

Measuring the interestingness of temporal logic behavioral specifications in process mining

Alessio Cecconi^{a,*}, Giuseppe De Giacomo^{b,1}, Claudio Di Ciccio^{b,1},
Fabrizio Maria Maggi^{c,1}, Jan Mendling^{d,a,1}

^a Vienna University of Economics and Business, Vienna, Austria

^b Sapienza University of Rome, Rome, Italy

^c Free University of Bozen-Bolzano, Bolzano, Italy

^d Humboldt-Universität zu Berlin, Berlin, Germany

ARTICLE INFO

Article history:

Received 14 March 2021

Received in revised form 21 September 2021

Accepted 13 October 2021

Available online 1 November 2021

Recommended by Gottfried Bossen

Keywords:

Declarative process mining

Specification mining

Association rule mining

Interestingness measures

Temporal rules

ABSTRACT

The assessment of behavioral rules with respect to a given dataset is key in several research areas, including declarative process mining, association rule mining, and specification mining. An assessment is required to check how well a set of discovered rules describes the input data, and to determine to what extent data complies with predefined rules. Particularly in declarative process mining, Support and Confidence are used most often, yet they are reportedly unable to provide a sufficiently rich feedback to users and cause rules representing coincidental behavior to be deemed as representative for the event logs. In addition, these measures are designed to work on a predefined set of rules, thus lacking generality and extensibility. In this paper, we address this research gap by developing a measurement framework for temporal rules based on (LTL_p). The framework is suitable for any temporal rules expressed in a reactive form and for custom measures based on the probabilistic interpretation of such rules. We show that our framework can seamlessly adapt well-known measures of the association rule mining field to declarative process mining. Also, we test our software prototype implementing the framework on synthetic and real-world data, and investigate the properties characterizing those measures in the context of process analysis.

© 2021 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Measuring the degree to which process traces comply with behavioral rules is key in process analysis branches such as conformance checking [1], compliance assessment [2], and discovery of process constraints [3]. To date, several measures have been defined to this end, yet there are two major problems with their application.

First, measures adopted for process mining are defined inconsistently for specific applications. For example, among the most frequently used measures there are Support and Confidence. However, their definition has been customized to the specification languages in use and even for the specific mining algorithms under analysis. For instance, there is a significant difference in the definition of Support used in [3] (percentage of

traces fully compliant to a rule) and [4] (percentage of activations that lead to a fulfillment), in a way that the Support of rule “If *a* is executed, then *b* will be executed later” on a set of traces like $\{(a, b, c, d), (a, b, c, a), (a, c)\}$ is equal to 0.33 for [3] and 0.5 according to [4]. Furthermore, the definition of those measures are defined ad hoc for specific sets of rules, like DECLARE [5] templates. Such issues hinder the fair comparison and eventually the advancement of rule-based process mining.

Second, the opportunity to adopt available measures from association rule mining has been largely missed so far. A plethora of measures that are reportedly superior in comparison to Support and Confidence [6] have been proposed in this field. Support measures only the satisfaction frequency of a rule and Confidence its validity. Although those are crucial aspects in the assessment of a rule, they do not suffice to avoid spurious result [7]. Markedly, various directions have been explored in prior research, among others by revising Support and Confidence [8–10], and by defining complementary measures [11–13]. However, all such measures do not account for the temporal perspective, which is a first-class citizen dimension in process mining.

In this paper, we address the research challenge of defining a general and comprehensive measurement system. More

* Corresponding author.

E-mail addresses: alessio.cecconi@wu.ac.at (A. Cecconi), giuseppe.degiacomo@uniroma1.it (G. De Giacomo), claudio.diciccio@uniroma1.it (C. Di Ciccio), maggi@inf.unibz.it (F. M. Maggi), jan.mendling@hu-berlin.de, jan.mendling@wu.ac.at (J. Mendling).

¹ Authors in alphabetical order.

specifically, we propose a framework based on formal semantics grounded in Linear-time Temporal Logic with Past on Finite Traces (LTL_P) to express Reactive Constraints (RCons) [14] in a way that abstracts from specific rule-specification languages. Such constraints are rules in the form of “IF *A* THEN *B*”, thus binding the satisfaction of an antecedent *A* to the occurrence of a consequent *B*, wherein both *A* and *B* are temporal formulas. We show that a probabilistic interpretation of the fine-grained temporal logic evaluation of any such formulas allows us to employ all available association rule mining measures as-is for temporal rules. Markedly, the framework has linear time and space complexity with respect to the input size.

Our contribution extends concepts from association rule mining to temporal logic specifications. In this way, we define a foundation upon which the fitness between measures and data analysis scenarios can be discussed in future research. We conduct an extensive set of simulation experiments, the results of which demonstrate that, driven by known properties, the measures respond differently to changes in the behavior evidenced by event logs. This is an important finding that highlights the need to select measures according to the application context, confirming previous findings for association rules [15] in the realm of temporal logic specifications.

This paper is an extension of our previous conference paper [16] presented at the 2nd International Conference on Process Mining (ICPM 2020). We extend the contribution in the following aspects:

- We extend the framework to provide measures at the level of the event log, and not only descriptive statistics at the trace level (Section 4);
- We revise the experiments based on the new theoretical extension. In particular, we score the proposed measures and rank them in order to identify the best candidates to be used for rule discovery (Section 6);
- We analyze the memory consumption of the framework along with the time performance (Section 5);
- We extend the discussion about the interestingness measures used and exploit their known properties for a better understanding of log behavior (Sections 3 and 6).

Additionally, we prove the linear-time performance of the RCons verification in [Appendix](#).

The remainder of this paper is structured as follows. Section 2 discusses prior research on measures for declarative process mining and specification mining. Section 3 defines preliminaries upon which we define our framework. Section 4 defines the measurement framework. Section 5 presents a computational study of the framework and Section 6 shows the results of an array of simulation experiments and discusses them. Finally, Section 7 summarizes the contribution of the paper and points to opportunities for future research.

2. Related work

Behavioral rules have been widely used to support application scenarios such as association rule mining in machine learning, process discovery and conformance checking in process mining, and specification mining in software engineering. The assessment of rules with respect to the available data is a key component of all these techniques.

In association rule mining, interestingness measures are used to discriminate candidate pairs of relevant co-occurring events. A common technique is to discover frequent rules above a certain Support threshold (frequency) and to prune the results below a certain Confidence threshold (validity). For example, [17] discovers associations between items through an Apriori algorithm

based on the downward-closure property of the Support measure. Nevertheless, the use of Support and Confidence alone is reportedly not sufficient to avoid a large number of spurious results [7], i.e., the discovery of rules which are frequently satisfied by the data although merely by chance, thus threatening their statistical validity. A plethora of new measures have been proposed in the literature to overcome the limits of using only Support and Confidence [11], yet the employment of Support and Confidence as the main interestingness measures remains widespread. The main goal driving the development of better measures is indeed the exclusion of spurious rules, so as to let the more crucial ones stand out. Several measures are directly improving on or refining the results of Support and Confidence (e.g., Lift [8] scales Confidence with the Support of the consequent of a rule), others combine different measures (e.g., Added Value [12] subtracts the Prevalence to the Confidence of a rule), and further ones show complementary information (e.g., Specificity [11] measures to what extent the absence of the consequent is related to the absence of the antecedent of a rule).

In declarative process discovery, interestingness measures are used to prune candidate rules based on user-defined thresholds. This pruning approach is used for DECLARE discovery in [3,4] and for DCR graphs discovery in [18]. These techniques are mainly based on Support and Confidence, which lead to the aforementioned limits [7]. In addition, the definitions of these metrics also differ depending on the techniques that use them. For instance, the Support measure presented in [3] does not correspond to the Support measure of [4], although both are expressly defined for the sole DECLARE constraints.

In the area of conformance checking, interestingness measures are used to check the degree of conformance of a rule with respect to an execution trace. In [19], Linear Temporal Logic (LTL) rules are checked against each trace in a given event log. This is highly generic as it supports any custom LTL formula, but it reports only binary results, i.e., whether a rule holds in a trace or not. In [20], Burattin et al. use measures like fulfillment ratio and violation ratio, based on the evaluation of the number of activations of a rule (intuitively, the occurrences of its antecedent) that lead to a fulfillment and the number of activations that lead to a violation in an event log. However, these metrics are specifically bound to the set of DECLARE rules, thus not providing a general measurement framework that can be applied to general type of rule.

In specification mining, interestingness measures are also used to prune candidate temporal specifications based on user-defined thresholds. Interestingly, specification mining and declarative process discovery are two largely overlapping concepts from distinct fields. Yang et al. [21] discover 2-value temporal patterns using a trace measure that quantifies partial satisfactions of a rule. Yet, the technique lacks generality as it is limited only to alternation patterns (similar to ALTERNATERESPONSE and ALTERNATEPRECEDENCE in [Table 1](#)) and the adopted computation heuristics are tailored to the software domain. Le et al. [6] emphasize the limits of using only Support and Confidence measures and investigate properties of other measures reviewed in [11]. Their results demonstrate that there are several measures outperforming Support and Confidence, and that the combination of different measures yields better results. However, they limit their study to 2-value temporal patterns (specifically, RESPONSE and PRECEDENCE in [Table 1](#)). Furthermore, their computation of the probability for a temporal specification is based on a sliding window technique [22]: traces are read in chunks of the size of a given window, then the probability of a rule is the percentage of windows in which it is satisfied. They test the effect of different window sizes, showing that their results depend not only on the input rules and the data, but also on this parameter selection.

Lemieux et al. [23] extend specification mining to arbitrary LTL specifications (implicitly on finite traces) beyond 2-value templates. However, they resort to the sole Support and Confidence measures to prune uninteresting results, thus incurring in the already mentioned statistical limits [7].

The aforementioned shortcomings of quality measures are also discussed in the field of sequence mining [24] when dealing with discovering patterns (specifically subsequences) to classify sequential data. Eggho et al. [25] highlight how measures like Confidence and Lift alone lead to unstable classification results of subsequences, proposing a probabilistic Bayesian-based measure to overcome such an instability and avoiding the requirement of setting thresholds for measures. It falls under the family of techniques based on the minimum description length principle, like [26], where an encoding scheme is used to discover a minimal set of subsequences which can reproduce the original data. Notably, subsequence interleaving patterns are only a subset of patterns expressible with LTL formulae. Works adopting behavioral rules for classification like [27], on the other hand, fall back to the sole employment of Support.

In summary, despite the discussion in different fields on measures and the problem of spurious relations, there is no technique that supports at the same time a comprehensive and extensible multi-measurements assessment of rules and its applicability on general temporal logic specifications.

3. Preliminaries

To develop our framework, we build on the sound foundations of $LTLp_f$. In this section, we introduce the fundamentals of $LTLp_f$ formulae (Section 3.1) and interestingness measures for association rules (Section 3.2).

3.1. Linear-time Temporal Logic with Past on Finite Traces ($LTLp_f$)

As the formal foundations of our framework, we consider the rules specified in Linear Temporal Logic on Finite Traces (LTL_f) [28] as used in DECLARE [5,29]. LTL_f has the same syntax as LTL [30]. Its semantics is interpreted on finite traces (here abstracted as finite sequences of symbols), and thus take into account that business processes are assumed to eventually terminate [31]. DECLARE focuses on a set of specific LTL_f formulae. Table 1 illustrates some of the most important rules for business process specifications in DECLARE.

$LTLp_f$ is an extension of LTL_f supporting the expression of properties of the past (hence the “p” suffix) [14]. Well-formed Linear-time Temporal Logic with Past on Finite Traces ($LTLp_f$) formulae are built from an alphabet $\Sigma \supseteq \{a\}$ of propositional symbols and are closed under the boolean connectives, the unary temporal operators \circ (next) and \ominus (previous), and the binary temporal operators \mathbf{U} (Until) and \mathbf{S} (Since):

$$\varphi := a | (\neg\varphi) | (\varphi_1 \wedge \varphi_2) | (\circ\varphi) | (\varphi_1 \mathbf{U} \varphi_2) | (\ominus\varphi) | (\varphi_1 \mathbf{S} \varphi_2).$$

From these basic operators, the following can be derived: Classical boolean abbreviations *True*, *False*, \vee , \rightarrow ; Constant $t_{\text{End}} \equiv \neg\circ\text{True}$, denoting the last instant of a trace; Constant $t_{\text{Start}} \equiv \neg\ominus\text{True}$, denoting the first instant of a trace; $\diamond\varphi \equiv \text{True} \mathbf{U} \varphi$ indicating that φ holds true eventually before t_{End} ; $\varphi_1 \mathbf{W} \varphi_2 \equiv (\varphi_1 \mathbf{U} \varphi_2) \vee \square\varphi_1$, which relaxes \mathbf{U} as φ_2 may never hold true; $\diamond\varphi \equiv \text{True} \mathbf{S} \varphi$ indicating that φ holds true eventually in the past, after t_{Start} ; $\square\varphi \equiv \neg\diamond\neg\varphi$ indicating that φ holds true from the current instant till t_{End} ; $\boxplus\varphi \equiv \neg\diamond\neg\varphi$ indicating that φ holds true from t_{Start} to the current instant.

Given a finite trace t of length $n \in \mathbb{N}$, an $LTLp_f$ formula φ is satisfied in a given instant i ($1 \leq i \leq n$) by induction of the following:

$$\begin{aligned} t, i &\models \text{True}; t, i \not\models \text{False}; \\ t, i &\models a \text{ iff } t(i) \text{ is assigned with } a; \\ t, i &\models \neg\varphi \text{ iff } t, i \not\models \varphi; \\ t, i &\models \varphi_1 \wedge \varphi_2 \text{ iff } t, i \models \varphi_1 \text{ and } t, i \models \varphi_2; \\ t, i &\models \circ\varphi \text{ iff } i < n \text{ and } t, i+1 \models \varphi; \\ t, i &\models \ominus\varphi \text{ iff } i > 1 \text{ and } t, i-1 \models \varphi; \\ t, i &\models \varphi_1 \mathbf{U} \varphi_2 \text{ iff } t, j \models \varphi_2 \text{ with } i \leq j \leq n, \text{ and } t, k \models \varphi_1 \text{ for all } k \\ &\text{s.t. } i \leq k < j; \\ t, i &\models \varphi_1 \mathbf{S} \varphi_2 \text{ iff } t, j \models \varphi_2 \text{ with } 1 \leq j \leq i, \text{ and } t, k \models \varphi_1 \text{ for all } k \\ &\text{s.t. } j < k \leq i. \end{aligned}$$

A formula φ is satisfied by a trace t , written $t \models \varphi$ iff $t, 1 \models \varphi$. One of the central properties of $LTLp_f$ and LTL_f is that a deterministic finite state automaton (DFS) A_φ can be computed such that for every trace t we have $t \models \varphi$ iff t is in the language recognized by A_φ , as illustrated in [14,32,33].

Without loss of generality, in this paper we abstract traces as finite strings of symbols representing *events*. We assume that every event reports on the execution of exactly one task and $LTLp_f$ formulae use those tasks as their propositional symbols – the so-called DECLARE *assumption* [32]. A trace extracted from the real-world Sepsis event log [34] is, e.g., $t_{\text{Sepsis}} = \langle \text{ER Registration, ER Triage, ER Sepsis Triage, CRP, Lactic Acid, IV Liquid, IV Antibiotics} \rangle$. Notice that this trace complies with the constraints that Mannhardt et al. identified as normative for the Sepsis treatment process [35], including the following ones: (i) $\text{INIT}(\text{ER Registration})$, i.e., every trace begins with the registration at the emergency department; (ii) $\text{ATMOSTONE}(\text{ER Triage})$, i.e., the triage in the emergency room occurs at most once in a process run; (iii) $\text{RESPONSE}(\text{ER Triage, ER Sepsis Triage})$, i.e., the ER Triage procedure should be eventually followed by the Sepsis-specific triage; (iv) $\text{PRECEDENCE}(\text{ER Sepsis Triage, IV Antibiotics})$, i.e., the intravenous injection of antibiotics must be preceded by the ER Sepsis Triage procedure.

Table 17 contains an extended set of DECLARE constraints that a correct execution of the Sepsis treatment process should fulfill. For the sake of succinctness, we may use single-letter identifiers in place of full-length task names whenever suitable in the following examples.

An *event log* is a multi-set of traces, i.e., traces can recur multiple times in an event log. The *cardinality* of the event log is the sum of the multiplicities of its traces. Considering the Sepsis event log, the multiplicity of t_{Sepsis} is 13 (i.e., t_{Sepsis} occurs 13 times in the event log). Other exemplary traces of that log are $t'_{\text{Sepsis}} = \langle \text{ER Registration, ER Triage, ER Sepsis Triage} \rangle$ (occurring 35 times) and $t''_{\text{Sepsis}} = \langle \text{ER Registration, ER Triage, ER Sepsis Triage, Leucocytes, CRP} \rangle$ (with a multiplicity of 24). The cardinality of the event sub-log consisting of the traces above is thus 72. The cardinality of the whole Sepsis event log is 1050.

3.2. Interestingness measures for association rules

In this section, we revisit key findings of research on quality measures in association rule mining. More specifically, we build on prior research that surveys measures in the area of software engineering and association rule mining, namely [11,15,36,37]. These works are specifically suited as a foundation due to the wide coverage of measures and their comparative study of both formal and user-perceived measure properties. The rules under study are in the form “IF A THEN B”, where A is called the *antecedent* of the rule, and B its *consequent*. We refer to A and B as *variables* of the rule.

Specifically, we consider only probability-based *objective* measures, i.e., measures depending only on the data as opposed to those requiring user-provided parameters. Objective measures are based on the probabilities derived from the contingency table of the occurrences of the variables, as depicted in Table 2. Table 3 presents the list of measures covered in this study. In the following, we provide a brief description of each measure.

Table 1
Some DECLARE constraints expressed as RCons.

Constraint	LTL _f expression [28]	RCon
PARTICIPATION(a)	$\diamond a$	$t_{\text{Start}} \square \rightarrow \diamond a$
INIT(a)	a	$t_{\text{Start}} \square \rightarrow a$
END(a)	$\square \diamond a$	$t_{\text{End}} \square \rightarrow a$
ATMOSTONE(a)	$\square(a \rightarrow \bigcirc(\neg \diamond a))$	$a \square \rightarrow \bigcirc(\neg \diamond a)$
RESPONDEXISTENCE(a, b)	$\diamond a \rightarrow \diamond b$	$a \square \rightarrow (\diamond b \vee \diamond b)$
RESPONSE(a, b)	$\square(a \rightarrow \diamond b)$	$a \square \rightarrow \diamond b$
ALTERNATERESPONSE(a, b)	$\square(a \rightarrow \diamond b) \wedge \square(a \rightarrow \bigcirc(\neg a \mathbf{W} b))$	$a \square \rightarrow \bigcirc(\neg a \mathbf{U} b)$
CHAINRESPONSE(a, b)	$\square(a \rightarrow \diamond b) \wedge \square(a \rightarrow \bigcirc b)$	$a \square \rightarrow \bigcirc b$
PRECEDENCE(a, b)	$\neg b \mathbf{W} a$	$b \square \rightarrow \diamond a$
ALTERNATEPRECEDENCE(a, b)	$(\neg b \mathbf{W} a) \wedge \square(b \rightarrow \bigcirc(\neg b \mathbf{W} a))$	$b \square \rightarrow \ominus(\neg b \mathbf{S} a)$
CHAINPRECEDENCE(a, b)	$(\neg b \mathbf{W} a) \wedge \square(\bigcirc b \rightarrow a)$	$b \square \rightarrow \ominus a$
SUCCESSION(a, b)	$\square(a \rightarrow \diamond b) \wedge (\neg b \mathbf{W} a)$	$(a \vee b) \square \rightarrow (a \wedge \diamond b) \vee (b \wedge \diamond a)$
ALTERNATESUCCESSION(a, b)	$\square(a \rightarrow \diamond b) \wedge \square(a \rightarrow \bigcirc(\neg a \mathbf{W} b))$	$(a \vee b) \square \rightarrow (a \wedge \bigcirc(\neg a \mathbf{U} b)) \vee (b \wedge \ominus(\neg b \mathbf{S} a))$
CHAINSUCCESSION(a, b)	$\square(a \rightarrow \diamond b) \wedge \square(a \rightarrow \bigcirc b)$	$(a \vee b) \square \rightarrow (a \wedge \bigcirc b) \vee (b \wedge \ominus a)$
COEXISTENCE(a, b)	$\wedge(\neg b \mathbf{W} a) \wedge \square(\bigcirc b \rightarrow a)$	$(a \vee b) \square \rightarrow (a \wedge \diamond b) \vee (a \wedge \diamond b) \vee (b \wedge \diamond a) \vee (b \wedge \diamond a)$
NOTCOEXISTENCE(a, b)	$\square(a \rightarrow \neg \diamond b) \wedge \square(b \rightarrow \neg \diamond a)$	$(a \vee b) \square \rightarrow (a \wedge \neg \diamond b \wedge \neg \diamond b) \vee (b \wedge \neg \diamond a \wedge \neg \diamond a)$

Table 2
Contingency tables to which the “IF A THEN B” rules and their variables ($x \in \{A, B\}$) comply with. On the left-hand side, the contingency table is based on probabilities ($P(x)$); on the right-hand side, the contingency table is based on frequencies (where $|x|$ is the number of occurrences of x and N is the total number of occurrences).

	A	$\neg A$		A	$\neg A$	
B	$P(AB)$	$P(\neg AB)$	$P(B)$	$ AB $	$ \neg AB $	$ B $
$\neg B$	$P(A\neg B)$	$P(\neg A\neg B)$	$P(\neg B)$	$ A\neg B $	$ \neg A\neg B $	$ \neg B $
	$P(A)$	$P(\neg A)$	1	$ A $	$ \neg A $	N

Support [17] measures the frequency of the co-occurrence of the antecedent and the consequent. The Support of the sole antecedent of the rule is also called **Coverage**, while the Support of the consequent is called **Prevalence**.

Confidence [17] measures the co-occurrences of the antecedent and the consequent in the fraction of data containing the antecedent.

Recall [11] measures the co-occurrences of the antecedent and the consequent in the fraction of data containing the consequent.

Specificity [11] measures the co-absences of the antecedent and the consequent in the fraction of data not containing the antecedent.

Accuracy [11] measures the fraction of the data either containing both the consequent and the antecedent or neither of the two.

Lift [8] scales the Confidence by the probability of the consequent, to check if the co-occurrence of the antecedent and consequent is more likely than their independence.

Leverage [9] measures the difference between the Confidence of the rule and the independent occurrence of its variables.

Added Value [12] measures the difference between the Confidence of the rule and the probability of the consequent alone, to check if the conditioned occurrence of the consequent differ from its unconditioned occurrence.

Relative Risk [38] measures the ratio of the conditional probability of the consequent given the antecedent to the conditional probability of the consequent given the negation of the antecedent.

Jaccard's Coefficient [39] measures the similarity between the variables using the ratio of their co-occurrence to the union of all their independent occurrences.

Certainty Factor [10] measures the ratio of the Added Value of the rule to the Added Value of the consequent alone, in order to see the variation of probability in the data containing the antecedent.

Odds Ratio [40] measures the ratio of the probability of having the consequent when the antecedent is present to the probability of having the consequent when the antecedent is not present.

Odds Multiplier [40] measures the ratio of the probability of having the antecedent when the consequent is present to the probability of having the antecedent when the consequent is not present.

Yules's Q and Yules's Y [41,42] are normalization of the Odds Ratio to have it centered around 0 and ranging between -1 and 1.

Klosgen's Measure [43] weights the Support of the rule using its Added Value.

Conviction [13] measures the occurrences of the antecedent without the consequent in comparison to their independence.

Interestingness Weighting Dependency [44] combines Support and Lift of a rule and explicitly gives weights to each of them to let the user decide their relative importance.

Collective Strength [12] measures the ratio of the agreement ratio (number of non-violations per expected number of non-violations) to the violation ratio (number of violations per expected number of violations).

Table 3
Probabilistic measures for association rules with respective range and properties.

Measure	Formula	Range	P1	P2	P3	P4	P5	P6
Support	$P(AB)$	[0, 1]	-	sym	↘	var	var	↘
Confidence/Precision	$P(B A)$	[0, 1]	✓	asym	→	var	const	↘
Coverage	$P(A)$	[0, 1]	-	asym	↘	var	var	↘
Prevalence	$P(B)$	[0, 1]	-	asym	↗	var	var	↘
Recall	$P(A B)$	[0, 1]	✓	asym	↘	var	var	↘
Specificity	$P(\neg B \neg A)$	[0, 1]	-	asym	↘	const	var	↘
Accuracy	$P(AB) + P(\neg A\neg B)$	[0, 1]	-	sym	↘	var	var	↘
Lift/Interest	$\frac{P(B A)}{P(B)}$ or $\frac{P(AB)}{P(A)P(B)}$	[0, +∞)	-	sym	↘	const	var	↘
Leverage	$P(B A) - P(A)P(B)$	[-1, 1]	-	sym	↗	var	var	↘
Added Value/ Change of Support/ Centered Confidence	$P(B A) - P(B)$	[-1, 1]	-	asym	↘	const	var	↘
Relative risk	$\frac{P(B A)}{P(B \neg A)}$	[0, +∞)	-	asym	↘	const	var	↘
Jaccard	$\frac{P(AB)}{P(A) + P(B) - P(AB)}$	[0, 1]	✓	sym	↘	var	var	↘
Certainty factor	$\frac{P(B A) - P(B)}{1 - P(B)}$	[-1, 1]	-	asym	↘	const	const	↘
Odds ratio/ Bayes Factor	$\frac{P(AB)P(\neg A\neg B)}{P(A\neg B)P(\neg BA)}$	[0, +∞)	-	sym	↘	const	const	↘
Yule's Q	$\frac{P(AB)P(\neg A\neg B) - P(A\neg B)P(\neg AB)}{P(AB)P(\neg A\neg B) + P(A\neg B)P(\neg AB)}$	[-1, 1]	-	sym	↘	const	const	↘
Yule's Y	$\frac{\sqrt{P(AB)P(\neg A\neg B)} - \sqrt{P(A\neg B)P(\neg AB)}}{\sqrt{P(AB)P(\neg A\neg B)} + \sqrt{P(A\neg B)P(\neg AB)}}$	[-1, 1]	-	sym	↘	const	const	↘
Klosgen	$\sqrt{P(AB)} \times \max(P(B A) - P(B), P(A B) - P(A))$	[-1, 1]	-	asym	↘	const	var	↘
Conviction	$\frac{P(A)P(\neg B)}{P(A\neg B)}$	[0, +∞)	-	asym	↘	const	const	↘
Interestingness Weighting Dependency	$\left(\left(\frac{P(AB)}{P(A)P(B)} \right)^k - 1 \right) \times P(AB)^m$	[0, +∞)	-	sym	↘	const	var	↘
Collective Strength	$\frac{P(AB) + P(\neg B \neg A)}{P(A)P(B) + P(\neg A)P(\neg B)} \times \frac{1 - P(A)P(B) - P(\neg A)P(\neg B)}{1 - P(AB) - P(\neg B \neg A)}$	[0, +∞)	-	asym	?	var	var	↘
Laplace Correction	$\frac{N(AB) + 1}{N(A) + 2}$	[0.5, 1]	✓	asym	→	var	var	↘
Gini Index	$\frac{P(A) \times (P(B A)^2 + P(\neg B A)^2) + P(\neg A) \times (P(B \neg A)^2 + P(\neg B \neg A)^2) - P(B)^2 - P(\neg B)^2}{2}$	[0, 1]	-	asym	↘	const	var	↘
J-Measure	$P(AB) \log_2 \frac{P(B A)}{P(B)} + P(A\neg B) \log_2 \frac{P(\neg B A)}{P(\neg B)}$	(-∞, +∞)	-	asym	↘	const	const	↘
One-Way Support	$P(B A) \log_2 \frac{P(AB)}{P(A)P(B)}$	(-∞, +∞)	-	asym	↘	const	var	↘
Two-Way Support	$P(AB) \log_2 \frac{P(AB)}{P(A)P(B)}$	(-∞, +∞)	-	sym	↘	const	var	↘
Two-Way Support Variation	$P(AB) \log_2 \frac{P(AB)}{P(A)P(B)} + P(A\neg B) \log_2 \frac{P(A\neg B)}{P(A)P(\neg B)}$ $+ P(\neg AB) \log_2 \frac{P(\neg AB)}{P(\neg A)P(B)} + P(\neg A\neg B) \log_2 \frac{P(\neg A\neg B)}{P(\neg A)P(\neg B)}$	(-∞, +∞)	-	sym	↘	const	const	↘

(continued on next page)

Laplace Correction [45] is a variation of Confidence to take into account small data.

Gini index [46] measures if the entropy introduced by a rule brings a marked difference.

J-measure [47] is an entropy based measure for the information content of a rule.

One-way Support and Two-way Support [48] combine respectively Confidence and Support of a rule with the degree of independence between the variables.

Table 3 (continued).

Measure	Formula	Range	P1	P2	P3	P4	P5	P6
\emptyset -Coefficient (Pearson's Linear Correlation Coefficient)	$\frac{P(AB) - P(A)P(B)}{\sqrt{P(A)P(B)P(\neg A)P(\neg B)}}$	$(-\infty, +\infty)$	-	sym	\searrow	const	var	\searrow
Piatetsky-Shapiro	$P(AB) - P(A)P(B)$	$[-1, 1]$	-	sym	\searrow	const	var	\searrow
Cosine	$\frac{P(AB)}{\sqrt{P(A)P(B)}}$	$[0, +\infty)$	\checkmark	sym	\searrow	var	var	\searrow
Loevinger	$1 - \frac{P(A)P(\neg B)}{P(A \rightarrow B)}$	$(-\infty, 1]$	-	asym	\nearrow	const	const	\curvearrowright
Information Gain	$\log \frac{P(AB)}{P(A)P(B)}$	$(-\infty, +\infty)$	-	sym	\searrow	const	var	\curvearrowright
Sebag-Schoenauer	$\frac{P(AB)}{P(A \rightarrow B)}$	$[0, +\infty)$	\checkmark	asym	\rightarrow	var	const	\curvearrowright
Least Contradiction	$\frac{P(AB) - P(A \rightarrow B)}{P(B)}$	$(-\infty, +\infty)$	\checkmark	asym	\searrow	var	var	\searrow
Odd Multiplier	$\frac{P(AB)P(\neg B)}{P(B)P(A \rightarrow B)}$	$[0, +\infty)$	-	asym	\searrow	const	const	\curvearrowright
Example and Counterexample Rate	$1 - \frac{P(A \rightarrow B)}{P(AB)}$	$(-\infty, 1]$	\checkmark	asym	\rightarrow	var	const	\curvearrowright
Zhang	$\frac{P(AB) - P(A)P(B)}{\max(P(AB)P(\neg B), P(B)P(A \rightarrow B))}$	$(-\infty, +\infty)$	-	asym	\searrow	const	const	\curvearrowright

Two-way Support Variation [48] measures the change in the Two-way-Support.

Linear Correlation Coefficient [49] measures the Pearson's correlation between the variables.

Piatetsky-Shapiro [9] measures the difference between the co-occurrences of antecedent and consequent and their independent frequency.

Cosine [37] measures the geometric mean between Lift and Support of a rule.

Loevinger [50] measures the homogeneity between antecedent and consequent.

Information Gain [51] is the logarithm of the Lift.

Sebag-Schoenauer [52] measures the proportion of positive and negative occurrences of the antecedent.

Least Contradiction [53] measures the difference between positive and negative occurrences of the antecedent weighted by its frequency.

Example and Counterexample Rate [11] measures the proportion of the antecedent occurrences with and without consequent.

Zhang [54] measures the positive or negative association between the antecedent and the consequent.

Different studies have been dedicated to the analysis of general properties for measures [9,11,15,37]. Properties show the response of measures under certain conditions. Therefore, properties can be used to group similar measures and decide the proper ones to be employed depending on the context. For example, we will analyze the sensitivity of measures to the increase of noise in the data as an important selection criterion for rule monitoring or discovery. We will delve deeper into this aspect in Section 6. In this paper, we focus specifically on a subset of the properties proposed in [37] and [15], as their meaning and effects are reportedly recognizable in a clear manner by the final user. The selected properties are explained below and associated to each measure \mathcal{M} in Table 3.

P1. Null invariance [37]. The measure is unaffected by traces not containing neither A or B . Therefore, it assesses whether the traces not related to the rule affect the measurement or not. To satisfy this property, the measure should not vary when $|\neg A \neg B|$ increases in the contingency table, while the other values remain fixed. In Table 3, we use the ' \checkmark ' or '-' symbols to indicate whether the property holds or not, respectively. For example, for Confidence and Recall this property holds, whereas for Support, Leverage and Collective Strength it does not.

P2. Asymmetric processing of variables [15]. The measure is asymmetric under variable swap, i.e., the measure of IF A THEN B differs from that of IF B THEN A . The measure enjoys this property if it does not vary upon the swapping of the values of $|\neg AB|$ and $|A \rightarrow B|$ in the contingency table. In Table 3, every measure is marked with "asym" or "sym" to indicate whether the property is enjoyed or not, respectively. For instance, Support and Leverage are symmetric under variable swap, whereas Confidence, Recall and Collective Strength provide an asymmetric processing of the variables.

P3. Variation with occurrences of B in the absence of A [15]. The value of the measure varies when the occurrences of B in the absence of A increase. In other words, this property focuses on whether the independent occurrence of B influences the measure. Given an IF A THEN B rule, if B is very likely to occur regardless of A , the influence of A on B may be questioned. To verify this, the value of the measure varies when $|\neg AB|$ increases in the contingency values (and the other values remain fixed). In Table 3, measures are marked with ' \searrow ' if the variation is a decrease (as in the case of Support), ' \nearrow ' if it is an increase (e.g., Leverage), '?' if the variation can be either a decrease or an increase depending on the values of B (e.g., Collective Strength), ' \rightarrow ' if the value does not vary at all (e.g., Confidence).

P4. Reference situations: independence [15]. If the variables are independent, then the measure exhibits a known value. The variables are considered independent when their joint probability is equal to the product of their respective probabilities, i.e., $P(AB) = P(A)P(B)$. The measure should have a constant and known value in that case. In Table 3, measures

are labeled as “const” (e.g., Lift) if this property holds, and “var” otherwise (e.g., Support).

P5. Reference situations: logical rule [15]. If the rule is always satisfied, then the measure exhibits a known value. An IF A THEN B rule is always satisfied if $P(A \rightarrow B) = 0$. In other words, if there are no counterexamples in the data, the value of the measure that enjoys this property is a known constant (let it be a number or tendency to infinite). In Table 3, measures are labeled as “const” if this property holds (as in the case of Confidence), and “var” otherwise (see, e.g., Lift).

P6. Trend with $P(A \rightarrow B)$ [15]. If the number of counterexamples to the rule rises, the value of the measure reacts exhibiting a decreasing trend that denotes a higher or lower sensitivity. For an IF A THEN B rule, a higher number of counterexamples translates into an increase of $P(A \rightarrow B)$. Against that increase, the measure may show a fast (convex), linear, or slow (concave) decrease. Measures in Table 3 are labeled either as ‘ \curvearrowright ’ (convex, e.g., Conviction), ‘ \searrow ’ (linear, e.g., Support), or ‘ \curvearrowleft ’ (concave, e.g., Recall) accordingly.

These properties, according to [15], can be divided into *normative* (i.e., always desirable: P2, P3, P4, P5) and *subjective* (i.e., depending on the user needs: P1, P6).

We will resort to these properties to examine the quality measures in the context of process mining.

In the following section, therefore, we extend the aforementioned measures to temporal process rules.

4. Temporal-extended measurement framework

Our framework addresses the limits of Support and Confidence measurements by building on LTL_p formal semantics and the spectrum of measures defined in different areas of computer science. Furthermore, it is generic as it allows for the usage of any probabilistic measure (including those of Table 3) on any temporal-logic-based rules specification. To this end, Section 4.1 formalizes the reactive temporal specification of rules, Section 4.2 discusses their probabilistic interpretation, and Section 4.4 defines the overall framework.

4.1. Reactive temporal specification

Our first building block is the concept of Reactive Constraint (RCon), originally introduced in [14], the paper which we extend here. A rule typically expresses that the occurrence of given preconditions (activator) implies certain consequences (target). The reactive nature of this kind of rule lies in the fact that the condition on the target is exerted only if the activator is verified. We codify this intuition in RCons, whose semantics is based on LTL_p.

Definition 4.1 (Reactive Constraint (RCon)). Given an alphabet of propositional symbols $\Sigma \cup \{t_{\text{Start}}, t_{\text{End}}, \text{True}, \text{False}\}$, let φ_α and φ_τ be LTL_p formulae over Σ . A *Reactive Constraint (RCon)* Ψ is a pair $(\varphi_\alpha, \varphi_\tau)$ hereafter denoted as $\Psi \triangleq \varphi_\alpha \square \rightarrow \varphi_\tau$.

An RCon is interpreted as follows: Each time the activator is true, the target should be true at that point of the trace. For example, $a \square \rightarrow \diamond c$ is an RCon stating that every time a (the activator, φ_α) is *True*, then also $\diamond c$ (the target, φ_τ) must evaluate to *True*. That RCon corresponds to $\text{RESPONSE}(a, c)$ in DECLARE as it requires that if a occurs in a trace, it must be eventually followed by c . $c \square \rightarrow \diamond d$ corresponds to $\text{PRECEDENCE}(d, c)$ in DECLARE because it requires that every time c (the activator) occurs in a trace, then it has to be preceded by d (the target). Table 1 provides a list of standard DECLARE constraints expressed in the form of RCons.

Table 4

Evaluation (0 is *False* and 1 is *True*) and probabilistic interpretation of RCon $a \square \rightarrow \diamond c$.

Trace $t_1 =$	(a,	b,	c,	d,	f,	c,	e,	c,	h)
$\varphi_\alpha: a$		1	0	0	0	0	0	0	0	0	
$\varphi_\tau: \diamond c$		1	1	1	1	1	1	1	1	1	0
$P(\varphi_\alpha, t) = 1/9$		$P(\neg\varphi_\alpha \cap \varphi_\tau, t) = 7/9$			$P(\neg\varphi_\alpha \cap \neg\varphi_\tau, t) = 1/9$						
$P(\varphi_\tau, t) = 5/9$		$P(\varphi_\alpha \cap \varphi_\tau, t) = 0/9$			$P(\varphi_\alpha \cap \neg\varphi_\tau, t) = 1/9$						
		$P(\varphi_\tau \varphi_\alpha, t) = 1/1$			$P(\varphi_\tau \neg\varphi_\alpha, t) = 7/8$						
		$P(\neg\varphi_\tau \varphi_\alpha, t) = 0/1$			$P(\neg\varphi_\tau \neg\varphi_\alpha, t) = 1/8$						

Table 5

Evaluation (0 is *False* and 1 is *True*) and probabilistic interpretation of RCon $c \square \rightarrow \diamond d$.

Trace $t_1 =$	(a,	b,	c,	d,	f,	c,	e,	c,	h)
$\varphi_\alpha: c$		0	0	1	0	0	1	0	1	0	
$\varphi_\tau: \diamond d$		0	0	0	1	1	1	1	1	1	
$P(\varphi_\alpha, t) = 3/9$		$P(\neg\varphi_\alpha \cap \varphi_\tau, t) = 4/9$			$P(\neg\varphi_\alpha \cap \neg\varphi_\tau, t) = 2/9$						
$P(\varphi_\tau, t) = 6/9$		$P(\varphi_\alpha \cap \varphi_\tau, t) = 1/9$			$P(\varphi_\alpha \cap \neg\varphi_\tau, t) = 2/9$						
		$P(\varphi_\tau \varphi_\alpha, t) = 2/3$			$P(\varphi_\tau \neg\varphi_\alpha, t) = 4/6$						
		$P(\neg\varphi_\tau \varphi_\alpha, t) = 1/3$			$P(\neg\varphi_\tau \neg\varphi_\alpha, t) = 2/6$						

Table 6

Evaluation (0 is *False* and 1 is *True*) and probabilistic interpretation of RCon $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$.

Trace $t_1 =$	(a,	b,	c,	d,	f,	c,	e,	c,	h)
$\varphi_\alpha: (\diamond b \wedge \diamond e)$		0	1	1	1	1	1	1	0	0	
$\varphi_\tau: (\neg c \vee \diamond f)$		1	1	1	1	1	0	1	0	1	
$P(\varphi_\alpha, t) = 6/9$		$P(\neg\varphi_\alpha \cap \varphi_\tau, t) = 2/9$			$P(\neg\varphi_\alpha \cap \neg\varphi_\tau, t) = 1/9$						
$P(\varphi_\tau, t) = 7/9$		$P(\varphi_\alpha \cap \varphi_\tau, t) = 1/9$			$P(\varphi_\alpha \cap \neg\varphi_\tau, t) = 5/9$						
		$P(\varphi_\tau \varphi_\alpha, t) = 5/6$			$P(\varphi_\tau \neg\varphi_\alpha, t) = 2/3$						
		$P(\neg\varphi_\tau \varphi_\alpha, t) = 1/6$			$P(\neg\varphi_\tau \neg\varphi_\alpha, t) = 1/3$						

An RCon that goes beyond the standard repertoire of DECLARE is $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$: Its activator is the formula $\varphi_\alpha = \diamond b \wedge \diamond e$, satisfied between the occurrence of b and the occurrence of e in a trace; its target is the formula $\varphi_\tau = \neg c \vee \diamond f$, which evaluates to *True* if either c is *False*, or c occurs and is eventually followed by f . Because at every event of the trace (i.e., any point in time) both the activator and target can be either *True* or *False*, the possible evaluation of an RCon can result in either of the following four combinations.

Definition 4.2 (RCon Evaluation). Given an RCon $\Psi \triangleq \varphi_\alpha \square \rightarrow \varphi_\tau$ and a trace t of length $n \in \mathbb{N}$, let i denote the i th event in the trace ($1 \leq i \leq n$). For each $t_i \in t$ the possible evaluations of Ψ are:

$$\begin{aligned} \varphi_\alpha = \text{False}, \varphi_\tau = \text{False} & \quad \text{if } t, i \not\models \varphi_\alpha \text{ and } t, i \not\models \varphi_\tau; \\ \varphi_\alpha = \text{False}, \varphi_\tau = \text{True} & \quad \text{if } t, i \not\models \varphi_\alpha \text{ and } t, i \models \varphi_\tau; \\ \varphi_\alpha = \text{True}, \varphi_\tau = \text{False} & \quad \text{if } t, i \models \varphi_\alpha \text{ and } t, i \not\models \varphi_\tau; \\ \varphi_\alpha = \text{True}, \varphi_\tau = \text{True} & \quad \text{if } t, i \models \varphi_\alpha \text{ and } t, i \models \varphi_\tau. \end{aligned}$$

For example, the second and third rows of Tables 4–6 show the evaluation of RCons $a \square \rightarrow \diamond c$ (i.e., $\text{RESPONSE}(a, c)$ in Table 4), $c \square \rightarrow \diamond d$ (i.e., $\text{PRECEDENCE}(d, c)$ in Table 5) and $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$ (Table 6) on trace $\langle a, b, c, d, f, c, e, c, h \rangle$. Notice that φ_α and φ_τ are evaluated separately at every event of a trace.

The RCon evaluation can be performed efficiently based on the automaton-based techniques defined in [14], adapting it for offline verification. The full discussion on this aspect can be found in Appendix, but we briefly outline the rationale here. Intuitively, we resort to [14, Theorem 4]: An RCon can be separated in pure-past, pure-present and pure-future components. The respective sub-formulae contain only past temporal operators, none,

Table 7
Contingency table of the probabilities of an RCon $\varphi_\alpha \square \rightarrow \varphi_\tau$ in a trace.

	φ_α	$\neg\varphi_\alpha$	
φ_τ	$P(\varphi_\alpha \cap \varphi_\tau, t)$	$P(\neg\varphi_\alpha \cap \varphi_\tau, t)$	$P(\varphi_\tau, t)$
$\neg\varphi_\tau$	$P(\varphi_\alpha \cap \neg\varphi_\tau, t)$	$P(\neg\varphi_\alpha \cap \neg\varphi_\tau, t)$	$P(\neg\varphi_\tau, t)$
	$P(\varphi_\alpha, t)$	$P(\neg\varphi_\alpha, t)$	1

or only future ones, respectively. As they are $LTLp_f$ formulae, all components correspond to finite state automata (FSAs). The key point is that, by mirroring pure-past formulae and reversing their automata, a single replay of the sub-trace from the beginning to the activator event keeps track of the truth value of the pure-past formula till that point. As we have knowledge of the whole trace, and thus of the suffix too a fortiori, we can apply the same principle to pure-future formulae too: A single replay from the end of the trace to the activator event keeps track of the truth value of the pure-future formula from that point onwards.

From this optimization, it follows that any $LTLp_f$ formula can be evaluated at each event reading the trace only twice (as in [4, 55] though for any RCon and not just DECLARE constraints): Once from t_{start} to t_{end} (past components) and once from t_{end} to t_{start} (future components). This result implies that the computational cost depends linearly on the number of events in the event log and in the number of rules to verify. Specifically, given an event log L of cardinality $|L|$, assuming that (i) every trace $t \in L$ has a length of up to n , and (ii) $|R|$ rules are under analysis, the cost to evaluate all rules on L is: $O(|L| \times n \times |R|)$.

4.2. Probabilistic interpretation on a trace

The evaluation of RCons indicates whether a rule holds true or false within a trace. In real life, traces often contain noise or partially deviate from desired process specifications. In those occasions wherein the trace may contain also events that do not satisfy the rule, we are interested in understanding *to what degree* a rule is satisfied. As we have previously defined the notion of satisfaction for φ_α and φ_τ on single events (Definition 4.2), we can devise a probabilistic interpretation for RCons over traces.

Definition 4.3 (Probability of an $LTLp_f$ Formula in a Trace). Given an $LTLp_f$ formula φ and a trace t of length $|t| = n$, we define the probability of φ in t^2 as the proportion of the events in t that satisfy φ :

$$P(\varphi, t) = \frac{|\{i \in [1, n] : t, i \models \varphi\}|}{n}.$$

Definition 4.4 (Joint Probability of $LTLp_f$ Formulae in a Trace). Given two $LTLp_f$ formulae φ_1 and φ_2 and a trace t of length n , we define the probability of the intersection of φ_1 and φ_2 in t (joint probability) as the proportion of the events in t that satisfy both φ_1 and φ_2 :

$$P(\varphi_1 \cap \varphi_2, t) = \frac{|\{i \in [1, n] : t, i \models \varphi_1 \text{ and } t, i \models \varphi_2\}|}{n}.$$

The probabilities of the evaluations of activator and target of an RCon follow from the above definitions (Table 7 shows the resulting contingency table):

$$P(\neg\varphi_\alpha \cap \neg\varphi_\tau, t) = \frac{|\{i \in [1, n] : t, i \not\models \varphi_\alpha \text{ and } t, i \not\models \varphi_\tau\}|}{n};$$

$$P(\neg\varphi_\alpha \cap \varphi_\tau, t) = \frac{|\{i \in [1, n] : t, i \not\models \varphi_\alpha \text{ and } t, i \models \varphi_\tau\}|}{n};$$

² Notice that we use the comma in $P(\varphi, t)$ and similar following expressions to separate the parameters, namely the formula to be evaluated (here, φ) and the structure on which the formula is analyzed (here, t).

Table 8
Contingency table of the probabilities of RCon $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$ in trace $\langle a, b, c, d, f, c, e, c, h \rangle$ (based on the results illustrated in Table 6).

	$\diamond b \wedge \diamond e$	$\neg(\diamond b \wedge \diamond e)$	
$\neg c \vee \diamond f$	5/9	2/9	7/9
$\neg(\neg c \vee \diamond f)$	1/9	1/9	2/9
	2/3	1/3	1

Table 9
Trace measures computation and event log statistics of a sample of measures for RCon $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$. The statistics are computed skipping divisions of zero by zero (marked with "NaN"), whenever they occur.

Event log	Support	Confidence	Specificity	Lift
	$P(\varphi_\alpha \cap \varphi_\tau)$	$P(\varphi_\tau \varphi_\alpha)$	$P(\neg\varphi_\tau \neg\varphi_\alpha)$	$\frac{P(\varphi_\alpha \cap \varphi_\tau)}{P(\varphi_\alpha)P(\varphi_\tau)}$
$t_1 = \langle a, b, c, d, f, c, e, c, h \rangle$	0.56	0.83	0.33	1.07
$t_2 = \langle b, d, a, f, g, d, e, d \rangle$	0.88	1.00	0.00	1.00
$t_3 = \langle a, c, d, b, c, e, f, c \rangle$	0.38	1.00	0.20	1.14
$t_4 = \langle b, c, c, e, a \rangle$	0.40	0.50	0.00	0.83
$t_5 = \langle b, c, d, a \rangle$	0.00	NaN	0.25	NaN
Mean	0.44	0.83	0.16	1.01
Standard deviation	0.32	0.24	0.15	0.13
Variance	0.10	0.06	0.02	0.02

$$P(\varphi_\alpha \cap \neg\varphi_\tau, t) = \frac{|\{i \in [1, n] : t, i \models \varphi_\alpha \text{ and } t, i \not\models \varphi_\tau\}|}{n};$$

$$P(\varphi_\alpha \cap \varphi_\tau, t) = \frac{|\{i \in [1, n] : t, i \models \varphi_\alpha \text{ and } t, i \models \varphi_\tau\}|}{n}.$$

For example, Tables 4–6 show the probabilities resulting from the evaluation of RCons $a \square \rightarrow \diamond c$, $c \square \rightarrow \diamond d$, and $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$, respectively, on trace $\langle a, b, c, d, f, c, e, c, h \rangle$. Table 8 summarizes the results of Table 6 in a contingency table.

In association rule mining, rules are in the form "IF A THEN B ", given an antecedent A and a consequent B . Probabilities defined as above allow for the application of measures defined for association rule mining [11] to the context of temporal logic specifications over finite traces. To that extent, it suffices to map φ_α to A and φ_τ to B , thus having $P(A)$ as $P(\varphi_\alpha, t)$, $P(B)$ as $P(\varphi_\tau, t)$, and $P(AB)$ as $P(\varphi_\alpha \cap \varphi_\tau, t)$. For example, Table 9 shows some measures computed from the probabilities associated to the activator and the target of RCon $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$. These probabilities pertain to the events in the traces, intuitively answering the question: "How likely is it that an event satisfies the constraint?". It follows that also the measures based on them pertain to events with respect to traces, and that their statistics over the entire event logs will preserve the focus on the singles events.

4.3. Probabilistic interpretation on an event log

Following the probability definition for $LTLp_f$ formulae over traces, it is of interest to define similar probabilities over event logs. Intuitively, if the trace probabilities assess the likelihood of the rule correctness in events within a trace, event log probabilities should question the likelihood of the rule correctness in the traces of an event log. As previously mentioned, the descriptive statistics of trace measures across an event log are suitable for this purpose because they preserve the focus on the events. In order to achieve this goal, we have to first derive the conditional probability of the target given the activator in a trace, i.e., $P(\varphi_\tau | \varphi_\alpha, t)$. Intuitively, this is the probability for the target to hold true when the activator holds true. Notice that this viewpoint is conceptually closer to the notion of Reactive Constraint than the joint probability of activator and target. Furthermore, the conditional interpretation of rules is also more in line with their

human interpretation [56]. This makes the conditional probability a suitable means for the probabilistic analysis of a constraint in a trace as a whole.

Definition 4.5 (Conditional Probability of LTL_{p_f} Formulae in a Trace). Given two LTL_{p_f} formulae φ_1 and φ_2 and a trace t of length n , we define the conditional probability of φ_2 given φ_1 over t as the proportion of events satisfying φ_2 among those that satisfy φ_1 :

$$P(\varphi_2|\varphi_1, t) = \frac{|\{i \in [1, n] : t, i \models \varphi_1 \text{ and } t, i \models \varphi_2\}|}{|\{i \in [1, n] : t, i \models \varphi_1\}|}.$$

From the above definition, it follows that:

$$\begin{aligned} P(\varphi_\tau|\varphi_\alpha, t) &= \frac{|\{i \in [1, n] : t, i \models \varphi_\alpha \text{ and } t, i \models \varphi_\tau\}|}{|\{i \in [1, n] : t, i \models \varphi_\alpha\}|}; \\ P(\neg\varphi_\tau|\varphi_\alpha, t) &= \frac{|\{i \in [1, n] : t, i \models \varphi_\alpha \text{ and } t, i \not\models \varphi_\tau\}|}{|\{i \in [1, n] : t, i \models \varphi_\alpha\}|}; \\ P(\varphi_\tau|\neg\varphi_\alpha, t) &= \frac{|\{i \in [1, n] : t, i \not\models \varphi_\alpha \text{ and } t, i \models \varphi_\tau\}|}{|\{i \in [1, n] : t, i \not\models \varphi_\alpha\}|}; \\ P(\neg\varphi_\tau|\neg\varphi_\alpha, t) &= \frac{|\{i \in [1, n] : t, i \not\models \varphi_\alpha \text{ and } t, i \not\models \varphi_\tau\}|}{|\{i \in [1, n] : t, i \not\models \varphi_\alpha\}|}. \end{aligned}$$

Tables 4–6 show the conditional probabilities of RCon ($\diamond_b \wedge \diamond_e$) $\square \rightarrow (\neg c \vee \diamond f)$ on trace $\langle a, b, c, d, f, c, e, c, h \rangle$. Notably, the conditional probability is not influenced by the total amount of events in the trace, but only by the events of interest.

To devise the probability of an RCon in an event log L (henceforth, *event log probability*), we have to detect the portion of the event log satisfying an LTL_{p_f} formula. To this end, we split the event log into a sub-log that has only the traces in which the activator occurs at least once (i.e., every $t \in L$ such that $P(\varphi_\alpha, t) > 0$), and the complementary sub-log consisting of the traces in which the activator does not occur (i.e., every $t \in L$ such that $P(\varphi_\alpha, t) = 0$). Given the above considerations and the definition of conditional probability for RCons in single traces (Definition 4.5), we devise a probabilistic interpretation for RCons over event logs as follows.

Definition 4.6 (Conditional Probability of LTL_{p_f} Formulae in an Event Log). Let φ_1 and φ_2 be two LTL_{p_f} formulae and L an event log of cardinality $|L|$. We say that φ_1 is *non-null* in a trace $t \in L$ if and only if $P(\varphi_1, t) > 0$. If $P(\varphi_1, t) = 0$, we say that φ_1 is *null* in t . The *conditional probability of φ_2 given φ_1 in L* is the portion of the event log that consists of traces for which φ_1 is non-null and satisfies φ_2 , given the satisfaction of φ_1 :

$$P(\varphi_2|\varphi_1, L) = \frac{\sum_{t \in L: P(\varphi_1, t) > 0} P(\varphi_2|\varphi_1, t)}{|L|}.$$

The *conditional probability of φ_2 given $\neg\varphi_1$ in L* is the portion of the event log that consists of traces for which φ_1 is null and satisfies φ_2 , given the satisfaction of $\neg\varphi_1$:

$$P(\varphi_2|\neg\varphi_1, L) = \frac{\sum_{t \in L: P(\varphi_1, t) = 0} P(\varphi_2|\neg\varphi_1, t)}{|L|}.$$

Table 10 shows the resulting contingency table. In the following, we provide the proof of the correctness of our approach.

Theorem 4.1 (Contingency of Event Log Conditionals). Given two LTL_{p_f} formulae φ_1 and φ_2 and an event log L of cardinality $|L|$, let $|L|_{P(\varphi_1) > 0}$ be the number of traces in which φ_1 is non-null and $|L|_{P(\varphi_1) = 0}$ the number of traces in which φ_1 is null. It follows that $P(\varphi_2|\varphi_1, L) + P(\neg\varphi_2|\varphi_1, L) + P(\varphi_2|\neg\varphi_1, L) + P(\neg\varphi_2|\neg\varphi_1, L) = 1$.

Proof. In light of the fact that there cannot be a trace where $P(\varphi_1)$ is both 0 and not 0 at the same time, the proof of Theorem 4.1

proceeds as follows.

$$P(\varphi_2|\varphi_1, L) + P(\neg\varphi_2|\varphi_1, L) + P(\varphi_2|\neg\varphi_1, L) + P(\neg\varphi_2|\neg\varphi_1, L) = 1 \quad (1)$$

$$\begin{aligned} & \frac{\sum_{t \in L: P(\varphi_1) > 0} P(\varphi_2|\varphi_1, t)}{|L|} + \frac{\sum_{t \in L: P(\varphi_1) > 0} P(\neg\varphi_2|\varphi_1, t)}{|L|} \\ & + \frac{\sum_{t \in L: P(\varphi_1) = 0} P(\varphi_2|\neg\varphi_1, t)}{|L|} + \frac{\sum_{t \in L: P(\varphi_1) = 0} P(\neg\varphi_2|\neg\varphi_1, t)}{|L|} = 1 \end{aligned} \quad (2)$$

$$\begin{aligned} & \sum_{t \in L: P(\varphi_1) > 0} P(\varphi_2|\varphi_1, t) + \sum_{t \in L: P(\varphi_1) > 0} P(\neg\varphi_2|\varphi_1, t) \\ & + \sum_{t \in L: P(\varphi_1) = 0} P(\varphi_2|\neg\varphi_1, t) + \sum_{t \in L: P(\varphi_1) = 0} P(\neg\varphi_2|\neg\varphi_1, t) = |L| \end{aligned} \quad (3)$$

$$\begin{aligned} & \sum_{t \in L: P(\varphi_1) > 0} (P(\varphi_2|\varphi_1, t) + P(\neg\varphi_2|\varphi_1, t)) \\ & + \sum_{t \in L: P(\varphi_1) = 0} (P(\varphi_2|\neg\varphi_1, t) + P(\neg\varphi_2|\neg\varphi_1, t)) = |L| \end{aligned} \quad (4)$$

$$\begin{aligned} & \sum_{t \in L: P(\varphi_1) > 0} \left(\frac{P(\varphi_1 \cap \varphi_2, t)}{P(\varphi_1, t)} + \frac{P(\varphi_1 \cap \neg\varphi_2, t)}{P(\varphi_1, t)} \right) \\ & + \sum_{t \in L: P(\varphi_1) = 0} \left(\frac{P(\neg\varphi_1 \cap \varphi_2, t)}{P(\neg\varphi_1, t)} + \frac{P(\neg\varphi_1 \cap \neg\varphi_2, t)}{P(\neg\varphi_1, t)} \right) = |L| \end{aligned} \quad (5)$$

$$\begin{aligned} & \sum_{t \in L: P(\varphi_1) > 0} \left(\frac{P(\varphi_1 \cap \varphi_2, t) + P(\varphi_1 \cap \neg\varphi_2, t)}{P(\varphi_1, t)} \right) \\ & + \sum_{t \in L: P(\varphi_1) = 0} \left(\frac{P(\neg\varphi_1 \cap \varphi_2, t) + P(\neg\varphi_1 \cap \neg\varphi_2, t)}{P(\neg\varphi_1, t)} \right) = |L| \end{aligned} \quad (6)$$

$$\sum_{t \in L: P(\varphi_1) > 0} \left(\frac{P(\varphi_1, t)}{P(\varphi_1, t)} \right) + \sum_{t \in L: P(\varphi_1) = 0} \left(\frac{P(\neg\varphi_1, t)}{P(\neg\varphi_1, t)} \right) = |L| \quad (7)$$

$$\sum_{t \in L: P(\varphi_1) > 0} 1 + \sum_{t \in L: P(\varphi_1) = 0} 1 = |L| \quad (8)$$

$$|L|_{P(\varphi_1) > 0} + |L|_{P(\varphi_1) = 0} = |L| \quad (9)$$

$$|L| = |L| \blacksquare \quad (10)$$

Probabilities defined as above permit the application of the association rule mining measures presented in Section 3.2 over an entire event log. In the light of the contingency table in Table 10, it suffices to map

$$P(A) \text{ to } P(\varphi_\alpha, L) = \frac{\sum_{t \in L: P(\varphi_\alpha, t) > 0} P(\varphi_\alpha, t)}{|L|},$$

$$\begin{aligned} P(B) \text{ to } P(\varphi_\tau, L) \\ = \frac{\sum_{t \in L: P(\varphi_\alpha, t) > 0} P(\varphi_\tau|\varphi_\alpha, t) + \sum_{t \in L: P(\varphi_\alpha, t) = 0} P(\varphi_\tau|\neg\varphi_\alpha, t)}{|L|}, \text{ and} \end{aligned}$$

$$P(AB) \text{ to } P(\varphi_\tau|\varphi_\alpha, L) = \frac{\sum_{t \in L: P(\varphi_\alpha, t) > 0} P(\varphi_\tau|\varphi_\alpha, t)}{|L|}.$$

We remark that the non-trivial mapping from $P(AB)$ to $P(\varphi_\tau|\varphi_\alpha, L)$ is intuitively rooted into the inherent nature of “IF A THEN B” rules such as the RCons, as evidenced in [56], and its soundness is evidenced by Theorem 4.1. For example, Table 11 shows a few measures computed from the probabilities of the RCon ($\diamond_b \wedge \diamond_e$) $\square \rightarrow (\neg c \vee \diamond f)$ over a log composed of five traces. We remark that this result is distinct from the mere aggregation of trace measures. For example, comparing the average of the Support values in Table 9 (0.44) and the Support value presented in Table 11 (0.67), we observe that the former is the average

Table 10
Contingency table of conditional event log probabilities.

	$P(\varphi_\alpha, t) > 0$	$P(\varphi_\alpha, t) = 0$	
φ_τ	$\frac{\sum_{t \in L} P(\varphi_\tau \varphi_\alpha, t)}{ L }$	$\frac{\sum_{t \in L} P(\varphi_\tau \neg \varphi_\alpha, t)}{ L }$	$\frac{\sum_{t \in L: P(\varphi_\alpha, t) > 0} P(\varphi_\tau \varphi_\alpha, t) + \sum_{t \in L: P(\varphi_\alpha, t) = 0} P(\varphi_\tau \neg \varphi_\alpha, t)}{ L }$
$\neg \varphi_\tau$	$\frac{\sum_{t \in L} P(\neg \varphi_\tau \varphi_\alpha, t)}{ L }$	$\frac{\sum_{t \in L} P(\neg \varphi_\tau \neg \varphi_\alpha, t)}{ L }$	$\frac{\sum_{t \in L: P(\varphi_\alpha, t) > 0} P(\neg \varphi_\tau \varphi_\alpha, t) + \sum_{t \in L: P(\varphi_\alpha, t) = 0} P(\neg \varphi_\tau \neg \varphi_\alpha, t)}{ L }$
	$\frac{\sum_{t \in L} P(\varphi_\alpha, t)}{ L }$	$\frac{\sum_{t \in L} P(\neg \varphi_\alpha, t)}{ L }$	1

Table 11
Event log probabilities and measures of a sample of measures for the RCon $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$.

Event log	$P(\varphi_\alpha, t)$	$P(\varphi_\tau \varphi_\alpha, t)$	$P(\neg \varphi_\tau \varphi_\alpha, t)$	$P(\varphi_\tau \neg \varphi_\alpha, t)$	$P(\neg \varphi_\tau \neg \varphi_\alpha, t)$
$t_1 = \langle a, b, c, d, f, c, e, c, h \rangle$	>0	0.83	0.17	0.67	0.33
$t_2 = \langle b, d, a, f, g, d, e, d \rangle$	>0	1.00	0.00	1.00	0.00
$t_3 = \langle a, c, d, b, c, e, f, c \rangle$	>0	1.00	0.00	0.80	0.20
$t_4 = \langle b, c, c, e, a \rangle$	>0	0.50	0.50	1.00	0.00
$t_5 = \langle b, c, d, e \rangle$	=0	NaN	NaN	0.75	0.25
$P(\varphi_\alpha, L) = 0.80$		$P(\varphi_\tau \varphi_\alpha, L) = 0.67$		$P(\varphi_\tau \neg \varphi_\alpha, L) = 0.13$	
$P(\varphi_\tau, L) = 0.82$		$P(\neg \varphi_\tau \varphi_\alpha, L) = 0.15$		$P(\neg \varphi_\tau \neg \varphi_\alpha, L) = 0.05$	
Support: 0.67	Confidence: 0.83	Specificity: 0.25	Lift: 1.02		

proportion of events in a trace satisfying both the target and the activator, while the latter represents the proportion of traces of the event log satisfying the RCon.

4.4. Measurement system

Given an event log L , a set of RCons R , and a set of probabilistic measures M as input, our framework returns the measurement of every measure in M for each constraint in R both over every single trace $t \in L$ and over the entire event log L . More specifically, the output can be reported at three different levels of detail:

Event level: distinct evaluation of φ_α and φ_τ of each constraint in R on every event of every trace in L ;

Trace level: measurement of each measure in M for each constraint in R for every trace in L ;

Aggregated view: descriptive statistics over the event log of all the trace-level measures.

Event log level: measurement of every measure in M for every constraint in R for the entire event log L .

For example, Table 9 shows some trace level measures together with their descriptive statistics and Table 11 shows the corresponding event log level measures for the RCon $(\diamond b \wedge \diamond e) \square \rightarrow (\neg c \vee \diamond f)$. Since being able to perceive the overall status of a constraint is as important as the possibility to analyze its details in single traces, we report the entire statistical distribution of a measure across the event log to provide a complete information spectrum. Fig. 1 depicts the pipeline of the framework from the input to the output. In the first stage, an RCon is evaluated on each trace of the event log. Then, the evaluation result is used to compute the probabilities of the rule. On top of them, the measures of the rule in each trace and in the entire event log are computed. Also, descriptive statistics over the event log are reported for each trace measure.

We remark that the design of the RCons is crucial for the evaluation and the computation of the measures especially in terms of definition of their activator. Let us take as an example the RESPONDEDEXISTENCE(a, b) constraint from the repertoire of DECLARE (see Table 1). The classical LTL_f formula underlying RESPONDEDEXISTENCE(a, b) for whole-trace evaluations is

$\neg \diamond a \vee \diamond b$ [28]. However, the formulation of the rule as an RCon can lead to different interpretations:

- $a \square \rightarrow (\diamond b \vee \diamond b)$: If a occurs, b is expected to occur somewhere in the trace;
- $(\diamond a \vee \diamond a) \square \rightarrow (\diamond b \vee \diamond b)$: For every event in the trace such that a occurs either in the past or in the future, also b should occur somewhere in the trace;
- $True \square \rightarrow \neg(\diamond a \vee \diamond a) \vee (\diamond b \vee \diamond b)$: For every event in the trace, if a occurs in the trace, also b is expected to occur;
- $t_{\text{Start}} \square \rightarrow (\neg \diamond a \vee \diamond b)$: At the beginning of the trace, if a occurs in the trace, also b should occur.

All the formulations above are legitimate as they entail that the occurrence of a in the trace demands the occurrence of b . However, the difference in the way the activator is represented turns out to be crucial. The activator, indeed, encodes when the rule is of interest. For example: Are we interested in each occurrence of task a or only in its eventual occurrence in the trace? Do we want the rule to be satisfied in every point of the trace or just at the beginning of the trace? These choices have a clear impact on the measures. Table 12 presents the evaluation of a trace with the different formulations seen above and their trace measurements for Confidence and Support. Confidence is equal to 1 as each time the activator holds true, also the target holds true. Support (i.e., the frequency of $\varphi_\alpha \cap \varphi_\tau$) varies considerably instead. Notice that this phenomenon comes with neither a good nor with a bad connotation, but stresses the idea that a full control over the formula implies a mindful decision about its design and subsequently on picking the right measures for it. In Table 1, we devised the RCon formulae based on the activators of DECLARE templates described in [3,20] – hence, e.g., the choice for RESPONDEDEXISTENCE of the first option presented in Table 12. While DECLARE templates are reasonably simple and well-known standard cases, encoding the right activator is crucial for the design of custom RCons and their measures. Lastly, we would like to remark that all the variants in Table 12 take exactly the same amount of computational time to be checked as any formula requires a trace to be read only twice, as described in Section 4.1.

In summary, we have described in this section a novel measurement framework for reactive temporal specifications based on LTL_p_f, supporting probabilistic interestingness measures at

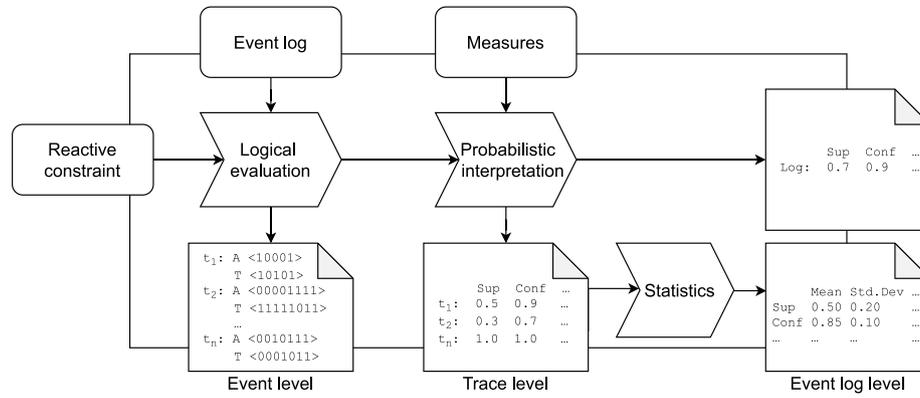


Fig. 1. Measurement framework architecture.

Table 12

Measurements of a constraint expressed with different formulations on trace (d, a, b, c, a) .

RCon formulation	Evaluation	Support $P(\varphi_\alpha \varphi_\tau, t)$	Confidence $P(\varphi_\tau \varphi_\alpha, t)$
$a \Box \rightarrow (\diamond b \vee \diamond c)$	$\varphi_\alpha: \langle 0, 1, 0, 0, 1 \rangle$ $\varphi_\tau: \langle 1, 1, 1, 1, 1 \rangle$	$2/5 = 0.4$	$2/2 = 1$
$(\diamond a \vee \diamond a) \Box \rightarrow (\diamond b \vee \diamond b)$	$\varphi_\alpha: \langle 1, 1, 1, 1, 1 \rangle$ $\varphi_\tau: \langle 1, 1, 1, 1, 1 \rangle$	$5/5 = 1$	$5/5 = 1$
$True \Box \rightarrow \neg(\diamond a \vee \diamond a) \vee (\diamond b \vee \diamond b)$	$\varphi_\alpha: \langle 1, 1, 1, 1, 1 \rangle$ $\varphi_\tau: \langle 1, 1, 1, 1, 1 \rangle$	$5/5 = 1$	$5/5 = 1$
$t_{start} \Box \rightarrow (\neg \diamond a \vee \diamond b)$	$\varphi_\alpha: \langle 1, 0, 0, 0, 0 \rangle$ $\varphi_\tau: \langle 1, 1, 1, 0, 0 \rangle$	$1/5 = 0.2$	$1/1 = 1$

both trace and event log levels. The framework is designed to be suitable for any custom formula in the form of a Reactive Constraint, and any measure that is based on the probability of the activator and target of the constraints. Therefore, it supports template sets like DECLARE and all the interestingness measures from association rule mining seen in Table 3, though not being limited to them. Next, we evaluate our approach through tests conducted with our implemented prototype.

5. Implementation and performance analysis

We have implemented our measurement framework as a proof-of-concept software prototype built upon the existing declarative process specification processor tool Janus [14,16]. The Java source-code can be found at github.com/Oneiroe/Janus. The core component of the software is the RCons verification engine, upon which are build independently a declarative process discovery module and the present declarative rules measurement module. All the process specifications used in the following experiments are discovered with this discovery module implementing the technique presented in [14]. In the remainder of this section, we first report on the results of a time and space analysis with simulated data. Then, we investigate the computational performance on real-world event log datasets. The results demonstrate the practical feasibility and applicability of our approach.

5.1. Time analysis

To assess the efficiency of our implemented technique, we measure its time performance against an increase in the data size (i.e., the cardinality of the event log and the length of its traces) and the model size (i.e., the number of rules) with synthetic event logs. We repeated every experiment 10 times to smooth random

Table 13

The set of DECLARE rules used in the experiments.

INIT(a)	RESPONSE(e, f)	CHAINRESPONSE(o, p)
END(b)	PRECEDENCE(g, h)	CHAINPRECEDENCE(q, r)
ATMOSTONE(c)	ALTERNATEPRECEDENCE(i, l)	RESPONDEXISTENCE(s, t)
PARTICIPATION(d)	ALTERNATERESPONSE(m, n)	COEXISTENCE(u, v)
SUCCESION(w, x)	ALTERNATESUCCESION(y, z)	CHAINSUCCESION(j, k)
NOTCOEXISTENCE(0, 1)		

factors. The reported results average over the ones of the single repetitions. The machine used for the experiments was equipped with an Intel Core i5-7300U CPU at 2.60 GHz, quad-core, 16 Gb of RAM and an Ubuntu 18.04 LTS operating system.

To test the response of our implemented framework against the input data size, we set up a controlled experiment in which we first generate logs of varying sizes that are compliant with a fixed set of rules, resorting to the simulation engine of MINERful [57]. Thereupon, we compute the measures listed in Table 3 against all the rules of a larger test specification (not fully compliant with the event log). For every run, we recorded the wall-clock time of our prototype.

The starting set of rules stems from the DECLARE repertoire of templates [5] and is provided in Table 13. Notice that the set contains all the rule templates seen in Table 1 and is designed in a way that every constraint insists on different tasks.

The test model consists of 649 constraints extracted by the discovery algorithm of Janus (setting the Support and Confidence threshold parameters to 0.05 and 0.8) from a synthetic event log of 834963 events, 500 traces and tasks in $\{a, \dots, z, 0, 1\}$ that is compliant with the initial model.³ Given the test specification obtained as described above, we performed two tests of 65 iterations (1) increasing the length of the traces (with a step of 100 events per iteration, keeping the number of traces per event log equal to 500), and (2) increasing the number of traces in the event log (with a step of 50 new traces per iteration, keeping the trace lengths between 900 and 1000 events). Fig. 2 illustrates the results of both experiments. We observe that the factor actually influencing the wall-clock time is the total amount of events rather than the trace length: indeed, Fig. 2 shows that the recorded timings of both experiments tend to lie on the same line. This experimental result confirms the linear relation between the total number of events in the log and the computational performance illustrated in Section 4.1.

Next, we investigate the response of the framework to an increase in the model size. To do so, we first generate an event log containing 1000 traces with a trace length between 100 and 500

³ Available at <https://oneiroe.github.io/DeclarativeMeasurements-static>.

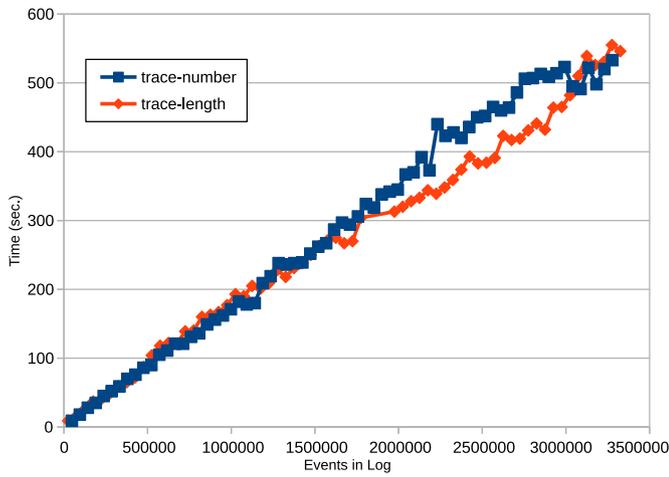


Fig. 2. The computation time is linearly dependent on the total number of events in the event log.

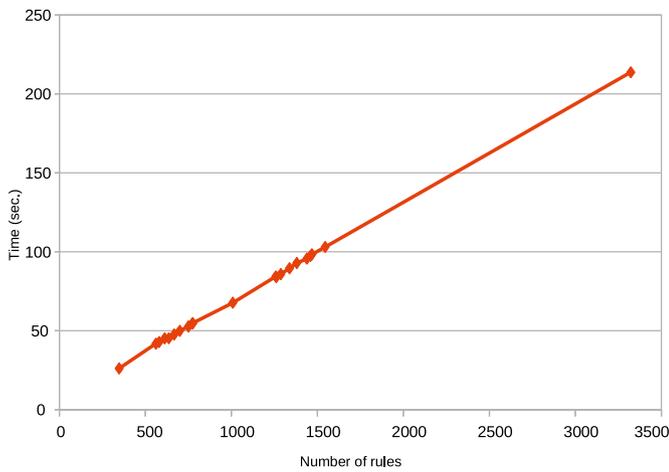


Fig. 3. The computation time is linearly dependent on the total number of rules to check.

events from the simulation of the rules in Table 13. Thereupon, we use the discovery algorithm of Janus to automatically retrieve different test models with varying levels of compliance. To that extent, we make the Confidence threshold range from 1.0 (full model compliance), down to 0.0 with a step of 0.05. The rationale is, the lower the Confidence threshold, the higher the number of constraints in the test model. Then, we calculate all the measures in Table 3 for every constraint of each test model. The time taken for the measurements are shown in Fig. 3. Notice that the computation time is linearly dependent on the number of rules to check, thus in line with the theoretical computational cost exposed in Section 4.1.

5.2. Space analysis

The space consumption of our technique depends on the data structures required to store the multi-level results. More specifically, four multidimensional matrices are used, containing respectively (1) the evaluation at the event level, (2) the measures at the trace level, (3) their descriptive statistics over the log, and finally (4) the measures at the event log level.

Considering $|L|$ the number of traces in the log, $|E|$ the total number of events in a log, $|R|$ the number of constraints, $|M|$ the number of measures, the sizes of the matrices are respectively the following:

- (1) Events evaluation: $|E| \times |R|$, wherein each cell contains two boolean values (i.e., the evaluation of the activator and target of the constraint);
- (2) Trace measures: $|L| \times |R| \times |M|$, having a real number in every cell;
- (3) Trace measures statistics: $|R| \times |M|$, containing seven real numbers per cell (for the mean, geometric mean, variance, population variance, standard deviation, maximum value, and minimum value, respectively);
- (4) Event log measures: $|R| \times |M|$, with a real number each.

The events matrix is optimized as a bit matrix, where two bits are sufficient to store the boolean results of the evaluation of both the activator and target for one event. We implemented our framework in Java, so we employ 1-byte `Byte` objects and 4-byte `Float` numbers (6 decimal digits are sufficiently accurate for our purpose). Taking these indicators into account, we can estimate the space consumption. For example, assuming that $|L| = 1000$, the maximum number of events in a trace n is 50, $|R| = 100$, $|M| = 30$, the expectation for the space demands are distributed as follows:

- (1) Events evaluation: $1000 \times 20 \times 260 \times 1 = 5\,200\,000$ bit = 5.2 Mb;
- (2) Trace measures: $1000 \times 260 \times 37 \times 4 = 38\,480\,000$ bit = 38.48 Mb;
- (3) Trace measures statistics: $260 \times 37 \times 7 \times 4 = 269\,360$ bit = 269.36 Kb;
- (4) Event log measures: $260 \times 37 \times 4 = 38\,480$ bit = 38.48 Kb.

Therefore, the most memory-demanding data structures are those that pertain to the events evaluation and the trace measurements matrices. The former is bigger than the latter only if the average number of events per trace is greater than 4 times the number of measures used, i.e., $\frac{|E|}{|L|} > 4 \times |M|$. In our experiments, even using all the 37 measures of Table 3, this has not occurred.

As with the experiments for the evaluation of time, we analyze empirically the space consumption through simulations, controlling the number of events n per trace, the number of traces $|L|$, and the number of constraints under analysis $|M|$. To measure the memory consumed by the data structures, we perform a Bit serialization of the matrix objects listed above. This allows us to have a precise measure of the space consumed by every object, though it unavoidably requires that the available memory is twice as much as the strictly necessary amount.

Fig. 4 illustrates the results of our experiments. As it can be seen, the resulting linear trends are in line with the expectations, modulo the constant factors introduced by the Java Virtual Machine objects. It can be noticed that the number of constraints to check, being a common factor among all the objects, increases the overall required memory quicker than the other parameters. We remark that depending on the desired outcome, not all the measures nor all the matrices are necessary. For example, if the log measures are desired, the trace measures and their statistics can be ignored and vice-versa. The events evaluation is the only mandatory object upon which all other computations are based.

At present, our implementation works in-memory, thus it is assumed that all the objects fit in main memory (which proved to be sufficient for all the real life log under analysis). However, we would remark that every measure calculation (both at log and trace levels) is independent from the other one, thus it is possible to either (i) compute one measure or constraint at the time in pipeline, in order to reduce the memory load, or (ii) to distribute the workload, by making each system compute independently one measure per constraint. Both are interesting directions for the future upgrades of our implementation.

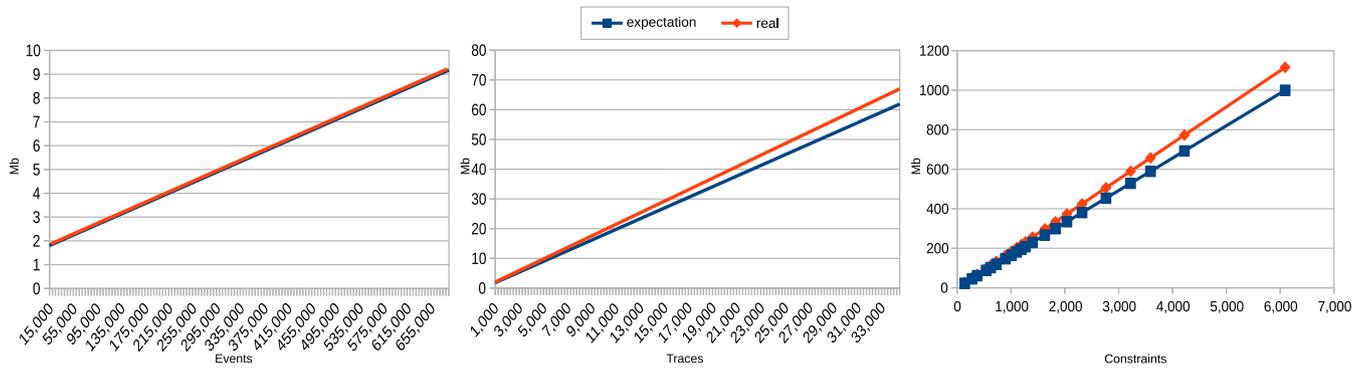


Fig. 4. The space consumption is linearly dependent on the total number of events in the log (left), number of traces in the log (center), and number of rules to check (right).

5.3. Application on real-world event logs

To test the performance also in real settings, we compute the rule measures on 13 event logs, whose characteristics are exposed in Table 14. Twelve of those event logs are openly available⁴ and belong to the Business Process Intelligence Challenge (BPIC) collection, a Road-Traffic Fines Management Process (RTFMP) and the aforementioned Sepsis event log. In addition, we analyze the performance of our prototype on an event log stemming from a partner of a smart-city project in which the authors are involved (labeled as “Smart city” in Table 14). We included the Smart city event log due to its considerable size: as it can be noticed from the table, it is the one bearing the largest amount of events in this experiment.

For each log, we ran the discovery algorithm of Janus [14] in order to extract a test model to check the event log against. We tuned the parameters of the discovery algorithm to obtain a set of rules to which the event log complies for the most part (Confidence threshold of 0.8), even though the constrained tasks are possibly infrequently co-occurring (Support threshold of 0.05). Table 14 illustrates the results. For each event log we report, along with the number of traces, the occurring tasks, the events, the number of constraints in the test model, the total time from the launch to the termination of the software (“Time”), the time to evaluate the rules on events (“Checks”), the time to compute the measures both at the trace and log levels (“Measures”), and the total space consumed by the data structures of our tool (“Space”) against their expected value (“Expectation”). We remark that the space consumption is consistent with our theoretical expectation and that the wall-clock time remains within acceptable ranges as the slowest run takes around 4.5 min to check about 600 constraints in a considerably big event log such as BPIC17 [58] handling around 2.5 Gb of data.

5.4. Analysis of custom rules

In order to demonstrate that our framework can handle any Reactive Constraints, beyond the standard DECLARE repertoire, we applied our approach to compute the discussed measures of a custom rule on the Sepsis event log. We name the custom rule `BIDIRECTIONALTIMECONSEQUENT(a, b, c)` as its RCon formulation is $a \square \rightarrow \diamond b \wedge \diamond c$. It states that if a occurs, it is expected that either c occurred before it or b will occur afterwards. Table 15 reports the measures at log and trace level for `BIDIRECTIONALTIMECONSEQUENT(Admission NC, CRP, IV Liquid)` calculated on the Sepsis real-life log [34]. As it can be noticed, Confidence and Recall are relatively high (0.82 and 0.79, respectively)

and the values of Coverage and Prevalence (0.76 and 0.79, respectively) suggest a frequent occurrence of activator and target. The value of Lift is greater than 1, which denotes dependency between activator and target (especially at a trace level). The detailed results of the evaluation on each trace can be found at oneiroe.github.io/DeclarativeMeasurements-static.

The capability of our framework to handle non-standard rules opens up new possibilities for the claimed extensibility of DECLARE as a declarative specification language, claimed from its very inception to be open to customization through the definition of new rules according to the process analyst needs [65].

6. Evaluation

In this section, we report on experiments that show interesting implications of having a vast availability of measures with customization options. Specifically, Section 6.1 investigates over which measures can be of interest in the scope of declarative specification discovery, and Section 6.2 shows how the properties of measures can be exploited to characterize the alterations of constraints when the underlying process changes.

All the experimental data (code, input data, results) can be found at <https://oneiroe.github.io/DeclarativeMeasurements-static>. In the following experiments, we resort to the following tools: (i) The Janus discovery algorithm [14] for the discovery of declarative models from events logs; (ii) The simulation engine of MINERful [57] for the generation of event logs complying with given declarative specifications; (iii) The error injection engine of MINERful [66] for the controlled insertion of noise into event logs; (iv) The declarative model simplification technique of MINERful [33] for the removal of redundancies from declarative specifications.

6.1. Ranking experiment

The objectives of this experiment are the following: (1) Empirically showing that relying on more measures than the sole Support and Confidence measures is effective to characterize process rules in an event log; (2) Highlighting insightful measures in a declarative process mining context.

To achieve both objectives we rank the measures according to how many correct and interesting rules they are able to recognize. We take inspiration from the *correct rules at N* experiment introduced by the seminal work of Le and Lo [6]. Given an event log L , a ground truth set of rules R_G satisfied in L , a set of rules $R_D \supset R_G$ containing also loosely satisfied rules in L , the set of interestingness measures M of Table 3, and a predefined threshold N , we compute the value of each measure $m \in M$ for all the rules $r \in R_D$ on L . Then, for every measure $m \in M$ we create sets of rules that are associated to a common value and

⁴ <https://data.4tu.nl/>.

Table 14
Performance records on real-life datasets.

Event log	Traces	Tasks	Events	Rules	Time [s]	Checks [ms]	Measures [ms]	Space [Mb]	Expectation [Mb]
BPIC12 [59]	13087	36	262200	519	106.7	19237.1	83160.6	1222.76	1141.93
BPIC13_cp [60]	1487	7	6660	20	1.2	221.2	338.3	5.10	4.56
BPIC13_i [60]	7554	13	65533	14	3.6	584.9	1064.3	18.09	16.59
BPIC14_f [61]	41353	9	369485	51	29	3977.1	21277.4	357.10	331.04
BPIC15_1f [62]	902	70	21656	3856	42.2	13686	25974.4	670.45	602.83
BPIC15_2f [62]	681	82	24678	5889	57.8	19790.6	35029.5	832.38	745.84
BPIC15_3f [62]	1369	62	43786	4098	81.2	26468	51764.6	1105.65	1014.59
BPIC15_4f [62]	860	65	29403	4690	60.4	18997.8	38492.4	818.55	740.39
BPIC15_5f [62]	975	74	30030	5164	69.7	22947.1	43684.2	999.51	906.35
BPIC17_f [58]	21861	41	714198	611	275.9	58909.8	207182.2	2561.93	2413.94
RTFMP [63]	150370	11	561470	49	80.1	7887.5	65214	1210.04	1118.05
Sepsis [34]	1050	16	15214	260	3.8	765.8	1842.6	49.80	44.67
Smart city [64]	4347	20	692333	292	61.5	23261.8	32301.8	398.74	390.37

Table 15
Measures resulting from the evaluation of constraint BIDIRECTIONALTIMECONSEQUENT(Admission NC, CRP, IV Liquid) on the Sepsis event log [34].

Measure	Log	Trace	Measure	Log	Trace	Measure	Log	Trace
Support	0.62	0.06	Prevalence	0.79	0.66	One-way Support	0.04	0.51
Confidence	0.82	0.82	Added Value	0.03	0.18	Two-way Support	0.03	0.05
Recall	0.79	0.09	Relative Risk	1.17	1.36	Two-Way Support	0.01	0.02
Lovinger	-0.16	0.05	Jaccard	0.67	0.09	Variation		
Specificity	0.30	0.36	Ylue Q	0.32	0.63	Linear Correlation Coefficient	0.13	0.13
Accuracy	0.70	0.39	Ylue Y	0.17	0.63	Piatetsky-Shapiro	0.02	0.01
Lift	1.04	1.24	Klogsen	0.02	0.07	Cosine	0.80	0.30
Leverage	0.22	0.76	Conviction	1.16	Infinity	Information Gain	0.04	0.30
Compliance	0.86	0.98	Interestingness	0.03	0.01	Sebag-Schoenauer	4.56	Infinity
Odds Ratio	1.95	Infinity	Weighting Dependency			Least Contradiction	0.62	0.00
Gini Index	0.01	0.03	Collective Strength	6.47	Infinity	Odd Multiplier	1.20	Infinity
Certainty factor	0.14	0.67	Laplace Correction	0.82	0.60	Example and Counterexample Rate	0.78	0.81
Coverage	0.76	0.07	J Measure	0.00	0.01	Zhang	0.17	0.63

sort those sets accordingly. This leads to a separate sorting of the rules for every measure. If, for instance, rules r_1 and r_2 have a Confidence of 1.0, rule r_3 has a Confidence of 0.9, and rules r_4 , r_5 and r_6 have a Confidence of 0.8, the top- N sets are: $\{r_1, r_2\}$ in the top-1, $\{r_1, r_2, r_3\}$ in the top-2, and $r_1, r_2, r_3, r_4, r_5, r_6$ in the top-3 for Confidence. Intuitively, a good measure should assign high scores to correct rules. Therefore, we finally count how many of the rules in R_C are within the top- N sets. We repeated the experiment 10 times and considered the average of the results to avoid fluctuations caused by the random factors of simulation. We performed the experiment with N set to 1, 5, 10, 25, 50, 100, 200, 500, 1000, and 1500, i.e., ranging from considering only the best-scoring rules to considering all the rules in R_D . The final ranking of a measure is computed as the average of its ranking for each N . Table 16(a) shows the final rankings for this experiment. Together with the ranking, for each measure we report also the number of correct rules (“Correct” column) and the average ratio of correct rules over the total number (“Ratio” column) in the top- N sets, averaged over the 10 repetitions of the experiment.

We run our experiments with three different event logs: (i) a simulation of a synthetic process specification (Section 6.1.1); (ii) a simulation of a synthetic process specification with random changes in the event log so as to mimic partial non-compliance (Section 6.1.2); (iii) a real-life event log (Section 6.1.3). At the end of this section, we draw some conclusions from the obtained results.

6.1.1. Process simulation

More specifically, we simulate the specification described in Table 13. Notice that the rules are designed to not interfere with one another and each of them constrains different tasks. The simulation produces an event log that is fully compliant with the rules in Table 13. From the simulated event log, we discover a new process specification with loose bounds (Support and Confidence thresholds set to 0.05 and 0.5, respectively) in order to discover also infrequent and seldom violated rules. We

simplify the resulting set of rules by removing those that do not match yet strictly subsume or are entailed by the ground-truth rules, in order to avoid misleading results – for example, if $\text{CHAINRESPONSE}(o, p)$ and $\text{RESPONSE}(o, p)$ both belong to the returned set of rules, the former is retained and the latter is removed because $\text{CHAINRESPONSE}(o, p)$ is part of the ground-truth specification (see Table 13) and is subsumed by $\text{RESPONSE}(o, p)$. The full detail of the technique that deals with the removal of redundant rules is out of scope for this paper. The interested reader can find a detailed description of the problem and the approach in [33]. The simplified discovered model consists of 1310 rules on average. Thereupon, we apply our measurement framework to compute the measures in Table 3 at the event log level for all the discovered rules. Finally, we sort the rules according to each measure, and rank the measures according to how many of the original rules are among in top N sets.

Notice that Support ranks only sixteenth. Confidence, by contrast, is at the top of the ranking. It should be observed, however, that the experiment considers by design never violated rules (i.e., those with maximal Confidence), hence the top position of this measure. Nevertheless, there are two measures that match Confidence, namely (i) Example and Counterexample Rate and (ii) the Sebag-Schoenauer measure. This is in accordance with P5, as the absence of counterexamples for the rules makes measures with known maximal values highly rank the original, correct rules. The “Ratio” column reported in Table 16(a) helps to distinguish the accuracy of the results. For example, Odd Multiplier and Odd Ratio have the same rank (i.e., the same amount of correct rules identified), although the Odd Ratio returns 3 times more rules than the Odd Multiplier within the top- N sets.

6.1.2. Process simulation with noise

The previous experiment tests a perfectly compliant setup where the reference rules are never violated. For this reason, we conduct a modified version of the experiment, this time injecting noise in the event log in order to check if the ranking is preserved

Table 16

Ranking of measures according to the simulation experiment with a fully compliant simulated event log (left), with an altered event log (center), and with a real-life event log (right).

(a) Simulation				(b) Simulation with noise				(c) Real-world based			
Rank	Measure	Correct	Ratio	Rank	Measure	Correct	Ratio	Rank	Measure	Correct	Ratio
1	Confidence	11.00	24.79%	1	Confidence	5.58	6.21%	1	Recall	77.67	14.64%
1	Example and Counterexample Rate	11.00	24.79%	1	Example and Counterexample Rate	5.58	6.21%	2	Confidence	76.67	51.00%
1	Sebag–Schoenauer	11.00	24.81%	3	Laplace Correction	5.53	6.72%	2	Example and Counterexample Rate	76.67	51.00%
4	Laplace Correction	8.65	17.70%	4	Odd Multiplier	5.42	11.16%	4	Sebag–Schoenauer	74.78	50.39%
5	Accuracy	6.62	10.69%	5	Conviction	5.43	11.20%	5	Least	70.56	56.43%
6	Least	6.35	10.55%	6	Lift	4.58	11.38%	6	Contradiction		
7	Contradiction			7	Prevalence	4.42	4.31%	6	Cosine	70.44	56.40%
7	Cosine	6.24	10.51%	8	Accuracy	4.61	4.06%	6	Jaccard	70.44	56.40%
8	Jaccard	6.23	10.51%	9	Least	4.31	3.98%	8	Odds Ratio	72.33	19.88%
8	Prevalence	6.15	10.17%	8	Contradiction			8	Ylue Q	72.33	19.88%
10	Conviction	7.00	24.90%	10	Recall	4.20	0.51%	8	Ylue Y	72.33	19.88%
10	Odd Multiplier	7.00	22.24%	11	Jaccard	4.22	3.97%	11	Accuracy	70.11	56.03%
10	Odds Ratio	7.00	7.57%	12	Cosine	4.18	3.97%	12	Laplace Correction	42.89	31.07%
10	Ylue Q	7.00	7.57%	13	One-way Support	4.09	3.12%	12	Relative Risk	71.33	20.46%
10	Ylue Y	7.00	7.57%	14	Certainty factor	4.50	2.95%	12	Specificity	71.33	20.46%
15	Support	5.18	6.15%	15	Added Value	4.01	2.21%	15	Conviction	65.33	54.59%
16	Lift	4.94	9.55%	16	Information Gain	4.05	3.40%	16	Odd Multiplier	65.00	54.50%
17	Certainty factor	4.48	3.66%	17	Support	3.85	3.88%	17	Certainty factor	52.56	36.07%
18	One-way Support	4.39	3.24%	18	Zhang	4.44	2.77%	18	Linear Correlation Coefficient	41.00	37.66%
18	Zhang	4.35	4.15%	19	Two-way Support	4.05	1.99%	19	Gini Index	34.78	39.72%
20	Two-way Support	4.26	1.79%	20	Klogsen	4.02	2.02%	19	One-way Support	30.44	27.55%
21	Interestingness Weighting	4.27	1.79%	21	Interestingness Weighting	4.04	1.80%	19	Interestingness Weighting	31.67	18.93%
22	Added Value	4.25	1.95%	22	Piatetsky–Shapiro	3.94	1.73%	21	Dependency		
22	Klogsen	4.24	2.25%	23	Sebag–Schoenauer	3.30	3.19%	22	Information Gain	29.78	25.97%
24	Piatetsky–Shapiro	4.19	1.65%	24	Coverage	3.28	0.88%	22	Piatetsky–Shapiro	30.78	19.64%
25	Recall	4.64	3.69%	25	Ylue Q	2.99	0.46%	24	Zhang	38.33	27.05%
26	Information Gain	4.33	3.91%	25	Ylue Y	2.99	0.46%	25	Added Value	29.67	17.07%
27	Linear Correlation Coefficient	3.92	2.44%	27	Leverage	3.28	0.97%	25	Leverage	29.78	19.62%
28	Leverage	3.40	1.20%	28	Two-Way Support Variation	3.96	2.70%	27	Prevalence	31.56	26.90%
29	Coverage	3.24	0.81%	29	Lovinger	3.87	0.93%	28	Two-way Support	30.22	16.00%
30	Specificity	2.68	0.38%	30	Linear Correlation Coefficient	3.58	1.38%	29	Lift	29.44	25.21%
31	Gini Index	3.01	0.84%	31	Gini Index	3.00	0.83%	30	Klogsen	29.56	14.75%
31	Lovinger	2.23	0.32%	31	J Measure	2.77	0.78%	31	Support	28.22	19.04%
33	Relative Risk	2.48	0.33%	32	J Measure	2.77	0.78%	32	Coverage	19.89	4.68%
34	Collective Strength	2.02	0.24%	33	Specificity	2.14	0.28%	33	Lovinger	27.00	7.44%
35	J Measure	0.00	0.00%	34	Odds Ratio	1.94	0.24%	34	Collective Strength	20.78	5.81%
35	Two-Way Support Variation	0.00	0.00%	35	Collective Strength	1.98	0.23%	35	J Measure	7.22	4.10%
				36	Relative Risk	2.06	0.25%	36	Two-Way Support Variation	2.11	1.48%

in non-optimal situations. Specifically, we randomly delete or duplicate 5% of the occurrences of every task. The results can be found in Table 16(b). It can be seen that also in the presence of partial non-compliance, the measures of Confidence and Example and Counterexample Rate continue to be on top of the ranking, while Sebag–Schoenauer measure drops in the second half of the list, together with Ylue Q, Ylue Y and Odds Ratio. This sudden change due to noise is motivated by the fact that these measures are convex, as we discussed in regard with P6, so they rapidly decrease in presence of counterexamples. Support persists in the middle of the ranking, with 17 measures scoring better than it. Lift and Information Gain perform by far better in presence of noise. Understanding how measures react to changes in the data (noise in this case), is key for process drift analysis [67], where the evolution of the process is the main objective. We study more in details the influence of noise on the measurements in Section 6.2.

6.1.3. Real-life event log

We replicate our ranking experiment based on a real-life event log, specifically the Sepsis data-set [34]. In [35], Mannhardt and Blinde illustrate a procedural model discovered with the help of domain experts representing the sepsis treatment process at a hospital. We manually translated the model into DECLARE rules

that we use as a ground-truth specification to replicate the ranking experiment. The model consists of 100 rules, listed in Table 17. Because it is a real life event log, we do not inject any noise. Table 16(c) illustrates the results. It can be seen that the measures of Confidence and Example and Counterexample Rate remain at the top of the ranking. Also, while Recall conquers the first position for the amount of correct rules reported, it also returns a high number of other incorrect rules as signaled by its low average ratio score. In this case, Support drops at the bottom of the ranking.

These experiments suggest possible candidate measures to be used for declarative process discovery. We tested different scenarios, thereby showing how the measures' scores vary in each of them and how the evaluation of rules can benefit from the perspectives of multiple measures that were not previously available for temporal specifications. However, we remark that some measures should be handled in a specific manner. For example, Coverage is a descriptive measure reporting on the portion of the event log in which the rule activator occurs: this characteristic determines its low ranking in all the previous experiments. A low score in these experiments does not imply that a measure is of scarce use in general, though. Indeed, we can only highlight the average top ranking measures of these experiments as excellent options, as we do in the next section while depicting the current

Table 17
 DECLARE model of SEPSIS data-set based on [35].

INIT(ER Registration)	ATMOSTONE(ER Registration)	ATMOSTONE(ER Sepsis Triage)
ATMOSTONE(ER Triage)	SUCCESSION(ER Registration, ER Triage)	SUCCESSION(ER Triage, ER Sepsis Triage)
PRECEDENCE(ER Registration, CRP)	PRECEDENCE(ER Registration, Leucocytes)	PRECEDENCE(ER Registration, LacticAcid)
ATMOSTONE(IV Liquid)	ATMOSTONE(IV Antibiotics)	PRECEDENCE(ER Sepsis Triage, IV Liquid)
PRECEDENCE(ER Sepsis Triage, IV Antibiotics)	COEXISTENCE(IV Liquid, IV Antibiotics)	PRECEDENCE(ER Sepsis Triage, Admission NC)
PRECEDENCE(ER Sepsis Triage, Admission IC)	NOTSUCCESSION(Admission IC, IV Antibiotics)	NOTSUCCESSION(Admission IC, IV Liquid)
NOTSUCCESSION(Admission NC, IV Antibiotics)	NOTSUCCESSION(Admission NC, IV Liquid)	PRECEDENCE(ER Sepsis Triage, Return ER)
ATMOSTONE(Return ER)	NOTSUCCESSION(Return ER, Admission IC)	NOTSUCCESSION(Return ER, Admission NC)
NOTSUCCESSION(Return ER, CRP)	NOTSUCCESSION(Return ER, ER Sepsis Triage)	NOTSUCCESSION(Return ER, IV Antibiotics)
NOTSUCCESSION(Return ER, IV Liquid)	NOTSUCCESSION(Return ER, LacticAcid)	NOTSUCCESSION(Return ER, Leucocytes)
NOTSUCCESSION(Return ER, Release A)	NOTSUCCESSION(Return ER, Release B)	NOTSUCCESSION(Return ER, Release C)
NOTSUCCESSION(Return ER, Release D)	NOTSUCCESSION(Return ER, Release E)	ATMOSTONE(Release A)
ATMOSTONE(Release B)	ATMOSTONE(Release C)	ATMOSTONE(Release D)
ATMOSTONE(Release E)	NOTCOEXISTENCE(Release A, Release B)	NOTCOEXISTENCE(Release A, Release C)
NOTCOEXISTENCE(Release A, Release D)	NOTCOEXISTENCE(Release A, Release E)	NOTCOEXISTENCE(Release B, Release A)
NOTCOEXISTENCE(Release B, Release C)	NOTCOEXISTENCE(Release B, Release D)	NOTCOEXISTENCE(Release B, Release E)
NOTCOEXISTENCE(Release C, Release A)	NOTCOEXISTENCE(Release C, Release B)	NOTCOEXISTENCE(Release C, Release D)
NOTCOEXISTENCE(Release C, Release E)	NOTCOEXISTENCE(Release D, Release A)	NOTCOEXISTENCE(Release D, Release C)
NOTCOEXISTENCE(Release D, Release B)	NOTCOEXISTENCE(Release D, Release E)	NOTCOEXISTENCE(Release E, Release B)
NOTCOEXISTENCE(Release E, Release C)	NOTCOEXISTENCE(Release E, Release D)	NOTCOEXISTENCE(Release E, Release A)
NOTSUCCESSION(Release A, Admission IC)	NOTSUCCESSION(Release A, Admission NC)	NOTSUCCESSION(Release A, CRP)
NOTSUCCESSION(Release A, ER Sepsis Triage)	NOTSUCCESSION(Release A, IV Antibiotics)	NOTSUCCESSION(Release A, IV Liquid)
NOTSUCCESSION(Release A, LacticAcid)	NOTSUCCESSION(Release A, Leucocytes)	NOTSUCCESSION(Release B, Admission IC)
NOTSUCCESSION(Release B, Admission NC)	NOTSUCCESSION(Release B, CRP)	NOTSUCCESSION(Release B, ER Sepsis Triage)
NOTSUCCESSION(Release B, IV Antibiotics)	NOTSUCCESSION(Release B, IV Liquid)	NOTSUCCESSION(Release B, LacticAcid)
NOTSUCCESSION(Release B, Leucocytes)	NOTSUCCESSION(Release C, Admission IC)	NOTSUCCESSION(Release C, Admission NC)
NOTSUCCESSION(Release C, CRP)	NOTSUCCESSION(Release C, ER Sepsis Triage)	NOTSUCCESSION(Release C, IV Antibiotics)
NOTSUCCESSION(Release C, IV Liquid)	NOTSUCCESSION(Release C, LacticAcid)	NOTSUCCESSION(Release C, Leucocytes)
NOTSUCCESSION(Release D, Admission IC)	NOTSUCCESSION(Release D, Admission NC)	NOTSUCCESSION(Release D, CRP)
NOTSUCCESSION(Release D, ER Sepsis Triage)	NOTSUCCESSION(Release D, IV Antibiotics)	NOTSUCCESSION(Release D, IV Liquid)
NOTSUCCESSION(Release D, LacticAcid)	NOTSUCCESSION(Release D, Leucocytes)	NOTSUCCESSION(Release E, Admission IC)
NOTSUCCESSION(Release E, Admission NC)	NOTSUCCESSION(Release E, CRP)	NOTSUCCESSION(Release E, ER Sepsis Triage)
NOTSUCCESSION(Release E, IV Antibiotics)	NOTSUCCESSION(Release E, IV Liquid)	NOTSUCCESSION(Release E, LacticAcid)
NOTSUCCESSION(Release E, Leucocytes)		

results. The informed proposal of best practices for the selection of measures and combinations thereof depending on the analysis purposes (e.g., discovery) constitutes a highly interesting outlook for research.

6.2. Sensitivity and resilience to noise

In this section, we study in details the effect of noise injection in the event log on the constraints log measurements. We experimentally observed in Section 6.1.2 that measures exhibit different changes upon the presence of alterations from the expected behavior in an event log. A measure may “sense” the alteration in the data with respect to a rule or remain unaffected. Also, if the alteration is perceived, the magnitude of the measures reaction may be different, in light of the different properties that the measures enjoy, as discussed in Section 3.2. Markedly, we will empirically demonstrate that the properties originally for association rules hold also in the context of temporal rules. The possibility to characterize a constraint evolution is crucial in continuous measurement settings such as streaming analysis [68] or drift analysis [67]. Informed decisions on the measures to monitor based on the characteristic they have and the properties they enjoy is, therefore, key. Providing a set of guidelines to support such decisions is out of scope for this paper and paves the path for future research avenues. Nevertheless, with this experiment, we hope to provide useful preliminary indications in that sense. We also remark that while we call uncommon or unexpected events as “noise” or “error” (reflecting our experimental setting), the same measures evolution would occur in the case of process improvements or changes in the normative aspects the process is subject to.

We conducted the experiment as follows. We took as a reference model the set of rules in Table 13 and we simulated it to generate a clean event log that is compliant with it. We remark again how the rules are designed to not interfere with

one another. In this way, it is possible to observe the response of measures at varying noise levels targeting one constraint at a time, thus limiting the effect of cross-interference.

Thereupon, we injected noise in the event log and calculated all the measures in Table 3 for the reference model. In particular, we made use of the following types of noise [66]:

Events insertion: Spurious events are included in the traces (mimicking, e.g., double records, alien events, etc.);

Events deletion: Events are expunged from the event log (mimicking, e.g., missing records, uncommitted transactions, etc.);

White noise: Events are randomly inserted and deleted.

Addressing one rule at a time, we studied (1) the direct effect of noise on that constraint by altering the occurrences of its activator and target via insertions and deletions, and (2) the indirect effect, by altering the occurrences of the other tasks in the event log with white noise. We made the noise spread all over the event log according to a controlled probability variable. For instance, setting the noise injection as the deletion of occurrences of task a with a probability of 20% results in the removal of 20% of the occurrences of task a from the event log, picked at random.

More specifically, for every rule in the set of Table 13, we ran a separate experiment for (i) event insertion noise affecting the activator or (ii) the target, (iii) event deletion noise affecting the activator or (iv) the target, (v) white noise affecting neither the activator nor the target. For each of the combinations above, we let the error-injection probability range from 0 to 100% with a step of 10%. Because of the random factor, we repeated each experiment 10 times and recorded the average results.

Fig. 5 shows the results of our experiments on constraint $\text{RESPONSE}(e, f)$. The RESPONSE constraint imposes that the target

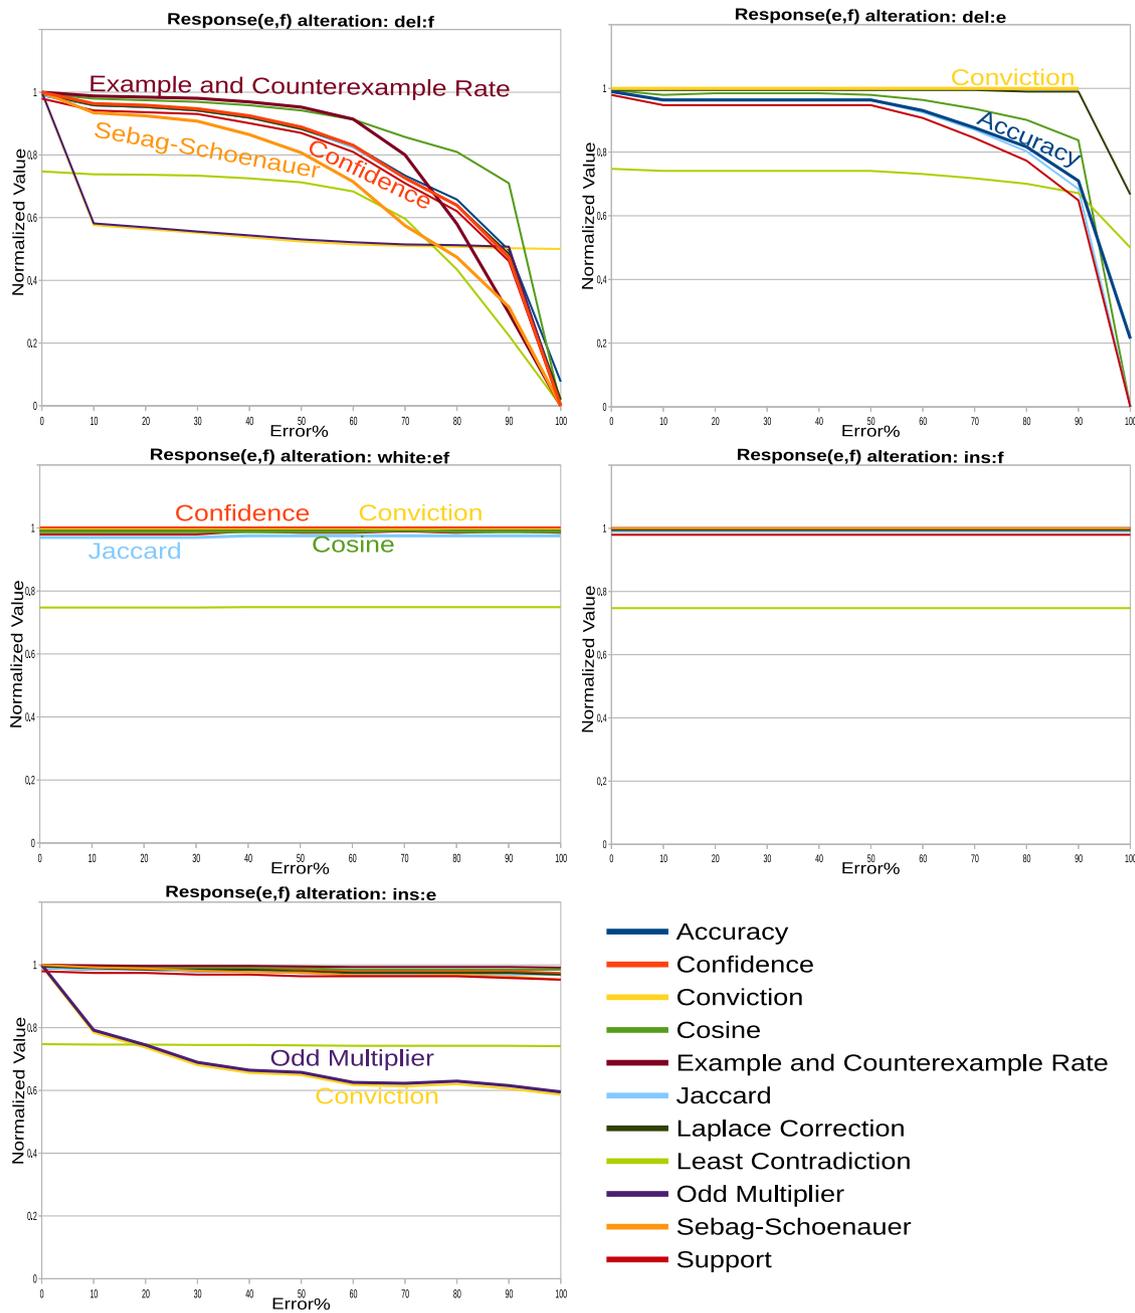


Fig. 5. Effect of error injection on constraint RESPONSE(e, f) for the best 10 scoring measures from the ranking experiment in Section 6.1.

occurs eventually after every occurrence of the activator. Plotting all the measures together in a static image within the boundaries of a page results in a complex intertwining of lines, hampering the readability of results. Thus we report here only the top 10 measures among the averaged results of the ranking experiment presented in Section 6.1, namely: Confidence, Example and Counterexample Rate, Laplace Correction, Least Contradiction, Accuracy, Cosine, Jaccard, Sebag–Schoenauer, Conviction, and Odd Multiplier plus Support as a baseline reference. Furthermore, as the measures have different ranges, we normalize them between 0 and 1 in order to compare the different trends. The normalization formula used is

$$\frac{\left(\frac{v_m}{1+|v_m|} - \frac{\min_m}{1+|\min_m|}\right)}{\left(\frac{\max_m}{1+|\max_m|} - \frac{\min_m}{1+|\min_m|}\right)}$$

where v_m is the value of measure, $|v_m|$ its absolute value, \max_m and \min_m the maximum and minimum values of the range of a measure, respectively (in case of infinite ranges, the maximum value supported by the software is considered). The full set of plots with all the measures and constraints can be interactively explored at <https://github.com/Oneiroe/DeclarativeMeasurement-static>.

For the RESPONSE constraint, the selected measures are overall particularly sensitive to the deletions of the target and influenced by both the deletions and insertions of the activator, while mostly insensitive to spurious insertions of the target and white noise. More specifically, we can derive the following observations.

- The deletion of events that satisfy the target leads to more violations of the rule, i.e., lower $P(\varphi_\tau|\varphi_\alpha, L)$ and higher $P(\neg\varphi_\tau|\varphi_\alpha, L)$. The negative effect is thus reflected in the constant decrease of all measures. The rapidity at which

this phenomenon occurs is particularly interesting. In accordance with property **P6**, the decrease of concave measures (e.g., Example and Counterexample Rate) is slower than that of linear ones (e.g., Confidence), and convex measures decrease faster than all the others (e.g., Sebag–Schoenauer).

- The deletion of events that satisfy the activator, instead, does not bring further violations, but a higher $P(\neg\varphi_\alpha, L)$. As a consequence, the measures are mostly stable until the error rate is close to 100%. As per **P5**, measures with a decreasing trend (e.g., Laplace Correction or Accuracy) are susceptible to the frequency of the rule activators in the event logs, while measures depending only on rule satisfaction (e.g., Conviction and Odd Multiplier) are constant at all the error rates except 100% (in which case, they are no longer defined due to a zero-by-zero division in their formula).
- The insertion of more events that satisfy the target lead to a higher $P(\varphi_\tau, L)$ but does not influence the constraint satisfactions, as $P(\varphi_\tau|\varphi_\alpha, L)$ remains constant. We remark that the target φ_τ of $\text{RESPONSE}(e, f)$ is $\diamond f$, thus injecting more occurrences of f does imply an increase of $P(\varphi_\tau, L)$ proportional to the error rate. Following from this consideration and in light of **P3**, all the selected measures remain constant.
- The insertion of more events that satisfy the activator is less definite, as the new tasks may bring either to more rule satisfactions, i.e., higher $P(\varphi_\tau|\varphi_\alpha, L)$, or violations, i.e., higher $P(\neg\varphi_\tau|\varphi_\alpha, L)$. As $\text{RESPONSE}(e, f)$ requires the occurrence of f eventually after e at any distance in the trace, in this experiment the injections mostly leads to satisfactions. That is why most of the measures show only a slightly decreasing trend. Notably, Conviction and Odd Multiplier measures sense this alteration in a marked way with respect to all the other rules.
- Lastly, the insertion of events not related to both the target and the activation mostly do not alter the measures. The satisfactions and violations remain constant, whilst the only increase is in $P(\neg\varphi_\alpha, L)$. Nevertheless, considering **P1** and **P5**, we can distinguish the fluctuating measures due to the frequencies change (e.g., Cosine and Jaccard) from the unaffected ones (e.g., Confidence and Conviction).

Because of space limits, we cannot illustrate the outcome for the other rules of [Table 13](#). The interested reader can find the full set of experimental data and results at <https://oneiroe.github.io/DeclarativeMeasurements-static> in a digital interactive format, which is more suitable for data exploration and browsing.

6.3. Discussion

The applications of our framework presented in Sections [6.1](#) and [6.2](#) allows for some reflections on the employment and availability of the measures. First, the measures evolution reflects whether the alteration in the data affects either the target or the activator of the rule (or both). This gives more insights into the analyzed phenomena than knowing whether a rule is satisfied or not at a given moment.

Nevertheless, given a constraint, not all the data alterations are perceived by all measures. This implies that depending on the requirements of the analysis the choice of the measure is crucial. We reported, for example, an experiment to identify measures suitable for discovery. Furthermore, given a certain alteration, it is possible to identify groups of measures with similar trends that focus on the same aspects of a rule, as it can be easily noticed in [Fig. 5](#). Within such groups, it is then possible to select one representative measure.

To this extent, the properties discussed in [Section 3.2](#) are clearly a guiding criterion for the selection of measures, markedly reflected in our experiments in which measures with similar trends could be distinctly identified by the properties they enjoy. Ultimately, this choice is strictly related to the requirements and goals of the event log analysis.

For instance, **P6** turned out to be particularly relevant. While measures may have similar trends, the magnitude of such trends, reflected in the steepness of the curves with which the measure evolves (i.e., concave, convex, or linear), indicates how quickly the measure react to changes in the data. Furthermore, it shows the range of tolerance before the alteration in the data becomes too large to recognize the specific constraint behavior. A resilient measure with a slower decrease is desirable to sense if the fundamental characteristics of a log are still visible despite the deviations (e.g., to implement discovery algorithms that are robust to noise). On the other hand, a sensitive measure with a faster decrease is desirable when exceptions to the rules are only accepted in a very limited manner (e.g., to monitor normative processes).

An approach enabling such a vast number of measures for temporal specifications is presented in the seminal work of Le and Lo [\[6\]](#). We extend their investigation along two main lines: first, our approach can handle arbitrary temporal formulae as activator and target, as opposed to single-task variables only (which also restrict the analysis to the sole RESPONSE and PRECEDENCE patterns). Second, our computation of the rule probabilities processes the whole trace at once, while the calculation scheme in [\[6\]](#) relies on a sliding window, whose size has to be manually set up thereby influencing the results.

7. Conclusion

In this paper, we presented a comprehensive measurement framework for declarative specifications modeled as Reactive Constraints. Given an event log and a set of custom probabilistic measures, the framework accepts in input any RCon and returns as output the evaluation of the rule for each event of the log, the computed measures for all the traces and their statistics over the entire event log, and the computed measures over the entire log. The framework goes beyond the current state of the art as it is not limited to a specific set of measures or rules. The experiments conducted reveal the possibility to characterize the behavior of a given constraint through the combination of different measures, which sense differently the behavior recorded in the log. Also, while the choice of the measures to employ is highly context-dependent, we showed how the measures properties can be used to guide the selection, as their effects are clearly visible in the results.

Future work Different possibilities are now open upon the foundations of this measurement framework. It is possible to exploit the possibility to characterize a phenomenon by studying the evolution of different measures for, e.g., the dynamic recognition of exceptions in process monitoring [\[69\]](#), the identification of process drifts [\[67\]](#) or the analysis of process variants [\[70\]](#). Also, it can be employed as a post-processing tool for multi-measure filtering of the results of declarative process discovery techniques like [\[3,4,14\]](#).

As measures react differently to diverse stimuli for distinct types of rules, a method to find the best combination of measures depending on the analysis context turns out to be key. To this extent, future research could resort to existing techniques like [\[71\]](#) or develop novel multi-measure heuristics. Also, the measures can be integrated for the assessment of multi-constraint specifications as a whole as in [\[1,20\]](#).

While the analysis of multiple measures at once may be overwhelming for a human, machine learning techniques could benefit from the availability of the great amount of information returned by the proposed framework, as they are designed to deal with large sets of multidimensional data. Therefore, it seems to be also promisingly exploitable for feature selection tasks in sequence classification [24].

Finally, we observe that the implementation of this framework can largely benefit from run-time optimization for the verification of the rules' automata, particularly for as far as the recognition of permanent violations and satisfactions is concerned [72]. The design and integration of such dedicated techniques serves as an impulse for future research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors want to thank Arik Senderovich for the fruitful discussion and useful observations on this research.

The work of Claudio Di Ciccio was partly supported by the Italian Ministry of University and Research (MUR) under grant "Dipartimenti di eccellenza 2018-2022" of the Department of Computer Science at Sapienza University of Rome and by the Sapienza SPECTRA research project, Italy.

The research by Jan Mendling was supported by the Einstein Foundation Berlin.

Appendix. Proof of linearity

The RCon evaluation can be performed efficiently based on the automaton-based technique defined in [14]. The verification is there defined for online settings (run-time analysis). In contrast, we can exploit the offline setting of the present contribution (post-mortem analysis) to enhance furthermore the performances.

Before proceeding, we need to extend the $LTLp_f$ notations presented in Section 3. In the following, we will classify $\circ, \square, \diamond, \mathbf{U}$ as *future operators*, $\ominus, \boxminus, \diamond, \mathbf{S}$ as *past operators*, and the following pairs of operators as *mirror images*: (i) \circ and \ominus , (ii) \square and \boxminus , (iii) \diamond and \diamond , (iv) \mathbf{U} and \mathbf{S} . A $LTLp_f$ formula φ is named: *pure past* ($\varphi^\blacktriangleleft$) if it contains only past operators; *pure present* ($\varphi^\blacktriangleright$) if it contains no temporal operators at all; *pure future* ($\varphi^\blacktriangleright$) if it contains only future operators [73]. For example, $\varphi^\blacktriangleleft = \diamond(a \mathbf{S} g)$, $\varphi^\blacktriangleright = a \wedge b \vee c$, and $\varphi^\blacktriangleright = \circ(\square e \vee (\diamond b) \mathbf{W} p)$ are pure past, pure present, and pure future formulae, respectively. Finally, the *mirror image* φ_M of formula φ is the temporal formula obtained by replacing all its operators with their mirror images [74].

In [14] it has been proven that

- any $LTLp_f$ formula φ can be decomposed in $\varphi \equiv \bigvee_{j=1}^m (\varphi^\blacktriangleleft \wedge \varphi^\blacktriangleright \wedge \varphi_j^\blacktriangleright)$, where $\varphi^\blacktriangleleft$, $\varphi^\blacktriangleright$, and $\varphi_j^\blacktriangleright$ are respectively pure past, present, and future formulae [14, Definition 2];
- each sub-formula of the decomposed formula can be evaluated through distinct automata [14, Theorem 3];
- The past-only automata can be reversed into future-only automata valuating the reverse of a trace [14, Theorem 4].

The advantage of the reversion is that it allows for the verification of past components in only one read of the trace. This was not possible for future components because the end of the trace is unknown in an online setting. In an offline setting, as the entire trace is given, it is possible to apply the same reasoning to reverse

future-only automata into past-only automata valuating a trace in one backward read from the end to the start of the trace.

In the following, we provide proof sketches of the intuitions above. We replicate the steps taken for past reversion [14] applying them for the future reversion. To this extent, we rely on mirror images and reversed automata.

Lemma A.1. *Let $t \in \Sigma^*$ be a trace of length n and t_R its reverse. Given a pure future formula $\varphi^\blacktriangleright$, and its mirror image $\varphi_M^\blacktriangleright$, then $t, 1 \models \varphi^\blacktriangleright$ iff $t_R, n \models \varphi_M^\blacktriangleright$.*

The proof follows from the semantics of future and past operators of $LTLp_f$ provided in Section 3. For instance, verifying $\varphi^\blacktriangleright = \diamond a$ on $t = \langle f, e, d, c, b, a, g, h, i \rangle$ at instant $i = 1$ is equivalent to verifying $\varphi_M^\blacktriangleright = \diamond a$ on $t_R = \langle i, h, g, a, b, c, d, e, f \rangle$ at $i = 9$. Notice that this holds for sub-traces too, thus verifying $\varphi^\blacktriangleright$ on t at instant $i = 5$ is equivalent to verifying $\varphi_M^\blacktriangleright$ over $t_{[5:9]R} = \langle i, h, g, a, b, f \rangle$ at $i = 9$.

It follows from Lemma A.1 that any pure future formula can be seen as a pure past one on a reversed trace. Therefore the automaton verifying the mirror image of $\varphi^\blacktriangleright$ can be used for verification on the reversed trace, as stated in the following.

Corollary A.1. *Let $A_{\varphi_M^\blacktriangleright}$ be the automaton verifying $\varphi_M^\blacktriangleright$. Then $t, n \models \varphi^\blacktriangleright$ iff $t_R \in \mathcal{L}(A_{\varphi_M^\blacktriangleright})$.*

Notice that $\varphi_M^\blacktriangleright$ is a pure past formula, therefore $A_{\varphi_M^\blacktriangleright}$ can be built by applying the technique of [75]. Furthermore, it is possible to transform the obtained automaton in order to read directly the original trace t thanks to the property of *closure under reversion* of regular languages [76].

From Lemmas A.1 and A.1 we derive the following.

Theorem A.1 (Valuation through $\overleftarrow{A}_{\varphi_M^\blacktriangleright}$). *Let $\varphi^\blacktriangleright$ be a pure future formula and $\varphi_M^\blacktriangleright$ its mirror image. Let $\overleftarrow{A}_{\varphi_M^\blacktriangleright} \in \mathcal{A}$ be the automaton verifying $\varphi_M^\blacktriangleright$. Given a trace $t \in \Sigma^*$ of length n , we have that: $t, 1 \models \varphi^\blacktriangleright$ iff $t \in \mathcal{L}(\overleftarrow{A}_{\varphi_M^\blacktriangleright})$.*

Consider $t = \langle f, e, d, c, b, a, g, h, i \rangle$ and the RCon $\Psi = a \circ \rightarrow (\ominus b \vee \diamond c)$ and the pure future formula of its separated automata set (sep.aut.set) $\varphi^\blacktriangleright = \diamond c$. The RCon is activated by $t(6)$. Its mirror image is $\varphi_M^\blacktriangleright = \diamond c$. It is possible to verify $\varphi_M^\blacktriangleright$ over trace $t_{[6:9]R} = \langle i, h, g, a \rangle$ at $i = 9$, thereby verifying $\varphi^\blacktriangleright$ over $t_{[6:9]}$ as per Lemma A.1. Thanks to Theorem A.1, $\varphi_M^\blacktriangleright$ can be verified on $t_{[6:9]}$ with the reversed automaton $\overleftarrow{A}_{\varphi_M^\blacktriangleright}$.

In conclusion, we proved that any pure future formula can be verified by parsing sub-traces from the end of the trace till the activator event. As the verification of future formulae is the only potentially exponential time element in [14], using the reversed verification provided in this section guarantees the optimization to linear time, which has been also empirically proved in Section 5.

References

- [1] M. de Leoni, F.M. Maggi, W.M.P. van der Aalst, An alignment-based framework to check the conformance of declarative process models and to preprocess event-log data, *Inf. Syst.* 47 (2015) 258–277, <http://dx.doi.org/10.1016/j.is.2013.12.005>.
- [2] L.T. Ly, F.M. Maggi, M. Montali, S. Rinderle-Ma, W.M.P. van der Aalst, Compliance monitoring in business processes: Functionalities, application, and tool-support, *Inf. Syst.* 54 (2015) 209–234.
- [3] F.M. Maggi, R.P.J.C. Bose, W.M.P. van der Aalst, Efficient discovery of understandable declarative process models from event logs, in: *CAiSE*, 2012, pp. 270–285, http://dx.doi.org/10.1007/978-3-642-31095-9_18.
- [4] C. Di Ciccio, M. Mecella, On the discovery of declarative control flows for artful processes, *ACM Trans. Manage. Inf. Syst.* 5 (4) (2015) 24:1–24:37, <http://dx.doi.org/10.1145/2629447>.

- [5] W.M.P. van der Aalst, M. Pesic, DecSerFlow: Towards a truly declarative service flow language, in: WS-FM, 2006, pp. 1–23, http://dx.doi.org/10.1007/11841197_1.
- [6] T.B. Le, D. Lo, Beyond support and confidence: Exploring interestingness measures for rule-based specification mining, in: SANER, 2015, pp. 331–340, <http://dx.doi.org/10.1109/SANER.2015.7081843>.
- [7] W. Hämmäläinen, G.I. Webb, A tutorial on statistically sound pattern discovery, *Data Min. Knowl. Discov.* 33 (2) (2019) 325–377, <http://dx.doi.org/10.1007/s10618-018-0590-x>.
- [8] S. Brin, R. Motwani, C. Silverstein, Beyond market baskets: Generalizing association rules to correlations, in: J. Peckham (Ed.), SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data, May 13–15, 1997, Tucson, Arizona, USA, ACM Press, 1997, pp. 265–276, <http://dx.doi.org/10.1145/253260.253327>.
- [9] G. Piatetsky-Shapiro, *Discovery, analysis, and presentation of strong rules*, in: G. Piatetsky-Shapiro, W.J. Frawley (Eds.), *Knowledge Discovery in Databases*, AAAI/MIT Press, 1991, pp. 229–248.
- [10] F.B. Galiano, I.J. Blanco, D. Sánchez, M.A.V. Miranda, Measuring the accuracy and interest of association rules: A new framework, *Intell. Data Anal.* 6 (3) (2002) 221–235.
- [11] L. Geng, H.J. Hamilton, Interestingness measures for data mining: A survey, *ACM Comput. Surv.* 38 (3) (2006) 9, <http://dx.doi.org/10.1145/1132960.1132963>.
- [12] C.C. Aggarwal, P.S. Yu, A new framework for itemset generation, in: A.O. Mendelzon, J. Paredaens (Eds.), Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 1–3, 1998, Seattle, Washington, USA, ACM Press, 1998, pp. 18–24, <http://dx.doi.org/10.1145/275487.275490>.
- [13] S. Brin, R. Motwani, J.D. Ullman, S. Tsur, Dynamic itemset counting and implication rules for market basket data, in: J. Peckham (Ed.), SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data, May 13–15, 1997, Tucson, Arizona, USA, ACM Press, 1997, pp. 255–264, <http://dx.doi.org/10.1145/253260.253325>.
- [14] A. Cecconi, C. Di Ciccio, G. De Giacomo, J. Mendling, Interestingness of traces in declarative process mining: The Janus LTL_f approach, in: BPM, 2018, pp. 121–138, http://dx.doi.org/10.1007/978-3-319-98648-7_8.
- [15] P. Lenca, P. Meyer, B. Vaillant, S. Lallich, On selecting interestingness measures for association rules: User oriented description and multiple criteria decision aid, *European J. Oper. Res.* 184 (2) (2008) 610–626, <http://dx.doi.org/10.1016/j.ejor.2006.10.059>.
- [16] A. Cecconi, G. De Giacomo, C. Di Ciccio, F.M. Maggi, J. Mendling, A temporal logic-based measurement framework for process mining, in: B.F. van Dongen, M. Montali, M.T. Wynn (Eds.), 2nd International Conference on Process Mining, ICPM 2020, Padua, Italy, October 4–9, 2020, IEEE, 2020, pp. 113–120, <http://dx.doi.org/10.1109/ICPM49681.2020.00026>.
- [17] R. Agrawal, R. Srikant, Fast algorithms for mining association rules in large databases, in: J.B. Bocca, M. Jarke, C. Zaniolo (Eds.), VLDB'94, Proceedings of 20th International Conference on Very Large Data Bases, September 12–15, 1994, Santiago de Chile, Chile, Morgan Kaufmann, 1994, pp. 487–499.
- [18] S. ren Debois, T.T. Hildebrandt, P.H. vsgaard Laursen, K.R. Ulrik, Declarative process mining for DCR graphs, in: SAC, 2017, pp. 759–764, <http://dx.doi.org/10.1145/3019612.3019622>.
- [19] W.M. van der Aalst, H. de Beer, B.F. van Dongen, Process mining and verification of properties: An approach based on temporal logic, in: OTM To Meaningful Internet Systems, 2005, pp. 130–147, http://dx.doi.org/10.1007/11575771_11.
- [20] A. Burattin, F.M. Maggi, W.M. van der Aalst, A. Sperduti, Techniques for a posteriori analysis of declarative processes, in: EDOC, 2012, pp. 41–50, <http://dx.doi.org/10.1109/EDOC.2012.15>.
- [21] J. Yang, D. Evans, D. Bhardwaj, T. Bhat, M. Das, Perracotta: mining temporal API rules from imperfect traces, in: ICSE, 2006, pp. 282–291, <http://dx.doi.org/10.1145/1134285.1134325>.
- [22] M. Gabel, Z. Su, Online inference and enforcement of temporal properties, in: ICSE, 2010, pp. 15–24, <http://dx.doi.org/10.1145/1806799.1806806>.
- [23] C. Lemieux, D. Park, I. Beschastnikh, General LTL specification mining (T), in: ASE, 2015, pp. 81–92, <http://dx.doi.org/10.1109/ASE.2015.71>.
- [24] Z. Xing, J. Pei, E.J. Keogh, A brief survey on sequence classification, *SIGKDD Explor.* 12 (1) (2010) 40–48, <http://dx.doi.org/10.1145/1882471.1882478>.
- [25] E. Egho, D. Gay, M. Boullé, N. Voisine, F. Clérot, A user parameter-free approach for mining robust sequential classification rules, *Knowl. Inf. Syst.* 52 (1) (2017) 53–81, <http://dx.doi.org/10.1007/s10115-016-1002-4>.
- [26] J.M. Fowkes, C. Sutton, A subsequence interleaving model for sequential pattern mining, in: B. Krishnapuram, M. Shah, A.J. Smola, C.C. Aggarwal, D. Shen, R. Rastogi (Eds.), Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016, ACM, 2016, pp. 835–844, <http://dx.doi.org/10.1145/2939672.2939787>.
- [27] J.D. Smedt, G. Deeva, J.D. Weerdt, Mining behavioral sequence constraints for classification, *IEEE Trans. Knowl. Data Eng.* 32 (6) (2020) 1130–1142, <http://dx.doi.org/10.1109/TKDE.2019.2897311>.
- [28] G. De Giacomo, M.Y. Vardi, Linear temporal logic and linear dynamic logic on finite traces, in: IJCAI, 2013, pp. 854–860.
- [29] M. Pesic, D. Bosnacki, W.M. van der Aalst, Enacting declarative languages using LTL: avoiding errors and improving performance, in: SPIN, in: LNCS, vol. 6349, Springer, 2010, pp. 146–161.
- [30] C. Baier, J.-P. Katoen, K. Guldstrand Larsen, *Principles of Model Checking*, The MIT Press, 2008.
- [31] M. Dumas, M. La Rosa, J. Mendling, H.A. Reijers, *Fundamentals of business process management*, second edition, Springer, 2018, <http://dx.doi.org/10.1007/978-3-662-56509-4>.
- [32] G. De Giacomo, R. De Masellis, M. Montali, Reasoning on LTL on finite traces: Insensitivity to infiniteness, in: AAAI Conference on Artificial Intelligence, AAAI Press, 2014, pp. 1027–1033.
- [33] C. Di Ciccio, F.M. Maggi, M. Montali, J. Mendling, Resolving inconsistencies and redundancies in declarative process models, *Inf. Syst.* 64 (2017) 425–446, <http://dx.doi.org/10.1016/j.is.2016.09.005>.
- [34] F. Mannhardt, Sepsis Cases - Event Log, 4TU.ResearchData, 2016, <http://dx.doi.org/10.4121/uuid:915d2bfb-7e84-49ad-a286-dc35f063a460>.
- [35] F. Mannhardt, D. Blinde, Analyzing the trajectories of patients with sepsis using process mining, in: J. Gulden, S. Nurcan, I. Reinhartz-Berger, W. Guédria, P. Bera, S. Guerreiro, M. Fellmann, M. Weidlich (Eds.), BPMDS, in: CEUR Workshop Proceedings, vol. 1859, CEUR-WS.org, 2017, pp. 72–80.
- [36] M. Ohsaki, S. Kitaguchi, K. Okamoto, H. Yokoi, T. Yamaguchi, Evaluation of rule interestingness measures with a clinical dataset on hepatitis, in: J. Boulicaut, F. Esposito, F. Giannotti, D. Pedreschi (Eds.), PKDD, in: Lecture Notes in Computer Science, vol. 3202, Springer, 2004, pp. 362–373, http://dx.doi.org/10.1007/978-3-540-30116-5_34.
- [37] P.-N. Tan, V. Kumar, J. Srivastava, Selecting the right objective measure for association analysis, *Inf. Syst.* 29 (4) (2004) 293–313, [http://dx.doi.org/10.1016/S0306-4379\(03\)00072-3](http://dx.doi.org/10.1016/S0306-4379(03)00072-3), Knowledge Discovery and Data Mining (KDD 2002).
- [38] C.L. Siström, C.W. Garvan, Proportions, odds, and risk, *Radiology* 230 (1) (2004) 12–19, <http://dx.doi.org/10.1148/radiol.2301031028>, PMID: 14695382.
- [39] C. van Rijsbergen, *Information Retrieval*, Butterworth, 1979.
- [40] F. Mosteller, Association and estimation in contingency tables, *J. Amer. Statist. Assoc.* 63 (321) (1968) 1–28, <http://dx.doi.org/10.1080/01621459.1968.11009219>.
- [41] G.U. Yule, K. Pearson, VII. On the association of attributes in statistics: with illustrations from the material of the childhood society, *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* 194 (252–261) (1900) 257–319, <http://dx.doi.org/10.1098/rsta.1900.0019>.
- [42] G.U. Yule, On the methods of measuring association between two attributes, *J. R. Stat. Soc.* 75 (6) (1912) 579–652.
- [43] W. Klösgen, Problems for knowledge discovery in databases and their treatment in the statistics interpreter *explora*, *Int. J. Intell. Syst.* 7 (7) (1992) 649–673, <http://dx.doi.org/10.1002/int.4550070707>.
- [44] B. Gray, M.E. Orłowska, CCAIA: clustering categorical attributed into interesting association rules, in: X. Wu, K. Ramamohanarao, K.B. Korb (Eds.), Research and Development in Knowledge Discovery and Data Mining, Second Pacific-Asia Conference, PAKDD-98, Melbourne, Australia, April 15–17, 1998, Proceedings, in: Lecture Notes in Computer Science, vol. 1394, Springer, 1998, pp. 132–143, http://dx.doi.org/10.1007/3-540-64383-4_12.
- [45] P. Clark, R. Boswell, Rule induction with CN2: some recent improvements, in: Y. Kodratoff (Ed.), Machine Learning - EWVL-91, European Working Session on Learning, Porto, Portugal, March 6–8, 1991, Proceedings, in: Lecture Notes in Computer Science, vol. 482, Springer, 1991, pp. 151–163, <http://dx.doi.org/10.1007/BFb0017011>.
- [46] L. Breiman, J. Friedman, R. Olshen, C. Stone, *Classification and Regression Trees*, Wadsworth, 1984.
- [47] P. Smyth, R.M. Goodman, An information theoretic approach to rule induction from databases, *IEEE Trans. Knowl. Data Eng.* 4 (4) (1992) 301–316, <http://dx.doi.org/10.1109/69.149926>.
- [48] Y. Yao, N. Zhong, An analysis of quantitative measures associated with rules, in: N. Zhong, L. Zhou (Eds.), Methodologies for Knowledge Discovery and Data Mining, Third Pacific-Asia Conference, PAKDD-99, Beijing, China, April 26–28, 1999, Proceedings, in: Lecture Notes in Computer Science, vol. 1574, Springer, 1999, pp. 479–488, http://dx.doi.org/10.1007/3-540-48912-6_64.
- [49] A. Agresti, *Categorical Data Analysis*, Wiley, 1990.
- [50] J. Loevinger (Ed.), A systematic approach to the construction and evaluation of tests of ability, *Psychol. Monograph.* 61 (4) (1947) i–49, <http://dx.doi.org/10.1037/h0093565>, Place: US Publisher: American Psychological Association.
- [51] A.A. Freitas, On objective measures of rule surprisingness, in: J.M. Zytzkow, M. Quafafou (Eds.), Principles of Data Mining and Knowledge Discovery, Second European Symposium, PKDD '98, Nantes, France, September 23–26, 1998, Proceedings, in: Lecture Notes in Computer Science, vol. 1510, Springer, 1998, pp. 1–9, <http://dx.doi.org/10.1007/BFb0094799>.

- [52] M. Sebag, M. Schoenauer, Generation of rules with certainty and confidence factors from incomplete and incoherent learning bases, in: Proc. EKAW, vol. 88, 1988, p. 28.
- [53] J. Azé, Y. Kodratoff, Extraction de "pépites" de connaissance dans les données: une nouvelle approche et une étude de la sensibilité au bruit, *Revue RNTI, E-1* (2004) 247–270.
- [54] H. Zhang, B. Padmanabhan, A. Tuzhilin, On the discovery of significant statistical quantitative rules, in: W. Kim, R. Kohavi, J. Gehrke, W. DuMouchel (Eds.), Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Seattle, Washington, USA, August 22–25, 2004, ACM, 2004, pp. 374–383, <http://dx.doi.org/10.1145/1014052.1014094>.
- [55] F.M. Maggi, C. Di Ciccio, C. Di Francescomarino, T. Kala, Parallel algorithms for the automated discovery of declarative process models, *Inf. Syst.* 74 (Part) (2018) 136–152, <http://dx.doi.org/10.1016/j.is.2017.12.002>.
- [56] A. Fugard, N. Pfeifer, B. Mayerhofer, G. Kleiter, How people interpret conditionals: Shifts toward the conditional event, *J. Experiment. Psychol.: Learn. Memory Cogn.* 37 (3) (2011) 635–648, <http://dx.doi.org/10.1037/a0022329>.
- [57] C. Di Ciccio, M.L. Bernardi, M. Cimitile, F.M. Maggi, Generating event logs through the simulation of declare models, in: EOMAS, 2015, pp. 20–36, http://dx.doi.org/10.1007/978-3-319-24626-0_2.
- [58] B. van Dongen, BPI Challenge 2017, 2017, <http://dx.doi.org/10.4121/uuid:5f3067df-f10b-45da-b98b-86ae4c7a310b>.
- [59] B. van Dongen, BPI Challenge 2012, 4TU.ResearchData, 2012, <http://dx.doi.org/10.4121/uuid:3926db30-f712-4394-aebc-75976070e91f>.
- [60] W. Steeman, BPI Challenge 2013, 4TU.ResearchData, 2014, <http://dx.doi.org/10.4121/uuid:a7ce5c55-03a7-4583-b855-98b86e1a2b07>.
- [61] B. van Dongen, BPI Challenge 2014, 4TU.ResearchData, 2014, <http://dx.doi.org/10.4121/uuid:c3e5d162-0cfd-4bb0-bd82-af5268819c35>.
- [62] B. van Dongen, BPI Challenge 2015, 4TU.ResearchData, 2015, <http://dx.doi.org/10.4121/uuid:31a308ef-c844-48da-948c-305d167a0ec1>.
- [63] M. de Leoni, F. Mannhardt, Road Traffic Fine Management Process, 4TU.ResearchData, 2015, <http://dx.doi.org/10.4121/uuid:270fd440-1057-4fb9-89a9-b699b47990f5>.
- [64] A. Azzam, P.R. Aryan, A. Cecconi, C. Di Ciccio, F.J. Ekaputra, J.D. Fernández, S. Karampatakis, E. Kiesling, A. Musil, M. Sabou, P. Shadlau, T. Thurner, The cityspin platform: A CPSS environment for city-wide infrastructures, in: A. Longo, M. Fazio, R. Ranjan, M. Zappatore (Eds.), CPSS@IoT, in: CEUR Workshop Proceedings, vol. 2530, CEUR-WS.org, 2019, pp. 57–64.
- [65] W.M. van der Aalst, M. Pesic, H. Schonenberg, Declarative workflows: Balancing between flexibility and support, *Comput. Sci. - R&D* 23 (2) (2009) 99–113, <http://dx.doi.org/10.1007/s00450-009-0057-9>.
- [66] C. Di Ciccio, M. Mecella, J. Mendling, The effect of noise on mined declarative constraints, in: SIMPDA, 2013, pp. 1–24, http://dx.doi.org/10.1007/978-3-662-46436-6_1.
- [67] A. Yeshchenko, C. Di Ciccio, J. Mendling, A. Polyvyanyy, Visual drift detection for event sequence data of business processes, *IEEE Trans. Vis. Comput. Graphics* (2021) <http://dx.doi.org/10.1109/TVCG.2021.3050071>.
- [68] A. Burattin, M. Cimitile, F.M. Maggi, A. Sperduti, Online discovery of declarative process models from event streams, *IEEE Trans. Serv. Comput.* 8 (6) (2015) 833–846, <http://dx.doi.org/10.1109/TSC.2015.2459703>.
- [69] F.M. Maggi, M. Montali, M. Westergaard, W.M.P. van der Aalst, Monitoring business constraints with linear temporal logic: An approach based on colored automata, in: BPM, 2011, pp. 132–147.
- [70] A. Cecconi, A. Augusto, C. Di Ciccio, Detection of statistically significant differences between process variants through declarative rules, in: A. Polyvyanyy, M.T. Wynn, A.V. Looy, M. Reichert (Eds.), Business Process Management Forum – BPM Forum 2021, Rome, Italy, September 06–10, 2021, Proceedings, in: Lecture Notes in Business Information Processing, vol. 427, Springer, 2021, pp. 73–91, http://dx.doi.org/10.1007/978-3-030-85440-9_5.
- [71] Z. Cao, Y. Tian, T.B. Le, D. Lo, Rule-based specification mining leveraging learning to rank, *Autom. Softw. Eng.* 25 (3) (2018) 501–530, <http://dx.doi.org/10.1007/s10515-018-0231-z>.
- [72] A. Bauer, M. Leucker, C. Schallhart, Runtime verification for LTL and TLTL, *ACM Trans. Softw. Eng. Methodol.* 20 (4) (2011) 14:1–14:64, <http://dx.doi.org/10.1145/2000799.2000800>.
- [73] D.M. Gabbay, The declarative past and imperative future: Executable temporal logic for interactive systems, in: B. Banieqbal, H. Barringer, A. Pnueli (Eds.), Temporal Logic in Specification, Altrincham, UK, April 8–10, 1987, Proceedings, in: Lecture Notes in Computer Science, vol. 398, Springer, 1987, pp. 409–448, http://dx.doi.org/10.1007/3-540-51803-7_36.
- [74] M. Reynolds, The complexity of temporal logic over the reals, *Ann. Pure Appl. Log.* 161 (8) (2010) 1063–1096, <http://dx.doi.org/10.1016/j.apal.2010.01.002>.
- [75] G. De Giacomo, A. Di Stasio, F. Fuggitti, S. Rubin, Pure-past linear temporal and dynamic logic on finite traces, in: C. Bessiere (Ed.), IJCAI, International Joint Conferences on Artificial Intelligence Organization, 2020, pp. 4959–4965, <http://dx.doi.org/10.24963/ijcai.2020/690>, Survey track.
- [76] J.E. Hopcroft, R. Motwani, J.D. Ullman, Introduction To Automata Theory, Languages, and Computation, third ed, in: Pearson international edition, Addison-Wesley, 2007.