, a) \in T, (v_1, v_2) \in E\} \cup$
 $\{((v_1, v_2), (v_2, a)) \mid (v_1, v_2) \in P, (v_2, a) \in E\} \cup$
 $\{(start, (v, a)) \mid (v, a) \in T \wedge \forall v_* \in V \pi_{t_{min}}^{L_U^S}(v) < \pi_{t_{min}}^{L_U^S}(v_*)\} \cup$
 $\{((v, a), end) \mid (v, a) \in T \wedge \forall v_* \in V \pi_{t_{max}}^{L_U^S}(v) > \pi_{t_{max}}^{L_U^S}(v_*)\}$
- $l = \{((v, a), a) \mid (v, a) \in T \wedge a \neq \tau\}$
- $M_{init} = \{start\}$
- $M_{final} = \{end\}$

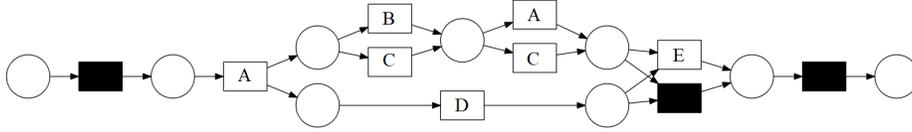


Fig. 5. The behavior net corresponding to the uncertain trace in Table 3.

In Figure 5 we can see the behavior net corresponding to the uncertain trace in Table 2. It is important to notice that every set of edges in the behavior graph with the same source vertex generate an AND split in the behavior net, and a set of edges with the same destination vertex generate an AND join. At the same time, the transitions which labels correspond to different possible activities in an uncertain event will appear in a XOR construct inside the behavior net.

This means that every set of events which timestamps allow for overlapping will be represented in the behavior net by transitions inside an AND construct, and will then allow to execute in the net all the possible sequences of events obtained choosing a possible value for the uncertain timestamp attribute. In the same fashion, an event with uncertainty on the activity will be represented by a number of transitions in a XOR construct, that allows to replay any possible choice for the activity attribute. It follows that, by construction, for a certain simple uncertain trace σ_U we have that $\phi(bn(\sigma_U)) = \mathcal{R}(\sigma_U)$.

We can use the behavior net of an uncertain trace σ_U in lieu of the event net to compute alignments with a model $SN \in \mathcal{U}_{SN}$; the search algorithm returns an optimal alignment, a sequence of moves $(x, (y, t))$ with $x \in \mathcal{U}_X$, $y \in \mathcal{U}_A$ and t transition of the model SN . After removing all “ \gg ” symbols, the sequence of first elements of the moves will describe a complete firing sequence σ_X of the behavior net. Since σ_X is complete, $\sigma_X \in \phi(bn(\sigma_U))$ and, thus, $\sigma_X \in \mathcal{R}(\sigma_U)$. It follows that σ_X is a realization of σ_U , and the search algorithm ensures that σ_X is a realization with optimal conformance cost for the model SN : $\delta(\lambda_{SN}(\sigma_X)) = \min_{\sigma_C \in \mathcal{R}(\sigma_U)} \lambda_{SN}(\sigma_C) = \delta_{min}(\sigma_U)$.

6 Experiments

The technique to compute conformance for strongly uncertain traces and to create the behavior net hereby described has been implemented for testing, using the code for alignments already provided in the process mining Python library PM4Py [1]. Uncertainty has been represented in the XES standard through meta-attributes and constructs such as lists, such that any XES importer can read an uncertain log file. The algorithm was designed to be fully compatible with non-uncertain XES event logs; the meta-attributes for uncertainty were designed to be partly compatible with other process mining algorithms – meta-attributes describing the possible values for an uncertain activity or the interval of an uncertain timestamp can also specify a “fallback value” that other process mining software will read as (certain) activity or timestamp value.

Two experiments were run: the first to inspect the bounds for conformance score as increasingly more uncertainty is added to an event log; the other test assesses the difference in performance between the bruteforce method and the behavior net. We ran the tests on synthetic event logs where we added uncertainty. This way we can control the amounts of uncertainty in event data. Through the ProM plugin “*Generate block-structured stochastic Petri nets*” we generated Petri nets of different sizes in terms of number of transitions n ; then, we used PM4Py in order to generate event logs adding uncertainty to attributes. Activities and timestamps are uncertain with probability p ; also, events have a chance to be indeterminate with probability p .

The pipeline for the first experiment was the following: we generated a model with $n = 10$; we executed the model to obtain 250 traces; we added deviations to events (every event has 20% chances to have the wrong activity, every pair of consecutive events have 20% chances of having the timestamps swapped, every trace has 40% chances to have an additional event). We then added uncertainty to the events: each event has probability p of having two possible values for the activity, probability p of having an uncertain timestamp, and probability p to be an indeterminate event. We calculated the bounds on conformance cost of the log so generated and repeated the procedure for increasing values of p .

Figure 6 shows the results. We can see that the cost shows the expected behavior: at $p = 0$ the two bounds coincide, since the traces are certain and have only one realization. Conversely, the log with $p = 0.6$ has a total of 1629 devia-

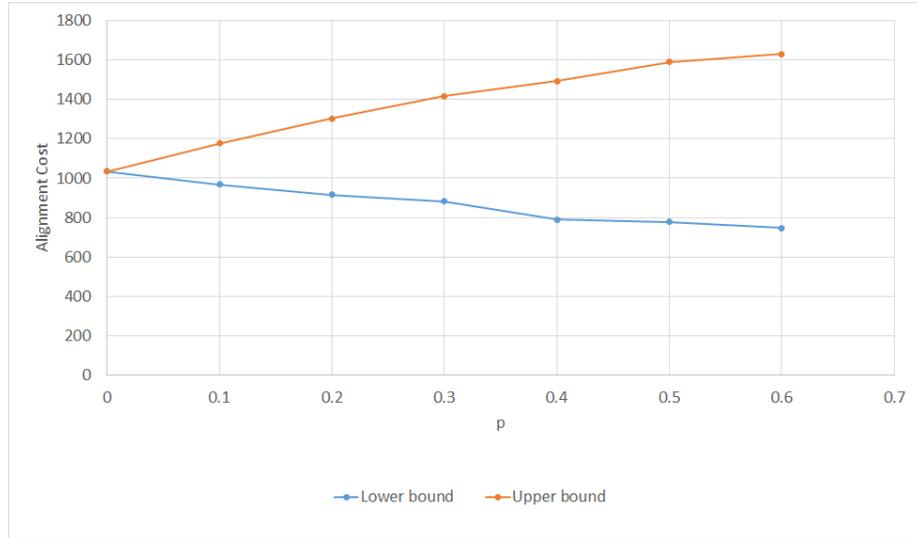


Fig. 6. The change in lower and upper bound for conformance checking of an event log with increasing probability of having uncertainty on event data.

tions on the worst case scenario (6.52 on average per trace), and 747 deviations in the best case scenario (2.99 on average per trace); a process that includes traces with comparable uncertainty has thus a huge difference in behavior between the best case scenario and worst case scenario. The evaluation of the best and worst case scenario for uncertain traces can give useful indications to a business user on the parts of the process where there is a high risk of deviation, in order to enhance them.

The second experiment concerns the performance of calculating the lower bound of the cost via the behavior net versus the bruteforce method of separately listing all the realizations of an uncertain trace, evaluating all of them through alignments, then picking the best value. We used a constant value of $p = 0.2$ and logs of 100 traces for this test, with progressively increasing values of n .

Figure 7 summarizes the results. As the diagram shows, the difference in time between the two methods tends to diverge quickly even on a logarithmic scale. With $n = 5$, the behavior net provides the lower bound in 40.4% of the time required by the bruteforce method. For $n = 20$, the largest model we could test, the behavior net takes 0.04% of the time needed by the bruteforce method. This shows a very large improvement in the computing time for the lower bound computation, so the best case scenario for the conformance cost of an uncertain trace can be obtained efficiently thanks to the structural properties of the behavior net.

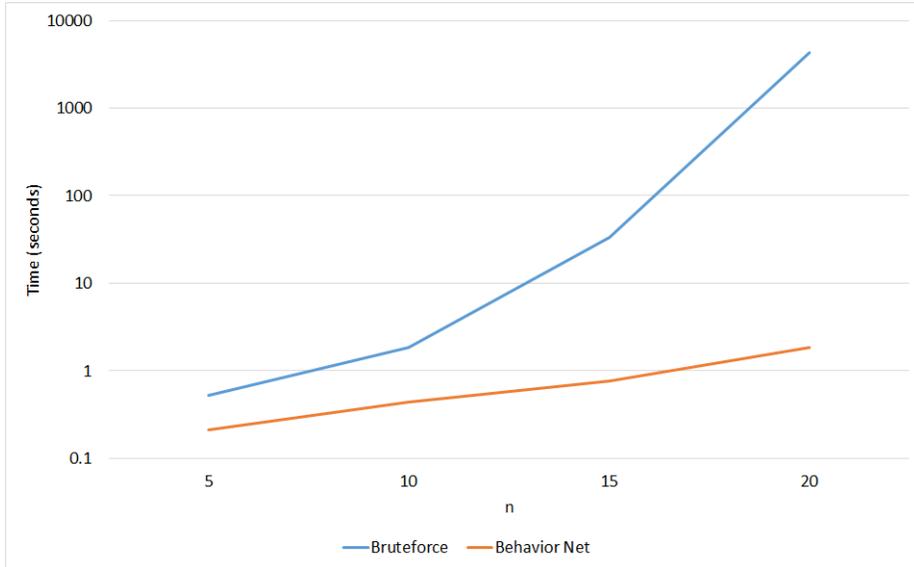


Fig. 7. Effect on time performance of calculating the lower bound for conformance cost with the bruteforce method vs. the behavior net.

7 Conclusion

As the need of quickly and effectively analyze process data has arisen in the recent past and is growing to this day, many new types of information regarding events are recorded; this calls for new techniques able to provide an adequate interpretation of the new data. In this paper we presented a new paradigm for process mining applied to event data: explicit uncertainty. We described the possible form it can assume, building a taxonomy of different types of uncertainty. We then designed a formal mathematical infrastructure to define the various flavors of uncertainty shown in the taxonomy. Then, in order to assess the practical applications of the uncertainty framework, we applied it to a well consolidated technique for conformance checking: aligning data to a reference Petri net. The results can provide insights on the possible violations of process instances recorded with uncertainty against a normative model. The behavior net provides an efficient way to compute the lower bound for the conformance cost – i.e. the best case scenario for conformity of uncertain process data – with a large improvement on time performance with respect to a bruteforce procedure.

The approaches shown here can be extended in a number of ways. An important step in this line of research is assessing the technique on real-life logs. From a performance perspective, to improve the usability of alignments over uncertainty we shall optimize the computation of the upper bound of the conformance cost. Another natural continuation of this work is extending the conformance checking technique to logs with weak uncertainty. Many possibilities can be pursued to

broaden the concept of uncertainty on different process mining methods: for example, discovering a Petri net from an uncertain event log, or develop techniques to mine Petri nets that embed uncertainty information about the process.

References

1. Pm4py, <http://pm4py.pads.rwth-aachen.de/>
2. Van der Aalst, W., Adriansyah, A., De Medeiros, A.K.A., Arcieri, F., Baier, T., Blickle, T., Bose, J.C., Van Den Brand, P., Brandtjen, R., Buijs, J., et al.: Process mining manifesto. In: International Conference on Business Process Management. pp. 169–194. Springer (2011)
3. Adriansyah, A.: Aligning observed and modeled behavior (2014)
4. Aho, A.V., Garey, M.R., Ullman, J.D.: The transitive reduction of a directed graph. *SIAM Journal on Computing* **1**(2), 131–137 (1972)
5. Carmona, J., van Dongen, B., Solti, A., Weidlich, M.: *Conformance Checking: Relating Processes and Models*. Springer (2018)
6. Conforti, R., La Rosa, M., ter Hofstede, A.: Timestamp repair for business process event logs (2018), <http://hdl.handle.net/11343/209011>, [preprint]
7. Conforti, R., La Rosa, M., ter Hofstede, A.H.: Filtering out infrequent behavior from business process event logs. *IEEE Transactions on Knowledge and Data Engineering* **29**(2), 300–314 (2017)
8. Han, J., Pei, J., Kamber, M.: *Data mining: concepts and techniques*. Elsevier (2011)
9. Sani, M.F., van Zelst, S.J., van der Aalst, W.M.: Improving process discovery results by filtering outliers using conditional behavioural probabilities. In: International Conference on Business Process Management. pp. 216–229. Springer (2017)
10. Sani, M.F., van Zelst, S.J., van der Aalst, W.M.: Repairing outlier behaviour in event logs. In: International Conference on Business Information Systems. pp. 115–131. Springer (2018)
11. Suriadi, S., Andrews, R., ter Hofstede, A.H., Wynn, M.T.: Event log imperfection patterns for process mining: Towards a systematic approach to cleaning event logs. *Information Systems* **64**, 132–150 (2017)
12. Wang, J., Song, S., Lin, X., Zhu, X., Pei, J.: Cleaning structured event logs: A graph repair approach. In: *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*. pp. 30–41. IEEE (2015)
13. van Zelst, S.J., Sani, M.F., Ostovar, A., Conforti, R., La Rosa, M.: Filtering spurious events from event streams of business processes. In: International Conference on Advanced Information Systems Engineering. pp. 35–52. Springer (2018)