

Abducing Compliance of Incomplete Event Logs

Federico Chesani¹, Riccardo De Masellis², Chiara Di Francescomarino²,
Chiara Ghidini²(✉), Paola Mello¹, Marco Montali³, and Sergio Tessaris³

¹ University of Bologna, Bologna, Italy

{federico.chesani,paola.mello}@unibo.it

² FBK-IRST, Via Sommarive 18, 38050 Trento, Italy

{r.demasellis,dfmchiara,ghidini}@fbk.eu

³ Free University of Bozen–Bolzano, Piazza Università, 1, 39100 Bozen–Bolzano, Italy

{montali,tessaris}@inf.unibz.it

Abstract. The capability to store data about business processes execution in so-called Event Logs has brought to the diffusion of tools for the analysis of process executions and for the assessment of the *goodness* of a process model. Nonetheless, these tools are often very rigid in dealing with Event Logs that include incomplete information about the process execution. Thus, while the ability of handling incomplete event data is one of the challenges mentioned in the process mining manifesto, the evaluation of compliance of an execution trace still requires an end-to-end complete trace to be performed. This paper exploits the power of abduction to provide a flexible, yet computationally effective, framework to deal with different forms of incompleteness in an Event Log. Moreover it proposes a refinement of the classical notion of compliance into *strong* and *conditional* compliance to take into account incomplete logs.

Keywords: Abductive logic programming · Formal verification · Compliance in business process · Incompleteness in business processes

1 Introduction

The proliferation of IT systems able to store process executions traces in so-called event logs has originated a quest towards tools that offer the possibility of discovering, checking the conformance and enhancing process models based on actual behaviors [1]. Focusing on conformance, that is, on a scenario where the aim is to assess how a *prescriptive* (or “de jure”) process model relates to the execution traces, a fundamental notion is the one of *trace compliance*. Compliance results can be used by business analysts to assess the goodness of a process model and understand how it relates to the actual behaviours exhibited by a company, consequently providing the basis for process re-design, governance and improvement.

The use of event logs to evaluate the goodness of a process model becomes hard and potentially misleading when the event log contains only partial information on the process execution. Thus, while the presence of non-monitorable activities (or errors in the logging procedure) makes the ability of handling incomplete

event data one of the main challenges of the BP community, as mentioned in the process mining manifesto [1], still trace compliance of an execution trace requires the presence of a complete end-to-end execution trace to be evaluated. Notable exceptions are [2,3] where trace incompleteness is managed in an algorithmic/heuristic manner using log repair techniques.

In this paper, we take an orthogonal approach and thoroughly address the problem of log incompleteness from a logic-based point of view, adopting an approach based on *abduction* [4]. Differently from techniques that focus on algorithmic/heuristic repairs of an incomplete trace, we are interested in characterising the notion of incomplete log compliance by means of a sound and complete inference procedure. We rely on abduction to combine the partial knowledge about the real executions of a process as reflected by a (potentially) incomplete event log, with the background knowledge captured in a process model. In particular, abductive reasoning handles different forms of missing information by formulating *hypotheses* that explain how the event log may be “completed” with the missing information, so as to reconcile it with the process model. This leads us to refine the classical notion of *conformance-by-alignment* [5] between an execution trace and a process model into **strong** and **conditional** compliance, to account for incompleteness. In detail, the paper provides: (i) a classification of different forms of incompleteness of an event log based on three dimensions: log incompleteness, trace incompleteness, and event description incompleteness (Sect. 2.1); (ii) a reformulation of the notion of compliance into strong and conditional compliance (Sect. 2.2); and (iii) an encoding of structured process models¹ and of event logs in the SCIFF abductive logic framework [8], and a usage of the SCIFF proof procedure to compute strong, conditional and non-compliance with incomplete event log (Sect. 3). The ideas are illustrated by means of a simple example, and related work is contained in Sect. 4.

2 Dealing with Incomplete Event Logs

We aim at solving the problem of the post mortem identification of compliant traces in the presence of incomplete event logs, given the *prescriptive* knowledge contained in a process model. To do this, we first investigate what incomplete event logs are (Sect. 2.1) and then how we can adapt the notion of compliance to deal with incomplete logs (Sect. 2.2). We perform this investigation with the help of a simple example, described using the BPMN (Business Process Modeling Notation) language².

Example 1 (Obtaining a Permit of Stay in Italy). Consider the BPMN process in Fig. 1, hereafter called the Permit-Of-Stay (POS) process, which takes inspiration from the procedure for the granting of a permit of stay in Italy.

Upon her arrival in Italy (AI), the person in need of a permit of stay has three different alternatives: if she is from a EU country and remains in Italy for at

¹ We follow previous work in the area of BPM and focus on structured process models and on models with no repeating tasks, in the spirit of [6,7], respectively.

² For the sake of clarity we use BPMN, but our framework is language-independent.




Fig. 1. A process for obtaining a permit of stay (in Italy).

most 30 days, then only indicating her presence in Italy (DP) is needed; if she is from the EU and must remain in Italy for more than 30 days, then she needs to get an identity certificate (GIC) and present it (PIC). In all the remaining cases she needs to fill a documentation (FD) which is then checked (CD). When the documentation is correct, it is presented (PD) and a certificate is received (RC). The procedure concludes with the provision of the permit of stay (SI).

2.1 Classifying Process Execution (In)Completeness

We assume that each execution of the POS process in Fig. 1 is (partially) monitored and logged by an information system. We also assume that activities are atomic, i.e., executing an activity results in an event associated to a single timestamp: event (A, t) indicates that activity A has been executed at time t . A sample trace³ that logs the execution of a POS instance is:

$$\{(AI, t_1), (FD, t_2), (CD, t_3), (PD, t_4), (RC, t_5), (SI, t_6)\} \quad (1)$$

where $t_i > t_j$ for $i, j \in \{1, \dots, 6\}$ such that $i > j$. This trace corresponds to the execution of the lower branch of the POC process, where the loop is never executed. A set of execution traces of the same process form an event log.

In many real cases, a number of difficulties may arise when exploiting the data contained in an information system in order to build an event log. Thus, instead of the extremely informative trace reported in (1), we may obtain something like:

$$\{(FD, -), (-, t_2), (SI, t_6), (-, -)\} \quad (2)$$

This trace does not completely describe an execution of the POS process. For example, the first event logged in the trace is FD. However, by looking at the process description, it is easy to see that the first event of every execution has to be AI. By assuming that the process executors indeed followed the prescriptions of the model, this suggests that the AI-related event has not been logged. Moreover, certain events have been only partially observed. For example, the FD-related event is incomplete, because its exact timestamp is unknown. In this paper, we use “-” to denote a missing information unit.

In accordance with the IEEE standard XES format for representing event logs [9], in general we can describe an event log as a set of execution traces. Each trace,

³ We often present the events in a trace ordered according to their execution time. This is only to enhance readability since the position of an event is fully determined by its timestamp, or unknown if the timestamp is missing.

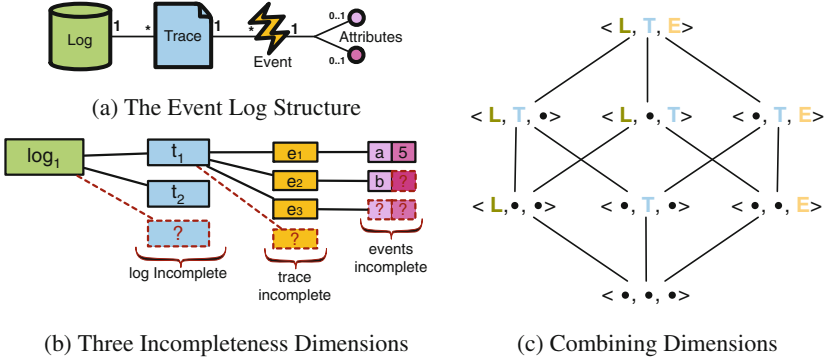


Fig. 2. Classifying (in)completeness

in turn, contains events, which are described by means of n -tuples, where each element of the tuple is the value of a given attribute (see Fig. 2a, where we restrict to two attributes as we do in the paper). Consequently, we can classify incompleteness along these three dimensions: incompleteness of the log, incompleteness of the trace, and incompleteness of the event description (see Fig. 2b).

(In)Completeness of the Log. Within this dimension we analyse whether the log contains at least one instance for each possible execution that is allowed in the model. Note that one can account for this form of (in)completeness only by: (a) limiting the analysis to the control flow, without considering complex data objects that may contain values from an unbounded domain; and (b) assuming that there is a maximum length for all traces, thus limiting the overall number of traces that may originate from the unbounded execution of loops. An example of complete log for the POS process is:

$$L_1 = \left\{ \left\{ (AI, t_{a1}), (DP, t_{a2}), (SI, t_{a3}) \right\}, \left\{ (AI, t_{b1}), (GIC, t_{b2}), (PIC, t_{b3}), (SI, t_{b4}) \right\}, \right. \\ \left. \left\{ (AI, t_{c1}), (FD, t_{c2}), (CD, t_{c3}), (PD, t_{c4}), (RC, t_{c5}), (SI, t_{c6}) \right\} \right\} \quad (3)$$

where we assume that each trace cannot contain more than 6 event, which intuitively means that the loop is never executed twice.

Assuming this form of completeness is not strictly required to have good process models, and could be unrealistic in practice. In fact, even under the assumption of a maximum trace length, the number of allowed traces could become extremely huge due to (bounded) loops, and the (conditional) interleavings generated by parallel blocks and or choices. Still, analysing the (in)completeness of an event log may be useful to discover parts of the control flow that never occur in practice.

(In)Completeness of the Trace. Within this dimension we focus on a single trace, examining whether it contains a sequence of events that corresponds to an execution foreseen by the process model from start to end. Trace (1) is an example of complete trace. An example of incomplete trace is:

$$\{(AI, t_1), (PIC, t_2)(SI, t_3)\} \quad (4)$$

This trace should also contain an event of the form (GIC, t) , s.t. $t_1 < t < t_2$.

(In)Completeness of the Event Description. Within this dimension we focus on the completeness of a single event. Events are usually described as complex objects containing data about the executed activity, its time stamp, and so on [9]. These data can be missing or corrupted. As pointed out before, we consider activity names and timestamps. Thus, incompleteness in the event description may concern the activity name, its timestamp, or both. This is reflected in trace (2): (i) event $(\text{FD}, -)$ indicates that activity FD has been executed, but at an unknown time; (ii) $(-, t_2)$ witnesses that an activity has been executed at time t_2 , but we do not know which one; (iii) $(-, -)$ attests that the trace contains some event, whose activity and time are unknown.

We can characterise the (in)completeness of an event log in terms of (any) combination of these three basic forms. At one extreme, we may encounter a log, such as (3), that is complete along all three dimensions. At the other extreme, we may have a log such as:

$$L_2 = \{ \{(\text{AI}, -), (-, t_{a2})\}, \{(\text{AI}, t_{b1}), (-, -), (-, t_{b2}), (\text{SI}, t_{b3})\} \} \quad (5)$$

that is incomplete along all the dimensions. Intermediate situations may arise as well, as graphically depicted in the lattice of Fig. 2c, where $\langle L, T, E \rangle$ indicates the top value (completeness for all three dimensions) and $\langle \bullet, \bullet, \bullet \rangle$ indicated the bottom value (incompleteness of all three dimensions).

2.2 Refining the Notion of Compliance

In our work we consider *prescriptive* process models, that is, models that describe the only acceptable executions. These correspond to the so-called “de jure” models in [5], and consequently call for a definition of *compliance*, so as to characterise the degree to which a given trace conforms/is aligned to the model. The notion of compliance typically requires that the trace represents an end-to-end, valid execution that can be fully replayed on the process model. We call this notion of compliance **strong** compliance. Trace (1) is an example of trace that is fully compliant to the POS process.

Strong compliance is too restrictive when the trace is possibly incomplete. In fact, incompleteness hinders the possibility of replaying it on the process model. However, conformance might be regained by assuming that the trace included additional information on the missing part; in this case we say that the trace is **conditionally** compliant, to reflect that compliance conditionally depends on how the partial information contained in the trace is complemented with the missing one. Consider the partial trace:

$$\{(\text{AI}, t_1), (\text{GIC}, -), (\text{SI}, t_3)\} \quad (6)$$

It is easy to see that (6) is compliant with POS, *if* we assume that

$$\text{GIC was executed at a time } t_i \text{ s.t. } t_1 < t_i < t_3 \quad (7)$$

$$\text{an execution of PIC was performed at a time } t_j \text{ s.t. } t_i < t_j < t_3 \quad (8)$$

Note that the set of assumptions needed to reconstruct full conformance is not necessarily unique. This reflects that alternative strongly compliant real process executions might have led to the recorded partial trace. On the other hand, there are cases in which no assumptions can lead to full conformance. In this case, the partial trace is considered **non-compliant**. For example, the following trace does not comply with POS, since it records that GIC and CD have been both executed, although they belong to mutually exclusive branches in the model:

$$\{(AI, t_1), (GIC, -), (CD, t_2), (SI, t_3)\} \quad (9)$$

3 Abduction and Incomplete Logs

Since the aim of this paper is to provide automatic procedures that identify compliant traces in the presence of incomplete event logs, given the *prescriptive* knowledge contained in a process model, we can schematise the input to our problem in three parts: (i) an instance-independent component, the process model, which in this paper is described using BPMN; (ii) an instance-specific component, that is, the (partial) log, and (iii) meta-information attached to the activities in the process model, indicating which are actually *always*, *never* or *possibly* observable (that is, logged) in the event log. The third component is an extension of a typical business process specification that we propose (following and extending the approach described in [10]) to provide *prescriptive* information about the (non-) observability of activities. Thus, for instance, a business analyst will have the possibility to specify that a certain manual activity is never observable while a certain interaction with a web site is always (or possibly) observable. This information can then be used to compute the compliance of a partial trace. In fact the presence of never observable activities will trigger the need to make hypothesis on their execution (as they will never be logged in the event log), while the presence of always observable activities will trigger the need to find their corresponding event in the execution trace (to retain compliance). Note that this extension is not invasive w.r.t. current approaches to business process modelling, as we can always assume that a model where no information on observability is provided is entirely possibly observable.

Given the input of our problem, in Sect. 3.1 we provide an overview on abduction and on how the SCIFF framework represents *always*, *never* or *possibly* observable activities; in Sect. 3.2 we show how to use SCIFF to encode a process model and a partial log; in Sect. 3.3 we show how we can formalize the informal different forms of compliance presented in Sect. 2.2; finally, in Sect. 3.4 we illustrate how SCIFF can be used to solve the different forms of incompleteness identified in Sect. 2.1.

3.1 The SCIFF in Short

Abduction is a non-monotonic reasoning process where hypotheses are made to explain observed facts [11]. While deductive reasoning focuses on deciding if a formula ϕ logically follows from a set Γ of logical assertions known to hold,

in abductive reasoning it is assumed that ϕ holds (as it corresponds to a set of observed facts) but it cannot be directly inferred by Γ . To make ϕ a consequence of Γ , abduction looks for a further set Δ of hypothesis, taken from a given set of abducibles \mathcal{A} , which completes Γ in such a way that ϕ can be inferred (in symbols $\Gamma \cup \Delta \models \phi$). The set Δ is called *abductive explanation* (of ϕ). In addition, Δ must usually satisfy a set of (domain-dependent) integrity constraints \mathcal{IC} (in symbols, $\Gamma \cup \Delta \models \mathcal{IC}$). A typical integrity constraint (IC) is a *denial*, which expresses that two explanations are mutually exclusive.

Abduction has been introduced in Logic Programming in [4]. There, an *Abductive Logic Program (ALP)* is defined as a triple $\langle \Gamma, \mathcal{A}, \mathcal{IC} \rangle$, where: (i) Γ is a logic program, (ii) \mathcal{A} is a set of abducible predicates, and (iii) \mathcal{IC} a set of ICs. Given a goal ϕ , abductive reasoning looks for a set of positive, atomic literals $\Delta \subseteq \mathcal{A}^4$ such that they entail $\phi \cup \mathcal{IC}$.

In this paper we leverage on the SCIFF abductive logic programming framework [8], an extension of the IFF abductive proof procedure [12], and on its efficient implementation in CHR [13]. Beside the general notion of abducible, the SCIFF framework has been enriched with the notions of *happened event*, *expectation*, and *compliance* of an observed execution with a set of expectations. This makes SCIFF suitable for dealing with event log incompleteness. Let a be an event corresponding to the execution of a process activity, and T (possibly with subscripts) its execution time⁵. Abducibles are used here to make hypotheses on events that are not recorded in the examined trace. They are denoted using $\mathbf{ABD}(a, T)$. Happened events are non-abducible, and account for events that have been logged in the trace. They are denoted with $\mathbf{H}(a, T)$. Expectations $\mathbf{E}(a, T)$, instead, model events that should occur (and therefore should be present in a trace). Compliance is described in Sect. 3.3.

ICs in SCIFF are used to relate happened events/abduced predicates with expectations/predicates to be abduced. Specifically, an IC is a rule of the form *body* \rightarrow *head*, where *body* contains a conjunction of happened events, general abducibles, and defined predicates, while *head* contains a disjunction of conjunctions of expectations, general abducibles, and defined predicates. Usually, variables appearing in the body are quantified universally, while variables appearing in the head are quantified existentially.

3.2 Encoding Structured Processes and Their Executions in SCIFF

Let us illustrate how to encode all the different components of an (incomplete) event log and a structured process model one by one.

Event Log. A log is a set of traces, each constituted by a set of observed (atomic) events. Thus trace (4) is represented in SCIFF as $\{\mathbf{H}(A1, t_1), \mathbf{H}(PIC, t_2), \mathbf{H}(S1, t_3)\}$.

⁴ We slightly abuse the notation of \subseteq , meaning that every positive atomic literal in Δ is the instance of a predicate in \mathcal{A} .

⁵ In the remainder of this paper we will assume that the time domain relies on natural numbers.




Fig. 3. Encoding always/never observable activities.

Always/never Observable Activities. Coherently with the representation of an execution trace, the logging of the execution of an observable activity is represented in SCIFF using an happened event, whereas the hypothesis on the execution of a never observable activity is represented using an abducible **ABD** (see Fig. 3a). Given an event **a** occurring at T , we use a function τ that represents the execution of **a** as:

$$\tau(\mathbf{a}, T) = \begin{cases} \mathbf{H}(\mathbf{a}, T) & \text{if } \mathbf{a} \text{ is observable} \\ \mathbf{ABD}(\mathbf{a}, T) & \text{if } \mathbf{a} \text{ is never observable} \end{cases}$$

As for expected occurrences, the encoding again depends on the observability of the activity: if the activity is observable, then its expected occurrence is mapped to a SCIFF expectation; otherwise, it is hypothesised using the aforementioned abducible **ABD** (see Fig. 3b). To this end we use a function ε that maps the expecting of the execution of **a** at time T as follows:

$$\varepsilon(\mathbf{a}, T) = \begin{cases} \mathbf{E}(\mathbf{a}, T) & \text{if } \mathbf{a} \text{ is observable} \\ \mathbf{ABD}(\mathbf{a}, T) & \text{if } \mathbf{a} \text{ is never observable} \end{cases}$$

Structured Process Model Constructs. A process model is encoded in SCIFF by generating ICs that relate the execution of an activity to the future, expected executions of further activities. In practice, each process model construct is transformed into a corresponding IC. We handle, case-by-case, all the single-entry single-exit block types of structured process models.

Sequence. Two activities **a** and **b** are in sequence if, whenever the first is executed, the second is expected to be executed at a later time:

$$\tau(\mathbf{a}, T_a) \rightarrow \varepsilon(\mathbf{b}, T_b) \wedge T_b > T_a. \quad (10)$$

Notice that T_a is quantified universally, while T_b is existentially quantified.

And-split activates parallel threads spanning from the same activity. In particular, the fact that activity **a** triggers two parallel threads, one expecting the execution of **b**, and the other that of **c**, is captured using an IC with a conjunctive consequent:

$$\tau(\mathbf{a}, T_a) \rightarrow \varepsilon(\mathbf{b}, T_b) \wedge T_b > T_a \wedge \varepsilon(\mathbf{c}, T_c) \wedge T_c > T_a.$$

And-join mirrors the and-split, synchronizing multiple concurrent execution threads and merging them into a single thread. When activities **a** and **b** are both executed, then activity **c** is expected next, is captured using an IC with a conjunctive antecedent:

$$\tau(\mathbf{a}, T_a) \wedge \tau(\mathbf{b}, T_b) \rightarrow \varepsilon(\mathbf{c}, T_c) \wedge T_c > T_a \wedge T_c > T_b.$$

The encoding of **Xor-split/Xor-join** and **Or-split/Or-join** can be found in [14].

Possibly Observable Activities. A possibly observable activity is managed by considering the disjunctive combination of two cases: one in which it is assumed to be observable, and one in which it is assumed to be never observable. This idea is used to refine ICs used to encode the workflow constructs in the case of partial observability. For instance, if a partially observable activity appears in the antecedent of an IC, two distinct ICs are generated, one where the activity is considered to be observable (**H**), and another in which it is not (**ABD**). Thus in the case of a sequence flow from **a** to **b**, where **a** is possibly observable and **b** is observable, IC (10) generates:

$$\mathbf{H}(\mathbf{a}, T_a) \rightarrow \varepsilon(\mathbf{b}, T_b) \wedge T_b > T_a. \quad \mathbf{ABD}(\mathbf{a}, T_a) \rightarrow \varepsilon(\mathbf{b}, T_b) \wedge T_b > T_a.$$

If multiple partially observable activities would appear in the antecedent of an IC (as, e.g., in the and-join case), then all combinations have to be considered.

Similarly, if a partially observable activity appears in the consequent of an IC, a disjunction must be inserted in the consequent, accounting for the two possibilities of observable/never observable event. If both the antecedent and consequent of an IC would contain a partially observable activity, a combination of the rules above will be used. For example, in the case of a sequence flow from **a** to **b**, where **b** is possibly observable, IC (10) generates:

$$\mathbf{H}(\mathbf{a}, T_a) \rightarrow \mathbf{E}(\mathbf{b}, T_b) \wedge T_b > T_a \vee \mathbf{ABD}(\mathbf{b}, T_b) \wedge T_b > T_a. \quad (11)$$

With this encoding, the SCIFF proof procedure generates firstly an abductive explanation Δ containing an expectation about the execution of **b**. If no **b** is actually observed, Δ is discarded, and a new abductive explanation Δ' is generated containing the hypothesis about **b** (i.e., $\mathbf{ABD}(\mathbf{b}, T_b) \in \Delta'$). Mutual exclusion between these two possibilities is guaranteed by the SCIFF declarative semantics (cf. Definition 3).

Finally, if both the antecedent and consequent of an IC would contain a possibly observable activity, a combination of the rules above will be used.

Start and End of the Process. We introduce two special activities **start** and **end** representing the entry- and exit-point of the process. Two specific ICs are introduced to link these special activities with the process. For example, if the first process activity is **a** (partially observable), the following IC is added:

$$\mathbf{ABD}(\mathbf{start}, 0) \rightarrow \mathbf{E}(\mathbf{a}, T_a) \wedge T_a > 0 \vee \mathbf{ABD}(\mathbf{a}, T_a) \wedge T_a > 0.$$

To ensure the IC triggering, $\mathbf{ABD}(\mathbf{start}, 0)$ is given as goal to the proof procedure.

3.3 Compliance in SCIFF: Declarative Semantics

We are now ready to provide a formal notion of compliance in its different forms. We do so by extending the SCIFF declarative semantics provided in [8] to incorporate log incompleteness (that is, observability features).

A structured process model corresponds to a SCIFF specification $\mathcal{S} = \langle \mathcal{KB}, \mathcal{A}, \mathcal{IC} \rangle$, where: (i) \mathcal{KB} is a Logic Program [15] containing the definition of accessory predicates; (ii) $\mathcal{A} = \{\mathbf{ABD}/2, \mathbf{E}/2\}$, possibly non-ground; (iii) \mathcal{IC} is a set of ICs constructed by following the encoding defined in Sect. 3.2. A (execution) trace and an abductive explanation Δ are defined as follows⁶:

Definition 1. A Trace \mathcal{T} is a set of terms of type $\mathbf{H}(e, T_i)$, where e is a term describing the happened event, and $T_i \in \mathbb{N}$ is the time instant at which the event occurred.

Definition 2. Given a SCIFF specification \mathcal{S} and a trace \mathcal{T} , a set $\Delta \subseteq \mathcal{A}$ is an abductive explanation for $\langle \mathcal{S}, \mathcal{T} \rangle$ if and only if $\text{Comp}(\mathcal{KB} \cup \mathcal{T} \cup \Delta) \cup \text{CET} \cup T_{\mathbb{N}} \models \mathcal{IC}$ where Comp is the (two-valued) completion of a theory [16], CET stands for Clark Equational Theory [17] and $T_{\mathbb{N}}$ is the CLP constraint theory [18] on finite domains.

The following definition fixes the semantics for observable events, and provides the basis for understanding the alignment of a trace with a process model.

Definition 3 (\mathcal{T} -Fulfillment). Given a trace \mathcal{T} , an abducible set Δ is \mathcal{T} -fulfilled if for every event e and for each time t , $\mathbf{E}(e, t) \in \Delta$ if and only if $\mathbf{H}(e, t) \in \mathcal{T}$.

The “only if” direction defines the semantics of expectation, indicating that an expectation is fulfilled when it finds the corresponding happening event in the trace. The “if” direction captures the prescriptive nature of process models, whose *closed* nature requires that only expected events may happen.

Given an abductive explanation Δ , fulfilment acts as a *compliance classifier*, which separates the legal/correct execution traces with respect to Δ from the wrong ones. Expectations however model the strong flavour of compliance: if \mathcal{T} -Fulfillment cannot be achieved because some expectations are not matched by happened events, a \mathbf{ABD} predicate is abduced as specified in the integrity constraints (see for example the IC 11).

Definition 4 (Strong/Conditional Compliance). A trace \mathcal{T} is compliant with a SCIFF specification \mathcal{S} if there exists an abducible set Δ such that: (i) Δ is an abductive explanation for $\langle \mathcal{S}, \mathcal{T} \rangle$, and (ii) Δ is \mathcal{T} -fulfilled. If Δ does not contain any \mathbf{ABD} (besides the special abducibles for start and end), then we say that it is strongly-compliant, otherwise it is conditionally-compliant.

If no abductive explanation that is also \mathcal{T} -fulfilled can be found, then \mathcal{T} is not compliant with the specification of interest. Contrariwise, the abductive explanation witnesses compliance. However, it may contain (possibly non-ground)

⁶ We do not consider the abductive goal, as it is not needed for our treatment.

ABD predicates, abduced due to the incompleteness of \mathcal{T} . In fact, the presence or absence of such predicates determines whether \mathcal{T} is conditionally or strongly compliant. To make an example let us consider traces (6), (1), and (9). In the case of partial trace (6), SCIFF will tell us that it is conditional compliant with the workflow model POS since Δ will contain the formal encoding of the two abducibles (7) and (8) which provide the *abductive explanation* of trace (6). In the case of trace (1), abduction will tell us that it directly follows from Γ without the need of any hypothesis. The case where Δ does not contain any **ABD** coincides in fact, with the classical notion of (deductive) compliance. Finally, in the case of trace (9) SCIFF will tell us that it is not possible to find any set of hypothesis Δ that explains it. This case coincides with the classical notion of (deductive) non-compliance.

We close this section by briefly arguing that our approach is indeed correct. To show correctness, one may proceed in two steps: (i) prove the semantic correctness of the encoding w.r.t. semantics of (conditional/strong) compliance; (ii) prove the correctness of the proof procedure w.r.t. the SCIFF declarative semantics. Step (i) requires to prove that a trace is (conditionally/strong) compliant (in the original execution semantics of the workflow) with a given workflow if and only if the trace is (conditionally/strong) compliant (according to the SCIFF declarative semantics) with the encoding of the workflow in SCIFF. This can be done in the spirit of [19] (where correctness is proven for declarative, constraint-based processes), by arguing that structured processes can be seen as declarative processes that only employ the “chain-response constraint” [19]. For step (ii), we rely on [8], where soundness and completeness of SCIFF w.r.t. its declarative semantics is proved by addressing the case of closed workflow models (the trace is closed and no more events can happen anymore), as well as that of open workflow models (future events can still happen). Our declarative semantics restricts the notions of fulfilment and compliance to a specific current time t_c , i.e., to open traces: hence soundness and completeness still hold.

3.4 Dealing with Process Execution (In)Completeness in SCIFF

We have already illustrated, by means of the POS example, how Definition 4 can be used to address compliance of a partial trace. In this section we illustrate in detail how SCIFF can be used to solve the three dimensions of incompleteness identified in Sect. 2.1.

Trace and event incompleteness are dealt by with SCIFF in a uniform manner. In fact, the **trace/event incompleteness problem** amounts to check if a given log (possibly equipped with incomplete traces/events), is compliant with a prescriptive process model. We consider as input the process model, together with information about the observability of its activities, a trace, and a maximum length for completed traces. The compliance is determined by executing the SCIFF proof procedure and evaluating possible abductive answers. We proceed as follows:

1. We automatically translate the process model with its observability meta-information into a SCIFF specification. If observability information is missing

for some/all the activities, we can safely assume that some/all activities are possibly observable.

2. The SCIFF proof procedure is applied to the SCIFF specification and to the trace under observation, computing *all* the possible abductive answers Δ_i . The maximum trace length information is used to limit the search, as in the unrestricted case the presence of loop may lead to nontermination.
3. If no abductive answer is generated, the trace is deemed as non-compliant. Otherwise, a set of abductive answers $\{\Delta_1, \dots, \Delta_n\}$ has been found. If there exists a Δ_i that does not contain any **ABD** predicate, then the trace is strongly compliant. The trace is conditionally compliant otherwise.

Note that, assessing strong/conditional compliance requires the computation of *all* the abductive answers, thus affecting the performances of the SCIFF proof procedure. If only compliance is needed (without classifying it in strong or conditional), it is possible to compute only the first solution.

A different scenario is provided by the **log incompleteness problem**, which instead focuses on an entire event log, and looks if some possible traces allowed by the model are indeed missing in the log. In this case we consider as input the process model, a maximum length for the completed traces, and a log consisting of a number of different traces; we assume each trace is trace- and event-complete. We proceed as follows:

1. We generate the SCIFF specification from the process model, considering all activities as never observable (i.e., their happening must be always hypothesised, so that we can generate all the possible traces).
2. The SCIFF proof procedure is applied to the SCIFF specification. *All* the possible abductive answers Δ_i are computed, with maximum trace length as specified. Each answer corresponds to a different execution instance allowed by the model. Since all the activities are never observable, the generated Δ_i will contain only **ABD**.
3. For each hypothesised trace in the set $\{\Delta_1, \dots, \Delta_n\}$, a corresponding, distinct trace is looked for in the log. If all the hypothesised traces have a distinct matching observed trace, then the log is deemed as complete.

Notice that, beside the completeness of the log, the proof procedure also generate the missing traces, defined as the Δ_i that do not have a corresponding trace in the log.

An evaluation of the algorithms above, and a study of how different inputs affect their performances is provided in [14] and omitted for lack of space. It shows that the performance of the abductive procedure to evaluate compliance ranges from few seconds when at most a single event description is completely unknown to about 4.5 min when up to 4 event descriptions are missing. A prototype implementation is currently available for download at <http://ai.unibo.it/AlpBPM>.

4 Related Work

The problem of incomplete traces has been tackled by a number of works in the field of process discovery and conformance. Some of them have addressed the

problem of aligning event logs and procedural/declarative process models [2,3]. Such works explore the search space of the set of possible moves to find the best one for aligning the log to the model. Our purpose is not managing generic misalignments between models and logs, but rather focus on a specific type of incompleteness: the model is correct and the log can be incomplete.

We can divide existing works that aim at constructing possible model-compliant “worlds” out of a set of incomplete observations in two groups: quantitative and qualitative approaches. The former rely on the availability of a probabilistic model of execution and knowledge. For example, in [20] the authors exploit stochastic Petri nets and Bayesian Networks to recover missing information. The latter stand on the idea of describing “possible outcomes” regardless of likelihood. For example, in [21] and in [10] the authors exploit Satisfiability Modulo Theory and planning techniques respectively to reconstruct missing information. A different line of work addresses problems of compliance through model checking techniques [22,23]. Here the focus is verifying a broad class of temporal properties rather than specific issues related to incompleteness, which we believe are more naturally represented by abductive techniques.

In this work, the notion of incompleteness has been investigated to take into account its different variants (*log*, *trace* and *event incompleteness*). Similarly, the concept of *observability* has been deeply investigated, by exploring activities *always*, *partially* or *never* observable. This has led to a novel refinement of the notion of compliance.

Abduction and the SCIFF framework have been previously used to model both procedural and declarative processes. In [24], a structured workflow language has been defined, with a formal semantics in SCIFF. In [25], SCIFF has been exploited to formalize and reason about the declarative workflow language Declare.

An interesting work where trace compliance is evaluated through abduction is presented in [26]. Compliance is defined as assessing if actions were executed by users with the right permissions (auditing), and the focus is only on incomplete traces (with complete events). The adopted abductive framework, CIFF [27], only supports ground abducibles, and ICs are limited to denials. The work in [26] explores also the dimension of human confirmation of hypotheses, and proposes a human-based refinement cycle. This is a complementary step with our work, and would be an interesting future direction.

5 Conclusions

We have presented an abductive framework to support business process monitoring (and in particular compliance checking) by attacking the different forms of incompleteness that may be present in an event log. Concerning future development, the SCIFF framework is based on first-order logic, thus paving the way towards (i) the incorporation of data [23], (ii) extensions to further types of workflows (e.g., temporal workflows as in [28]), and (iii) towards the investigation of probabilistic models to deal with incompleteness of knowledge.

References

1. Aalst, W.: Process mining manifesto. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) BPM 2011. LNBP, vol. 99, pp. 169–194. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-28108-2_19](https://doi.org/10.1007/978-3-642-28108-2_19)
2. Adriansyah, A., van Dongen, B.F., van der Aalst, W.M.P.: Conformance checking using cost-based fitness analysis. In: Proceedings of EDOC. IEEE Computer Society (2011)
3. Leoni, M., Maggi, F.M., Aalst, W.M.P.: Aligning event logs and declarative process models for conformance checking. In: Barros, A., Gal, A., Kindler, E. (eds.) BPM 2012. LNCS, vol. 7481, pp. 82–97. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-32885-5_6](https://doi.org/10.1007/978-3-642-32885-5_6)
4. Kakas, A.C., Kowalski, R.A., Toni, F.: Abductive logic programming. *J. Log. Comput.* **2**(6), 719–770 (1992). <http://dblp.uni-trier.de/rec/bibtex/journals/logcom/KakasKT92>
5. van der Aalst, W.M.P.: Process Mining - Discovery, Conformance and Enhancement of Business Processes. Springer, Heidelberg (2011)
6. Kiepuszewski, B., ter Hofstede, A.H.M., Bussler, C.J.: On structured workflow modelling. In: Bubenko, J., Krogstie, J., Pastor, O., Pernici, B., Rolland, C., Sølvsberg, A. (eds.) 25 Years of CAiSE. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-36926-1_19](https://doi.org/10.1007/978-3-642-36926-1_19)
7. van der Aalst, W.M.P., Weijters, T., Maruster, L.: Workflow mining: discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **16**, 1128–1142 (2004)
8. Alberti, M., Chesani, F., Gavanelli, M., Lamma, E., Mello, P., Torroni, P.: Verifiable agent interaction in abductive logic programming: the SCIFF framework. *ACM Trans. Comput. Log.* **9**(4), 29:1–29:43 (2008). <http://dl.acm.org/citation.cfm?doid=1380572.1380578>
9. On process mining, I.T.F.: XES standard definition (2015). <http://www.xes-standard.org/>
10. Di Francescomarino, C., Ghidini, C., Tessaris, S., Sandoval, I.V.: Completing workflow traces using action languages. In: Zdravkovic, J., Kirikova, M., Johannesson, J. (eds.) CAiSE 2015. LNCS, vol. 9097, pp. 314–330. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-19069-3_20](https://doi.org/10.1007/978-3-319-19069-3_20)
11. Kakas, A.C., Mancarella, P.: Abduction and abductive logic programming. In: Proceedings of ICLP (1994)
12. Fung, T.H., Kowalski, R.A.: The iff proof procedure for abductive logic programming. *J. Log. Program.* **33**(2), 151–165 (1997). <http://dblp.uni-trier.de/rec/bibtex/journals/jlp/FungK97>
13. Alberti, M., Gavanelli, M., Lamma, E.: The CHR-based implementation of the SCIFF abductive system. *Fundam. Inform.* **124**, 365–381 (2013)
14. Chesani, F., De Masellis, R., Di Francescomarino, C., Ghidini, C., Mello, P., Montali, M., Tessaris, S.: Abducing compliance of incomplete event logs. Technical report submit/1584687, arXiv (2016)
15. Lloyd, J.W.: Foundations of Logic Programming, 2nd edn. Springer, Heidelberg (1987)
16. Kunen, K.: Negation in logic programming. *J. Log. Program.* **4**(4), 289–308 (1987). <http://dblp.uni-trier.de/rec/bibtex/journals/jlp/Kunen87>
17. Clark, K.L.: Negation as Failure. In: Proceedings of Logic and Data Bases. Plenum Press (1978)

18. Jaffar, J., Maher, M.J., Marriott, K., Stuckey, P.J.: The semantics of constraint logic programs. *J. Log. Program.* **37**(1–3), 1–46 (1998). <http://dblp.uni-trier.de/rec/bibtex/journals/jlp/JaffarMMS98>
19. Montali, M.: Specification and Verification of Declarative Open Interaction Models: A Logic-Based Approach. LNBIP, vol. 56. Springer, Heidelberg (2010)
20. Rogge-Solti, A., Mans, R.S., Aalst, W.M.P., Weske, M.: Improving documentation by repairing event logs. In: Grabis, J., Kirikova, M., Zdravkovic, J., Stirna, J. (eds.) *PoEM 2013*. LNBIP, vol. 165, pp. 129–144. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-41641-5_10](https://doi.org/10.1007/978-3-642-41641-5_10)
21. Bertoli, P., Francescomarino, C., Dragoni, M., Ghidini, C.: Reasoning-based techniques for dealing with incomplete business process execution traces. In: Baldoni, M., Baroglio, C., Boella, G., Micalizio, R. (eds.) *AI*IA 2013*. LNCS (LNAI), vol. 8249, pp. 469–480. Springer, Heidelberg (2013). doi:[10.1007/978-3-319-03524-6_40](https://doi.org/10.1007/978-3-319-03524-6_40)
22. Bagheri Hariri, B., Calvanese, D., De Giacomo, G., Deutsch, A., Montali, M.: Verification of relational data-centric dynamic systems with external services, pp. 163–174. ACM Press (2013)
23. De Masellis, R., Maggi, F.M., Montali, M.: Monitoring data-aware business constraints with finite state automata. In: *Proceedings of ICSSP*. ACM Press (2014)
24. Chesani, F., Mello, P., Montali, M., Storari, S.: Testing careflow process execution conformance by translating a graphical language to computational logic. In: Bellazzi, R., Abu-Hanna, A., Hunter, J. (eds.) *AIME 2007*. LNCS (LNAI), vol. 4594, pp. 479–488. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-73599-1_64](https://doi.org/10.1007/978-3-540-73599-1_64)
25. Montali, M., Pesic, M., van der Aalst, W.M.P., Chesani, F., Mello, P., Storari, S.: Declarative specification and verification of service choreographies. *TWEB* **4**(1), 3:1–3:62 (2010). <http://dl.acm.org/citation.cfm?doid=1658373.1658376>
26. Mian, U.S., den Hartog, J., Etalle, S., Zannone, N.: Auditing with incomplete logs. In: *Proceedings of the 3rd Workshop on Hot Issues in Security Principles and Trust* (2015)
27. Mancarella, P., Terreni, G., Sadri, F., Toni, F., Endriss, U.: The CIFF proof procedure for abductive logic programming with constraints: theory, implementation and experiments. *TPLP* **9**(6), 691 (2009)
28. Kumar, A., Sabbella, S.R., Barton, R.R.: Managing controlled violation of temporal process constraints. In: Motahari-Nezhad, H.R., Recker, J., Weidlich, M. (eds.) *BPM 2015*. LNCS, vol. 9253, pp. 280–296. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-23063-4_20](https://doi.org/10.1007/978-3-319-23063-4_20)