Specification sketching for Linear Temporal Logic

Simon Lutz[1], Daniel Neider[2], and Rajarshi Roy[2]

[1] Technical University of Kaiserslautern, Kaiserslautern, Germany
[2] Max Planck Institute for Software Systems, Kaiserslautern, Germany

Abstract. Virtually all verification and synthesis techniques assume that the formal specifications are readily available, functionally correct, and fully match the engineer's understanding of the given system. However, this assumption is often unrealistic in practice: formalizing system requirements is notoriously difficult, error-prone, and requires substantial training. To alleviate this severe hurdle, we propose a fundamentally novel approach to writing formal specifications, named *specification sketching for Linear Temporal Logic (LTL)*. The key idea is that an engineer can provide a partial LTL formula, called an LTL *sketch*, where parts that are hard to formalize can be left out. Given a set of examples describing system behaviors that the specification should or should not allow, the task of a so-called *sketching algorithm* is then to complete a given sketch such that the resulting LTL formula is consistent with the examples. We show that deciding whether a sketch can be completed falls into the complexity class NP and present two SAT-based sketching algorithms. We also demonstrate that sketching is a practical approach to writing formal specifications using a prototype implementation.

Keywords: Formal Specifications · Linear Temporal Logic · Sketching

1 Introduction

Due to its unique ability to prove the absence of errors mathematically, formal verification is a time-tested method of ensuring the safe and reliable operation of safety-critical systems. Success stories of formal methods include application domains such as communication system [13,25], railway transportation [2,3], aerospace [16,9], and operating systems [37,21] to name but a few.

However, there is an essential and often overlooked catch with formal verification: virtually all techniques assume that the specification required for the design or verification of a system is available in a suitable format, is functionally correct, and expresses precisely the properties the engineers had in mind. These assumptions are often unrealistic in practice. Formalizing system requirements is notoriously difficult and error-prone [5,28,33]. Even worse, the training effort required to reach proficiency with specification languages can be disproportionate to the expected benefits [10], and the use of formalisms such as temporal logics require a level of sophistication that many users might never develop [19].

We address this severe, practical issue by introducing a fundamentally novel approach to writing formal specifications, named *specification sketching*. Inspired by recent advances in automated program synthesis [35,36], our new paradigm allows engineers to express their high-level insights about a system in terms of a partial specification, named *specification sketch*, where parts that are difficult or error-prone to formalize can be left out. Moreover, the engineer is asked to provide example executions of the system that the specification should allow or forbid. Based on this additional data, a so-called *sketching algorithm* then fills in the missing low-level details to obtain a complete specification.

While the concept of specification sketching is conceivable for a wide range of specification languages, we here focus on Linear Temporal Logic (LTL) [29]. The rationale for this choice is twofold. First, LTL is popular in academia and widely used in industry [15,16,18], making it the de facto standard for expressing (temporal) properties in verification and synthesis. Second, LTL is well-understood and enjoys good algorithmic properties [8,29]. Furthermore, its intuitive and variable-free syntax have recently prompted efforts to adopt LTL (over finite words) also in artificial intelligence (e.g., to specify reward functions in reinforcement learning [6]).

In the context of specification sketching for LTL, a sketch can leave logical operators or even entire subformulas unspecified, while examples are ultimately-periodic words. To illustrate this setting, let us consider the request-response property P expressing that every request p has to be answered eventually by a response q. This property can be formalized by the LTL formula $\varphi := \mathbf{G}(p \to \mathbf{X}\,\mathbf{F}\,q)$, which uses the standard temporal modalities $\mathbf{F}$ ("finally"), $\mathbf{G}$ ("globally"), and $\mathbf{X}$ ("next"). However, let us assume for the sake of this example that an engineer is unsure how to formalize P. In this situation, our sketching paradigm allows the engineer to express high-level insight in the form of a sketch, say $\mathbf{G}(p \to ?)$, where the question mark indicates which part of the specification is missing. Additionally, the engineer provides two examples: (i) a positive (infinite) trajectory $(\{p\}\{q\})^{\omega}$, expressing that q is the response that should be used to answer a request, and (ii) a negative (infinite) trajectory $\{q\}^{\omega}$, intended to disallow the system to send responses without requests. Our sketching algorithm then computes a substitution for the question mark such that the completed LTL formula is consistent with the examples (e.g., $\mathbf{X}\,\mathbf{F}\,q$). We set up all necessary definitions in Section 2 and formally define LTL sketching in Section 3.

It turns out that it is not always possible to find a substitution that is consistent with the given examples (see Example 1 on Page 8). However, we show in Section 4 that the problem of deciding whether such a substitution exists is in the complexity class NP. Moreover, we develop an effective decision procedure that reduces the original question to a satisfiability problem in propositional logic. This reduction permits us to apply highly-optimized, off-the-shelf SAT solvers to check whether a consistent substitution exists.

In Section 5, we develop two sketching algorithms for LTL. The first algorithm uses the decision procedure of Section 4 as a sub-routine and transforms the sketching problem into a series of LTL learning problems (i.e., in problems

of learning an LTL formula—without syntactic constraints—from positive and negative examples). This transformation allows us to apply a diverse array of learning algorithms for LTL, which have been proposed during the last five years [27,7,31]. Our second sketching algorithm, on the other hand, extends the LTL learning algorithm by Neider and Gavran [27] and uses a SAT-based technique as an effective means to search for solutions of increasing size.

Finally, Section 6 presents an empirical evaluation with a prototype implementation, named `LTL-Sketcher`. We observed that our algorithms are effective in completing a variety of sketches having different missing information. Further, comparing our algorithms, we concluded that our second algorithm outperforms our first algorithm in terms of running time and size of the solution. We conclude this paper in Section 7 with a discussion of promising directions for future work.

Related Work The general idea of allowing partial specifications is not entirely new, but it has not yet been investigated as general as in this work. The most similar setting is one in which templates are used to mine temporal specifications from system executions. In this context, a template is a partial formula, similar to a sketch. However, unlike a sketch, a template is typically completed with single atomic propositions or simple formulas, usually without temporal modalities (e.g., restricted Boolean combinations of atomic propositions). An example of this approach is Texada [22,23], a specification miner for LTL_f formulas (i.e., LTL over finite horizon). Texada accepts an arbitrary template (property type in their terminology) and a set of system executions as input and completes the template with atomic propositions such that the resulting LTL formula satisfies all of the system executions. Since LTL_f (and LTL for that matter) is defined over a finite, user-provided set of atomic propositions, there are only finitely many ways to fill in the missing parts of a template, and Texada amounts to a search over this finite search space. By contrast, our approach can complete a sketch with complex LTL formulas and, hence, has to search in an infinite space.

Various other techniques operate in settings where the templates are even more restricted. For example, Li et al. [24] mine LTL specification based on templates from the GR(1)-fragment of LTL (e.g., $\mathbf{G}\,\mathbf{F}?$, $\mathbf{G}(?_1 \rightarrow \mathbf{X}?_2)$, etc.), while Shah et al. [34] mine LTL formulas that are conjunctions of a set of common temporal properties as identified by Dwyer et al. [11]. In addition, the work by Kim et al. [20] uses a set of interpretable LTL templates, widely used in the development of software systems, to obtain LTL formulas robust to noise in the input data. In the context of CTL, on the other hand, Wasylkowski and Zeller [38] mine specifications using templates of the form $\mathbf{A}\,\mathbf{F}?$, $\mathbf{A}\,\mathbf{G}(?_1 \rightarrow \mathbf{F}?_2)$. However, all of the works above complete the templates only with atomic propositions (and their negations in some cases).

Another setting is the one in which general (and complex) temporal specifications are learned from system executions without any constraint on the structure of the specification. The most notable work in this setting is that by Neider and Gavran [27], who learn LTL formulas from system executions using a SAT solver. Similar to their Neider and Gavran is the work by Camacho et al. [7],

which proposes a SAT-based learning algorithm for LTL$_f$ formulas via Alternating Finite Automata as an intermediate representation. Raha et al. [30] present a scalable approach for learning formulas in a fragment of LTL$_f$ without the **U**-operator, while Roy, Fisman, and Neider [32] consider the Property Specification Language (PSL). However, none of these works can exploit insights about the structure of the specification to aid the learning/ mining process.

Finally, it is worth mentioning that LTL sketching can be seen as a particular case of syntax-guided synthesis, where syntactic constraints on the resulting formulas are expressed in terms of context-free grammars. While this approach is more expressive in that it also allows restricting the syntax of the formulas that can be used to fill in missing parts of a sketch, we are convinced that the concept of sketching is more natural and intuitive, and it allows for efficient algorithms. An example of a syntax-guided approach is SySLite [1], a CVC4-based tool for learning Past-time LTL (over finite executions). However, we are unaware of any syntax-guided algorithm for our setting: learning LTL formulas over infinite executions.

2   Preliminaries

We now formally introduce the basic notions that are used throughout the paper.

Finite and Infinite Words To model trajectories of a system, we use the notion of words over an alphabet that represents the system events. Formally, an *alphabet* Σ is a nonempty, finite set. The elements of this set are called *symbols*. A *finite word* over an alphabet Σ is a sequence $u = a_0 \ldots a_n$ of symbols $a_i \in \Sigma$ for $i \in \{0, \ldots, n\}$. The empty sequence, referred to as *empty word*, is denoted by ε. The length of a finite word u is denoted by $|u|$, where $|\epsilon| = 0$. Moreover, Σ^* denotes the set of all finite words over the alphabet Σ, while $\Sigma^+ = \Sigma^* \setminus \varepsilon$ is the set of all non-empty words.

An *infinite word* over Σ is an infinite sequence $\alpha = a_0 a_1 \ldots$ of symbols $a_i \in \Sigma$ for $i \in \mathbb{N}$. We denote the i-th symbol of an infinite word α by $\alpha[i]$ and the finite infix of α from position i up to (and excluding) position j with $\alpha[i, j] = a_i a_{i+1} \cdots a_{j-1}$. We use the convention that $\alpha[i, j] = \varepsilon$ for any $i \geq j$. Further, we denote the infinite suffix starting at position $j \in \mathbb{N}$ by $\alpha[j, \infty) = a_j a_{j+1} \cdots$. Given $u \in \Sigma^+$, the infinite repetition of u is the infinite word $u^\omega = uu \cdots \in \Sigma^\omega$. An infinite word α is called ultimately periodic if it is of the form $\alpha = uv^\omega$ for a $u \in \Sigma^*$ and $v \in \Sigma^+$. Finally, Σ^ω denotes the set of all infinite words over the alphabet Σ.

Propositional Boolean Logic Since the presented algorithms rely on the Satisfiability (SAT) problem, as a prerequisite, we introduce Propositional Logic. Let Var be a set of propositional variables, which take Boolean values from $\mathbb{B} = \{0, 1\}$ (0 representing *false* and 1 representing *true*). Formulas in propositional (Boolean) logic—which we denote by capital Greek letters—are inductively constructed as follows:

- each $x \in Var$ is a propositional formula; and
- if Ψ and Φ are propositional formulas, so are $\neg\Psi$ and $\Psi \vee \Phi$.

Moreover, we add syntactic sugar and allow the formulas *true*, *false*, $\Psi \wedge \Phi$, $\Psi \Rightarrow \Phi$, and $\Psi \Leftrightarrow \Phi$, which are defined as usual. A propositional valuation is a mapping $v : Var \to \mathbb{B}$, which maps propositional variables to Boolean values. The semantics of propositional logic is given by a satisfaction relation $\models$ that is inductively defined as follows: $v \models x$ if and only if $v(x) = 1$, $v \models \neg\Phi$ if and only if $v \nvDash \Phi$, and $v \models \Psi \vee \Phi$ if and only if $v \models \Psi$ or $v \models \Phi$. In the case that $v \models \Phi$, we say that v satisfies Φ and call it a *model* of Φ. A propositional formula Φ is *satisfiable* if there exists a model v of Φ. The *size* of a formula is the number of its subformulas (as defined in the usual way). The satisfiability problem of propositional logic is the problem of deciding whether a given formula is satisfiable. Although this problem is well-known to be NP-complete, modern SAT solvers implement optimized decision procedures that can check satisfiability of formulas with millions of variables [4]. Moreover, virtually all SAT solvers also return a model if the input-formula is satisfiable.

Linear Temporal Logic Linear Temporal Logic (LTL) is a logic to reason about sequences of relevant statements about a system by using temporal modalities. Formally, given set of propositions $\mathcal{P}$ that represent relevant statements about the system under consideration, an LTL formula—which we denote by small Greek letters—is defined inductively as follows:

- each proposition $p \in \mathcal{P}$ is an LTL formula; and
- if ψ and φ are LTL formulas, so are $\neg\psi$, $\psi \vee \varphi$, $\mathbf{X}\,\psi$ ("neXt"), and $\psi\,\mathbf{U}\,\varphi$ ("Until").

As syntactic sugar, we allow the formulas *true*, *false*, $\psi \wedge \varphi$ and $\psi \to \varphi$, which are defined as usual. Additionally, we allow the temporal formulas $\mathbf{F}\,\psi := true\,\mathbf{U}\,\psi$ ("Finally") and $\mathbf{G}\,\psi := \neg\,\mathbf{F}\,\neg\psi$ ("Globally"). Further, we use $\mathcal{F}_{\text{LTL}}$ to denote the set of all LTL formulas using.

LTL formulas are interpreted over infinite words $\alpha \in (2^{\mathcal{P}})^{\omega}$. To define how an LTL formula is interpreted on a word, we use a valuation function V. This function maps an LTL formula and a word to a Boolean value and is defined inductively as follows:

$$V(p, \alpha) = 1 \text{ if and only if } p \in \alpha(0),$$
$$V(\neg\varphi, \alpha) = 1 - V(\varphi, \alpha),$$
$$V(\varphi \vee \psi, \alpha) = \max\{V(\varphi, \alpha), V(\psi, \alpha)\},$$
$$V(\mathbf{X}\,\varphi, \alpha) = V(\varphi, \alpha[1, \infty))$$
$$V(\varphi\,\mathbf{U}\,\psi, \alpha) = \max_{i \geq 0}\{\min\{V(\psi, \alpha[i, \infty)), \min_{0 \leq j < i}\{V(\varphi, \alpha[j, \infty))\}\}\}$$

We call $V(\varphi, \alpha)$ the *valuation of φ on α* and say that α *satisfies φ* if $V(\varphi, \alpha) = 1$.

(a) Syntax Tree   (b) Syntax DAG

Fig. 1: Representation of LTL formulas

For a precise representation of LTL formulas, we use syntax trees and syntax DAGs. Syntax DAGs are directed acyclic graphs (DAG) that one obtains from a syntax tree by merging the common subformulas, providing a canonical representation of LTL formulas. Figure 1 illustrates the syntax tree and syntax DAG of the formula $(p\,\mathbf{U}\,\mathbf{G}\,q) \vee \mathbf{F}\,\mathbf{G}\,q$.

The size of an LTL formula $|\varphi|$ is defined as the number of its unique subformulas, which corresponds to the number of nodes in the syntax DAG of φ. For example, the size of the formula $(p\,\mathbf{U}\,\mathbf{G}\,q) \vee \mathbf{F}\,\mathbf{G}\,q$ is 6, as can be seen easily in Figure 1b.

We denote the set of all LTL operators as $\Lambda = \mathcal{P} \cup \Lambda_U \cup \Lambda_B$. Here, the propositions are the nullary operators, $\Lambda_U = \{\neg, \mathbf{X}, \mathbf{F}, \mathbf{G}\}$ are the unary operators and $\Lambda_B = \{\vee, \wedge, \mathbf{U}\}$ are the binary operators of LTL. Further, let $\mathcal{F}_{\mathrm{LTL}}$ denote the set of all LTL formulas.

3 Problem Formulation

Since the problem of *LTL sketching* relies heavily on LTL sketches, we begin with formalizing them first.

LTL Sketch An *LTL sketch* is incomplete LTL formulas in which parts that are difficult to formalize can be left out. The left-out parts are represented using placeholders, denoted by ?'s as can be seen in Figure 2. We comment on the superscripts on the placeholders in a moment.
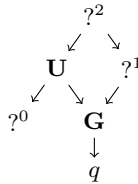

Fig. 2: An LTL sketch

Formally, an LTL sketch $\varphi^?$ is simply an LTL formula whose syntax is augmented with placeholders. The placeholders we allow can be of three types: placeholders of arity zero referred to as Type-0 placeholders, that replace missing LTL formulas; placeholders of arity one referred to as Type-1 placeholders, that replace missing unary operators; and placeholders of arity two referred to as Type-2 placeholders, that replace missing binary operators. In the Figure 2 (and in the rest of the paper), Type-i placeholders are represented using $?^i$.

Given (possibly empty) sets Π^0, Π^1 and Π^2 consisting of Type-0, Type-1 and Type-2 placeholders, respectively, we define LTL-sketches inductively as follows:

- each element of $\mathcal{P} \cup \Pi^0$ is an LTL-sketch; and
- if $\varphi_1^?$ and $\varphi_2^?$ are LTL-sketches, $\circ\,\varphi_1^?$ is an LTL-sketch for $\circ \in \Lambda_U \cup \Pi^1$ and so is $\varphi_1^? \circ \varphi_2^?$ for $\circ \in \Lambda_B \cup \Pi^2$.

Note that an LTL sketch in which $\Pi^0 = \Pi^1 = \Pi^2 = \emptyset$ is simply an LTL formula. Furthur, let $\Pi_{\varphi^?} = \Pi^0 \cup \Pi^1 \cup \Pi^2$ denote the set of all placeholders in an sketch $\varphi^?$. For the sketch in Figure 2, $\Pi_\varphi^? = \{?^0, ?^1, ?^2\}$. For brevity, in the rest of the paper, we refer an LTL sketch as a sketch.

The placeholders are abstract symbols that apriori do not have meaning. To assign meaning to a sketch, we need to substitute all Type-0 placeholders with LTL formulas, all Type-1 placeholders with unary operators and all Type-2 placeholders with binary operators. We do this using a so-called substitution function (or substitution for short).

Formally, a *substitution* function s maps placeholders and operators present in a sketch to LTL operators and LTL formulas in such a way that: $s(?) \in \mathcal{F}_{\mathrm{LTL}}$ if $? \in \Pi^0$; $s(?) \in \Lambda_U$ if $? \in \Pi^1$; $s(?) \in \Lambda_B$ if $? \in \Pi^2$; and $s(\lambda) = \lambda$ for any LTL operator $\lambda \in \Lambda$. Morever, a substitution s is said to be *complete* for an sketch $\varphi^?$ if s is defined for every element in $\Lambda \cup \Pi_{\varphi^?}$ in $\varphi^?$.

A complete substitution s can be applied to a sketch $\varphi^?$ to obtain an LTL formula. To make this mathematically precise, we define a function f_s, which is defined recursively on the structure of $\varphi^?$ as follows:

$$f_s(\varphi_1^? \,?^2\, \varphi_2^?) = f_s(\varphi_1^?) \circ f_s(\varphi_2^?), \text{ where } \circ = s(?^2);$$
$$f_s(?^1\varphi^?)) = \circ f_s(\varphi^?), \text{ where } \circ = s(?^1); \text{ and}$$
$$f_s(?^0) = s(?^0).$$

Input sample While there are many ways to fill a sketch, as alluded in the introduction, we rely on two finite, disjoint sets $P, N \subset (2^{\mathcal{P}})^\omega$ of ultimately periodic words, such that $P \cap N = \emptyset$. The words in P, referred to as the positive examples, represent the desirable system executions that must be allowed by the resulting specification. On the other hand, the words in N, referred to as the negative examples, represent the undesirable system executions that must be disallowed by the resulting specification. We call the pair $\mathcal{S} = (P, N)$ a sample. Since our sample consists of ultimately periodic words, we assume that they are

stored as pairs (u, v) of finite words, where $u \in (2^{\mathcal{P}})^*$ and $v \in (2^{\mathcal{P}})^+$. Moreover, let the size of a sample be $|\mathcal{S}| = \sum_{uv^\omega \in P \cup N} |uv|$.

We say that an LTL formula φ is *consistent* with a sample $\mathcal{S} = (P, N)$ if $V(\varphi, uv^\omega) = 1$ for each $uv^\omega \in P$ (i.e., all positive words satisfy φ) and $V(\varphi, uv^\omega) = 0$ for each $uv^\omega \in N$ (i.e., all negative examples do not satisfy φ);

The LTL sketching problem Having defined the setting, the main problem that we deal with in the paper is the following:

Problem 1 (LTL sketching). Given an LTL sketch $\varphi^?$ and a sample $\mathcal{S} = (P, N)$, find a complete substitution s for $\varphi^?$ such that $f_s(\varphi^?)$ is consistent with $\mathcal{S}$

Before we address the above (and the more general) problem, let us consider a simpler version of the problem:

Problem 2 (LTL learning). Given a sample $\mathcal{S} = (P, N)$, find a minimal LTL formula φ such that φ is consistent with $\mathcal{S}$

The above problem is a restricted version of Problem 1, where the sketch $\varphi^? = ?^0$ (i.e., a single Type-0 placeholder). Recently, it has attracted a lot of attention in the recent years due to its application in Explainable AI and Specification Mining. Along with its theoretical analysis [14], there have been a number of efficient algorithms [30,27,7] to solve it.

For the *LTL learning* problem, one can always find a solution due to the existence of a generic LTL formula that is consistent with a given sample, as indicated by the following remark:

Remark 1. Given sample $\mathcal{S}$, there exists an LTL formula φ such that φ is consistent with $\mathcal{S}$ and $|\varphi|$ is of size $\mathcal{O}(|S|^4)$.

To construct such a formula, one needs to perform the following two steps. First, construct formulas $\varphi_{\alpha,\beta}$ for all $\alpha \in P$ and $\beta \in N$, such that $V(\varphi_{\alpha,\beta}, \alpha) = 1$ and $V(\varphi_{\alpha,\beta}, \beta) = 0$, using a sequence of **X**-operators and an appropriate propositional formula to describe the first symbol where α and β differ. Second, construct the generic consistent formula $\varphi = \bigvee_{\alpha \in P} \bigwedge_{\beta \in N} \varphi_{\alpha,\beta}$.

For the *LTL sketching* problem, however, one may not always find a solution. This is illustrated using the following example.

Example 1. Consider the LTL-sketch $\mathbf{G}(?^0)$ and a sample consisting of a single positive word $\alpha = \{p\}\{q\}^\omega \in P$ and a single negative word $\beta = \{q\}^\omega$. For this LTL-sketch and sample, there does not exist any substitution that leads to an LTL specification consistent with the sample, which can be shown using contradiction. Towards contradiction, assume that there exists a specification $\mathbf{G}(\varphi)$ such that $V(\alpha, \mathbf{G}(\varphi)) = 1$ and $V(\beta, \mathbf{G}(\varphi)) = 0$. Based on the semantics of **G**-operator, $V(\alpha, \mathbf{G}(\varphi)) = V(\alpha[1, \infty), \mathbf{G}(\varphi)) = 1$. On the other hand, $V(\beta, \mathbf{G}(\varphi)) = V(\alpha[1, \infty), \mathbf{G}(\varphi)) = 0$ since $\beta = \alpha[1, \infty)$ leading to a contradiction.

Since there might not exist any complete substitution for a given LTL sketch and a sample, apriori, it is unclear whether the LTL sketching problem is solvable by any terminating algorithm. However, we show in Section 4 that it is indeed possible to find a terminating algorithm for the *LTL sketching* problem. Later in Section 5, we exploit this result to find suitable substitutions to sketches.

4   Existence of a Complete Sketch

To devise a terminating algorithm for the *LTL sketching* problem, we first introduce the related decision problem, which is the following:

Problem 3 (LTL sketch existence). Given an LTL-sketch $\varphi^?$ and a sample $\mathcal{S} = (P, N)$, does there exist a complete substitution s for $\varphi^?$ such that $f_s(\varphi^?)$ is consistent with $\mathcal{S}$.

In this section, we prove that this problem is indeed decidable and belongs to the complexity class NP. Later we also provide a decision procedure based on SAT solving to decide the problem.

4.1   The decidability result

For the decidability result, we introduce some concepts as a preparation. We begin by observing a key property of ultimately periodic words, which is as follows:

Observation 1 *Let $uv^\omega \in (2^\mathcal{P})^\omega$ and φ be an LTL formula. Then, $uv^\omega[|u| + i, \infty) = uv^\omega[|u| + j, \infty)$ for $j \equiv i \mod |v|$. Thus, $V(\varphi, uv^\omega[|u| + i, \infty)) = V(\varphi, uv^\omega[|u| + j, \infty))$.*

Observation 1 indicates that, for a word uv^ω, there exists only a finite number of distinct suffixes of uv^ω, all of which originate in the initial uv portion of uv^ω. We now define the set $suffix(uv^\omega) = \{uv^\omega[i, \infty) \mid 0 \le i < |uv|\}$ of all (possibly) distinct suffixes of a word uv^ω. Moreover, we define $suffix(\mathcal{S}) = \bigcup_{uv^\omega \in (P \cup N)} suffix(uv^\omega)$ to be the set of all suffixes of words in $\mathcal{S}$. Observation 1 further indicates that, to determine the evaluation of an LTL formula φ on an ultimately periodic word uv^ω, it is sufficient to determine its evaluation on the initial $|uv|$ suffixes of uv^ω.

Thus, for a compact representation of the evaluation of φ on uv^ω, we introduce a table notation $T^\varphi_{uv^\omega}$. A table $T^\varphi_{uv^\omega}$ is a $|\varphi| \times |uv|$ matrix that consists of the satisfaction of all the subformulas φ' of φ on the suffixes of $\{uv^\omega\}$. We define the entries of the matrix as follows:

$$T^\varphi_{uv^\omega}[\varphi', t] = V(\varphi', uv^\omega[t, \infty)) \text{ for all subformulas } \varphi' \text{ of } \varphi \text{ and } 0 \le t < |uv|$$

Based on the above definition of the table $T^\varphi_{uv^\omega}$, we identify three properties of these tables. These properties form important building blocks of the decidability proof (i.e., proof of Theorem 2) as we see later.

The first property, which we refer to as *Semantic* property, is that various rows of table are related to each other in a way that reflects the semantics of LTL. We formalize a row of a table using the notation $T^{\varphi}_{uv^{\omega}}[\varphi', \cdot]$, which refers to the row of $T^{\varphi}_{uv^{\omega}}$ corresponding to subformula φ'.

Let us first demonstrate this property on a running example. Consider the formula $\psi = p \vee \mathbf{X}\,q$ and the word $\alpha = \{p,q\}\{p\}\{q\}^{\omega}$. Figure 3 illustrates the table T^{ψ}_{α}. From the figure, one can see that the row $T^{\psi}_{\alpha}[p \vee \mathbf{X}\,q, \cdot]$ corresponds to the bitwise-OR of the rows $T^{\psi}_{\alpha}[p, \cdot]$ and $T^{\psi}_{\alpha}[\mathbf{X}\,q, \cdot]$.

	0	1	2
p	1	1	0
q	1	0	1
$\mathbf{X}\,q$	0	1	1
$p \vee \mathbf{X}\,q$	1	1	1

Fig. 3: Bit-vectors for the subformulas of $p \vee \mathbf{X}\,q$

To formalize the relation between the rows corresponding to different subformulas φ' of φ, we must uniquely identify the subformulas. To this end, given an LTL formula φ, we assign unique identifiers $i \in \{1, \cdots, n\}$ to each node of the syntax DAG of φ. This enables us to denote the subformula of φ rooted at Node i with $\varphi[i]$. For assigning identifiers, we follow the strategy that: we assign the root node with 1; and for every node, we assign an identifier smaller than its children, if it has any. Note that one can analogously assign labels to syntax DAGs of sketches. Figure 4 demonstrates identifiers for the formula $(p\,\mathbf{U}\,\mathbf{G}\,q) \wedge (\mathbf{F}(\mathbf{G}\,q))$. We further define a function $\ell\colon \{1, \cdots, n\} \mapsto \Lambda$ that maps the identifiers to the corresponding operators in the syntax DAG.

(a) Syntax DAG   (b) Labeling

Fig. 4: Labeling for formula $(p\,\mathbf{U}\,\mathbf{G}\,q) \wedge (\mathbf{F}(\mathbf{G}\,q))$.

We now describe the set of equations that formalize the relation between the rows. How a row $T^{\varphi}_{uv^{\omega}}[\varphi[i], \cdot]$ corresponding to Node i relates to the others depends the operator $\ell(i)$ in the root node of $\varphi[i]$. Thus, we list the relation separately for different LTL operators.

If $\ell(i) = p$ for some propostion p, then we have the following relation:

$$T^\varphi_{uv^\omega}[\varphi[i], t] = \begin{cases} 1 \text{ if } p \in uv^\omega[t] \\ 0 \text{ otherwise} \end{cases} \tag{1}$$

If $\ell(i)$ is a unary operator and Node j is the left child of Node i, we have the following relations:

$$\text{if } \ell(i) = \neg: \ T^\varphi_{uv^\omega}[\varphi[i], t] = 1 - T^\varphi_{uv^\omega}[\varphi[j], t] \ \text{ for } 0 \le t < |uv| \tag{2}$$

$$\text{if } \ell(i) = \mathbf{X}: \ T^\varphi_{uv^\omega}[\varphi[i], t] = \begin{cases} T^\varphi_{uv^\omega}[\varphi[j], t+1] \ \text{ for } 0 \le t < |uv| - 1 \\ T^\varphi_{uv^\omega}[\varphi[j], |u|] \ \text{ for } t = |uv| - 1 \end{cases} \tag{3}$$

While Equation 2 simply follows from the semantics of $\neg$-operator, Equation 3 for the $\mathbf{X}$-operator exploits Observation 1 along with its semantics. In particular, the entry $T^\varphi_{uv^\omega}[\varphi[i], |uv| - 1]$ relies on the evaluation of $\varphi[j]$ on $uv^\omega[|u|, \infty)$.

If $\ell(i)$ is a binary operator, and Node j and Node j' are the left and right children of Node i, respectively, then we have the following relation:

$$\text{if } \ell(i) = \vee : T^\varphi_{uv^\omega}[\varphi[i], t] = T^\varphi_{uv^\omega}[\varphi[j], t] \vee T^\varphi_{uv^\omega}[\varphi[j'], t] \text{ for } 0 \le t < |uv| \tag{4}$$

$$\text{if } \ell(i) = \mathbf{U}: \ T^\varphi_{uv^\omega}[\varphi[i], t] = \tag{5}$$

$$\begin{cases} \bigvee_{t \le t'' < |uv|} \left[T^\varphi_{uv^\omega}[\varphi[j'], t''] \wedge \bigwedge_{t \le t' < t''} T^\varphi_{uv^\omega}[\varphi[j], t'] \right] \text{ for } 0 \le t < |u| \\ \bigvee_{|u| \le t'' < |uv|} \left[T^\varphi_{uv^\omega}[\varphi[j'], t''] \wedge \bigwedge_{t' \in t \looparrowright_{u,v} t''} T^\varphi_{uv^\omega}[\varphi[j], t'] \right] \text{ for } |u| \le t < |uv| \end{cases}$$

Again, one can see that Equation 4 follows from the semantics of the $\vee$-operator. Equation 5 for the $\mathbf{U}$-operator consists of two cases: the first case provides the relation for entries $t \in \{0, \cdots, |u| - 1\}$ in the initial part u; the second case covers the entries $t \in \{|u|, \cdots, |uv| - 1\}$ in the periodic part v. Thereby, the second case uses the periodic nature of uv^ω to "loop back" into the periodic part v using the set $t \looparrowright_{u,v} t''$ defined as the follows:

$$t \looparrowright_{u,v} t'' = \begin{cases} \{t, \cdots, t'' - 1\} \text{ if } t < t''; \\ \{|u|, \cdots, t'' - 1, t, \cdots, |uv| - 1\} \text{ if } t'' \le t \end{cases}$$

Having defined the Semantic property, let us now describe the second property, the *Consistency* property. This property ensures that $T^\varphi_{uv^\omega}[\varphi, 0] = 1$ if and only if uv^ω satisfies φ. Thus, for an LTL formula φ consistent with $\mathcal{S}$, we have the following relation:

$$T^\varphi_{uv^\omega}[\varphi, 0] = 1 \text{ for all } uv^\omega \in P, \text{ and } T^\varphi_{uv^\omega}[\varphi, 0] = 0 \text{ for all } uv^\omega \in N \tag{6}$$

The final property we observe is called the *Suffix* property. This property originates from the fact that, LTL, being a future-time logic, has the same evaluation on equal suffixes, i.e., $V(\varphi, u_1 v_1^\omega[t, \infty)) = V(\varphi, u_2 v_2^\omega[t', \infty))$ for $u_1 v_1^\omega[t, \infty) = u_2 v_2^\omega[t', \infty)$. Formally, we state the property as follows:

$$T^\varphi_{u_1 v_1^\omega}[\varphi, t] = T^\varphi_{u_2 v_2^\omega}[\varphi, t'] \text{ for all } u_1 v_1^\omega[t, \infty) = u_2 v_2^\omega[t', \infty) \tag{7}$$

This property becomes significant later for constructing LTL formulas for Type-0 placeholders.

Having set up the prerequisites, we now proceed to provide an NP algorithm to decide the *LTL sketch existence* problem. For ease of presentation, we demonstrate the algorithm on a simple (but crucial) case in which $\varphi^?$ consists of only a single Type-0 placeholder $?^0$. For this case, one might assume that non-deterministically guessing a substitution for the placeholder should suffice. However, apriori one does not know the size of the LTL formula that is necessary to substitute the Type-0 placeholder.

Thus, in our NP algorithm, instead of guessing substitutions, we guess all the entries of the table $T_{uv^\omega}^{\varphi^?}$ for each uv^ω in $\mathcal{S}$. Note that the tables have a finite dimension, precisely $|\varphi^?| \times |uv|$, for each word uv^ω. Thus, the overall process of simply guessing all the table entries can be done in time $\mathcal{O}(\mathsf{poly}(|\varphi^?|, |\mathcal{S}|))$.

After guessing the table entries, we must verify that the guessed tables satisfy the three properties, Semantic, Consistency, and Suffix, discussed earlier in this section. It is easy to verify that checking the first two properties for the tables require time $\mathcal{O}(\mathsf{poly}(|\varphi^?|, |uv|))$ for each uv^ω in $\mathcal{S}$. For checking the Suffix property, one must identify the equal suffixes in $suffix(\mathcal{S})$. This fact that this can be also done in time $\mathsf{poly}(|\mathcal{S}|)$, is a consequence of Lemma 2, stated below. Intuitively, the result states that two suffixes are equal if they are equal only on a finite portion b of size $\mathsf{poly}(|u_1|, |u_2|, |v_1|, |v_2|)$.

Lemma 1. $u_1v_1^\omega[t, \infty) = u_2v_2^\omega[t', \infty)$ *if and only if* $u_1v_1^\omega[t, t+b] = u_2v_2^\omega[t', t'+b)$, *where* $b = \max(|u_1[t, |u_1|)|, |u_2[t', |u_2|)|) + lcm(|v_1|, |v_2|)$.

Proof. First, let us consider $u_1v_1^\omega[t, \infty) = u_2v_2^\omega[t', \infty)$. Clearly, all prefixes of $u_1v_1^\omega[t, \infty)$ and $u_2v_2^\omega[t', \infty)$ are equal, i.e., $u_1v_1^\omega[t, t+b) = u_2v_2^\omega[t', t'+b)$ for all $b \in \mathbb{N}$.

For the other direction, we consider $u_1v_1^\omega[t, t+b) = u_2v_2^\omega[t', t'+b)$, for $b = \max(|u_1[t, |u_1|)|, |u_2[t', |u_2|)|) + lcm(|v_1|, |v_2|)$. Without loss of generality, let us assume that $|u_1[t, |u_1|)| \geq |u_2[t', |u_2|)|$. To avoid clutter of notation, we denote $\mu = |u_1[t, |u_1|)|$ and $\nu = lcm(|v_1|, |v_2|)$. Thus, in this case, $b = \mu + \nu$.

The proof, now, is based on two main observations. First, we begin with the simple observation:

$$u_1v_1^\omega[t, t+\mu) = u_2v_2^\omega[t', t'+\mu); \text{ and}$$
$$u_1v_1^\omega[t+\mu, t+b) = u_2v_2^\omega[t'+\mu, t'+b)$$

Second, we have that

$$(u_1v_1^\omega[t+\mu, t+b))^\omega = v_1^\omega; \text{ and}$$
$$(u_2v_2^\omega[t'+\mu, t'+b))^\omega = (v_2^\omega[t'+\mu, t'+\mu+|v_2|))^\omega$$

The above observation is due to the fact that $u_1v_1^\omega[t+\mu, t+b) = v_1^\kappa$ for $\kappa = \nu/|v_1|$ and $u_2v_2^\omega[t'+\mu, t'+b) = (v_2^\omega[t'+\mu, t'+\mu+|v_2|))^\kappa$ for $\kappa = \nu/|v_2|$.

Now, combining the two observations, we have the following:

$$\begin{aligned}
u_1 v_1^\omega &= u_1 v_1^\omega[t, t+\mu) \cdot (u_1 v_1^\omega[t+\mu, t+b))^\omega \\
&= u_2 v_2^\omega[t', t'+\mu) \cdot (u_2 v_2^\omega[t'+\mu, t'+b))^\omega \\
&= u_2[t', |u_2|) \cdot u_2 v_2^\omega[|u_2|, t'+\mu) \cdot (v_2^\omega[t'+\mu, t'+\mu+|v_2|))^\omega \\
&= u_2 v_2^\omega
\end{aligned}$$

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

We now prove that if the guessed tables satisfy the three properties, then there exists a LTL formula ψ that one can replace in the Type-0 placeholder to obtain a consistent LTL formula. This fact is asserted by the following theorem.

Theorem 1. *Let $\mathcal{S} = (P, N)$ be a sample, $\varphi^?$ be a sketch with only Type-0 placeholders and tables $T_{uv^\omega}^{\varphi^?}$ be $|\varphi^?| \times |uv|$ matrices with $\{0,1\}$ entries for each $uv^\omega \in P \cup N$. Then, the following holds: the tables $T_{uv^\omega}^{\varphi^?}$ satisfy the Semantic, Consistency and Suffix properties if and only if there exists a substitution s such that LTL formula $f_s(\varphi^?)$ is consistent with $\mathcal{S}$.*

Proof. For simplicity, we again consider that $\varphi^?$ consists of only one Type-0 placeholder $?^0$. The proof can be seemlessly extended to multiple Type-0 placeholders.

For the forward direction, we show the existence of the substitution s by explicit construction of an LTL formula for $?^0$. Towards this, we first construct a sample $\mathcal{S}' = (P', N')$ as follows:

$$P' = \{uv^\omega[t, \infty) \in \textit{suffix}(\mathcal{S}) \mid T_{uv^\omega}^{\varphi^?}[?^0, t] = 1, \ uv^\omega \in P \cup N, \ 0 \le i < |uv|\}$$

$$N' = \{uv^\omega[t, \infty) \in \textit{suffix}(\mathcal{S}) \mid T_{uv^\omega}^{\varphi^?}[?^0, t] = 0, \ uv^\omega \in P \cup N, \ 0 \le i < |uv|\}.$$

Since the tables satisfy the Suffix property, we have that $P' \cap N' = \emptyset$. We can now construct the generic LTL formula ψ consistent with $\mathcal{S}'$ based on Remark 1. We claim that this formula ψ can be substituted in $?^0$ to obtain a consistent LTL formula.

Towards this, we first prove that $T_{uv^\omega}^{f_s(\varphi^?)}[\psi, \cdot] = T_{uv^\omega}^{\varphi^?}[?, \cdot]$ for all $uv^\omega \in P \cup N$. To prove this, we exploit two simple observations. First, using the definition of tables, we have $T_{uv^\omega}^{f_s(\varphi^?)}[\psi, t] = V(\psi, uv^\omega[t, \infty))$ for each $uv^\omega[t, \infty) \in \textit{suffix}(\mathcal{S})$. Second, since ψ is consistent with $\mathcal{S}'$, we know $V(\psi, uv^\omega[t, \infty)) = T_{uv^\omega}^{\varphi^?}[?, t]$. Together, we have $T_{uv^\omega}^{f_s(\varphi^?)}[\psi, t] = T_{uv^\omega}^{\varphi^?}[?, t]$.

Next, we prove that $T_{uv^\omega}^{f_s(\varphi^?)}[f_s(\varphi^?)[i], \cdot] = T_{uv^\omega}^{\varphi^?}[\varphi^?[i], \cdot]$ for each $0 \le i < |\varphi^?|$ and word $uv^\omega \in P \cup N$. (Note that we denote the same nodes in $f_s(\varphi^?)$ and $\varphi^?$ using the same identifiers.) Towards contradiction, we assume that there exists some $uv^\omega \in P \cup N$ and some $0 \le i < |\varphi^?|$ and such that $T_{uv^\omega}^{f_s(\varphi^?)}[f_s(\varphi^?)[i], \cdot] \ne T_{uv^\omega}^{\varphi^?}[\varphi^?[i], \cdot]$. Let i^* be the maximum row for which the tables become unequal.

The proof, in general, will proceed via a case analysis on the operator $\ell(i^*)$ labeled in Node i^*. However, since for proof is similar for all the operators, we assume $\ell(i) = \neg$ and Node j^* is the left child of Node i^*. Recall that $j^* > i^*$ based on our assignment of identifiers. Further, based on the Semantic property, $T_{uv^\omega}^{f_s(\varphi^?)}[f_s(\varphi^?)[i^*], t] = 1 - T_{uv^\omega}^{f_s(\varphi^?)}[f_s(\varphi^?)[j^*], t]$ and $T_{uv^\omega}^{\varphi^?}[\varphi^?[i^*], \cdot] = 1 - T_{uv^\omega}^{\varphi^?}[\varphi^?[j^*], \cdot]$ for each $0 \leq t \leq |uv|$ (Equation 2). This implies that $T_{uv^\omega}^{f_s(\varphi^?)}[f_s(\varphi^?)[j^*], \cdot] \neq T_{uv^\omega}^{\varphi^?}[\varphi^?[j^*], \cdot]$, contradicting the maximality of i^*.

Finally, observe that $T_{uv^\omega}^{f_s(\varphi^?)}[f_s(\varphi^?), \cdot] = T_{uv^\omega}^{\varphi^?}[\varphi^?, \cdot]$. As a consequence, since tables $T_{uv^\omega}^{\varphi^?}$ satisfy the Consistency property, so does tables $T_{uv^\omega}^{f_s(\varphi^?)}$. This implies that $f_s(\varphi^?)$ is consistent with $\mathcal{S}$.

For the other direction, we construct tables $T_{uv^\omega}^{\varphi^?}[\psi, \cdot]$ based on the tables $T_{uv^\omega}^{f_s(\varphi^?)}[\psi, \cdot]$. In particular, we have $T_{uv^\omega}^{f_s(\varphi^?)}[\psi, \cdot] = T_{uv^\omega}^{\varphi^?}[\psi, \cdot]$ for each $0 \leq i < |\varphi^?|$ and $uv^\omega \in P \cup N$. Since tables $T_{uv^\omega}^{f_s(\varphi^?)}[\psi, \cdot]$ satisfy the Semantic, the Consistency and the Suffix properties, so does the tables $T_{uv^\omega}^{\varphi^?}[\psi, \cdot]$. $\qquad\square$

With this, we conclude the NP algorithm for the case where $\varphi^?$ only has one Type-0 placeholder. We can easily extend the algorithm to the case where $\varphi^?$ additionally consists of Type-1 and Type-2 placeholders. In particular, we first guess the operators to be substituted for the Type-1 and Type-2 placeholders and substitute them. We then obtain a sketch consisting of only Type-0 placeholders. We can now apply our algorithm that relies on guessing tables, as described above. In total, we obtain the following result:

Theorem 2. *The LTL sketch existence problem is in* NP.

While we determine the complexity upper bound of the *LTL sketch existence* problem to be NP, the lower bound is open and adds to the list of open problems in the area of LTL inference (see Section 6 of [14]).

4.2   The decision procedure

Based on the NP algorithm described above, we now devise a decision procedure to decide the *LTL sketch existence* problem. The decision procedure relies upon solving an instance of SAT problem to check if there exists suitable tables $T_{uv^\omega}^{\varphi}$ which satisfy the three properties, discussed in Section 4.1.

To encode entries of the tables, we first introduce the following variables: $y_{i,t}^{u,v}$ for each $i \in \{1, \cdots n\}$, $t \in \{0, \cdots, |uv| - 1\}$, and $uv^\omega \in P \cup N$. A variable $y_{i,t}^{u,v}$ encodes the entry $T_{uv^\omega}^{\varphi}[\varphi[i], t]$. Further, to encode the operators to be substituted for the Type-1 and Type-2 placeholders in $\varphi^?$, we have the following variables: (i) $x_{i,\lambda}$ for each Node i where $\ell(i)$ is a Type-1 placeholder and each $\lambda \in \Lambda_U$; and (ii) $x_{i,\lambda}$ for each Node i where $\ell(i)$ is a Type-2 placeholder and each $\lambda \in \Lambda_B$

We now impose constraints on the introduced variables to ensure that they encode tables that satisfy the properties for inferring a consistent LTL formula.

We achieve this by constructing a propositional formula $\Phi^{\varphi^?,\mathcal{S}}$ using the introduced variables. This formula ensures that the variables $y_{i,t}^{u,v}$ encode tables that satisfy the three required properties.

Since $\Phi^{\varphi^?,\mathcal{S}}$ ensures the existence of a suitable tables, we have the following property of $\Phi^{\varphi^?,\mathcal{S}}$: there exists a complete substitution s for $\varphi^?$ such that $f_s(\varphi^?)$ is consistent with $\mathcal{S}$ if and only if $\Phi^{\varphi^?,\mathcal{S}}$ is satisfiable. We can now simply check the satisfiability of $\Phi^{\varphi^?,\mathcal{S}}$ using an off-the-shelf SAT solver to determine the existence of a complete substitution.

Internally, $\Phi^{\varphi^?,\mathcal{S}} = \Phi_?^{1,2} \wedge \Phi_{sem} \wedge \Phi_{con} \wedge \Phi_{suf}$ is a conjunction of four formulas. The first conjunct $\Phi_?^{1,2}$ ensures that the Type-1 and Type-2 placeholders are substituted by appropriate operators. The conjuncts Φ_{sem}, Φ_{con} and Φ_{suf} ensure that the variables $y_{i,t}^{u,v}$ encode entries of tables that satisfy the Semantic property (Equations 2 to 5), the Consistency property (Equation 6) and the Suffix property (Equation 7), respectively. We now describe briefly how each of these formulas are constructed.

Constraints for $\Phi_?^{1,2}$. For each Node i labeled with a Type-1 placeholder (i.e., $\ell(i) \in \Pi^1$), we have the following constraint:

$$\left[\bigvee_{\lambda \in \Lambda_U} x_{i,\lambda} \right] \wedge \left[\bigwedge_{\lambda \neq \lambda' \in \Lambda_U} \neg x_{i,\lambda} \vee \neg x_{i,\lambda'} \right], \tag{8}$$

which ensures that the Type-1 placeholder is substituted by a unique unary operator. For Type-2 placeholders, we have the exact same constraint except that the operators range from the set of binary operators Λ_B. Now, we construct $\Phi_?^{1,2}$ as the conjunction of all such constraints for the nodes labeled with Type-1 and Type-2 placeholders.

Constraints for Φ_{sem}. We define $\Phi_{sem} = \bigwedge_{uv^\omega \in P \cup N} \Phi^{u,v}$, where $\Phi^{u,v}$ is a propositional formula that ensures that the variables $y_{i,t}^{u,v}$ satisfy the Equations 2 to 5 for word uv^ω in $\mathcal{S}$. In $\Phi^{u,v}$, for instance, for each Node i labeled with $\mathbf{X}$-operator (i.e, $\ell(i) = \mathbf{X}$) and has Node j as its left child, we have the following constraint:

$$\left[\bigwedge_{0 \leq t \leq |uv|-1} \left[y_{i,t}^{u,v} \leftrightarrow y_{j,t+1}^{u,v} \right] \right] \wedge \left[y_{i,|uv|-1}^{u,v} \leftrightarrow y_{j,|u|}^{u,v} \right] \tag{9}$$

This constraints ensures that the variables $y_{i,t}^{u,v}$ satisfy Equation 3 for the word uv^ω. We construct similar constraints for the other operators based on their corresponding semantic relation.

For nodes labeled with Type-1 and Type-2 placeholders, we additionally rely on variables $x_{i,\lambda}$ to determine the operator λ to be substituted in Node i. For instance, for each Node i labeled with a Type-1 placeholder (i.e., $\ell(i) \in \Pi^1$) that has Node j as its left child, we have the following constraint:

$$x_{i,\mathbf{X}} \rightarrow \left[\bigwedge_{0 \leq t \leq |uv|-1} \left[y_{i,t}^{u,v} \leftrightarrow y_{j,t+1}^{u,v} \right] \right] \wedge \left[y_{i,|uv|-1}^{u,v} \leftrightarrow y_{j,|u|}^{u,v} \right] \tag{10}$$

The above constraint states that if Node i is substituted with a $\mathbf{X}$-operator (i.e., if $x_{i,\mathbf{X}}$ is true), the constraint on the $y_{i,t}^{u,v}$ variables is based on Equation 3. We construct similar constraints for when $x_{i,\lambda}$ is true for some other operator λ based on the corresponding equation for λ. Finally, we construct $\Phi^{u,v}$ as the conjunction of all such constraints that rely on the semantic equations.

Constraints for Φ_{sem}. We construct Φ_{con} as follows:

$$\Phi_{con} = \Big[\bigwedge_{uv^\omega \in P \cup N} \Phi^{u,v} \Big] \wedge \Big[\bigwedge_{uv^\omega \in P} y_{1,0}^{u,v} \wedge \bigwedge_{uv^\omega \in N} \neg y_{1,0}^{u,v} \Big] \qquad (11)$$

This formula ensures that Equation 6 is satisfied for the prospective bit-vectors, in addition to ensuring $\Phi^{u,v}$ holds for all uv^ω in $\mathcal{S}$.

Constraints for Φ_{suf}. In Φ_{suf}, for each Node i labeled with a Type-0 placeholder (i.e., $\ell(i) \in \Pi^0$), we have the following constraint:

$$\bigwedge_{u_1 v_1^\omega [t,\infty) = u_2 v_2^\omega [t',\infty) \in suffix(\mathcal{S})} \Big[y_{i,t}^{u_1,v_1} \leftrightarrow y_{j,t'}^{u_2,v_2} \Big], \qquad (12)$$

which ensures that the variables $y_{i,t}^{u,v}$ the entries of the tables satisfy Equation 7

The following lemma establishes the correctness of the decision procedure.

Lemma 2. *Let $\varphi^?$ be a sketch, $\mathcal{S}$ a sample and $\Phi^{\varphi^?,\mathcal{S}}$ the formula as defined above. Then, $\Phi^{\varphi^?,\mathcal{S}}$ is satisfiable if and only if there exists a complete substitution s such that $f_s(\varphi^?)$ is consistent with $\mathcal{S}$.*

Proof. If $\Phi^{\varphi^?,\mathcal{S}}$ is satisfiable, we can construct a suitable complete substitution s based on the model v of $\Phi^{\varphi^?,\mathcal{S}}$. To begin with, for Node i labeled with a Type-1 or Type-2 placeholder ?, we have $s(?) = \lambda$, where λ is the unique operator for which $v(x_{i,\lambda}) = 1$. Next, for nodes labeled with Type-0 placeholders, we first construct tables $T_{uv^\omega}^{\varphi^?}$ for each $uv^\omega \in P \cup N$ using the value of the variables $v(y_{i,t}^{u,v})$. Based on the construction of $\Phi^{\varphi^?,\mathcal{S}}$, all of these tables satisfy the Semantic, Consistency and Suffix properties. Now, we use Lemma 2 to find substitutions to Type-0 placeholders using the tables $T_{uv^\omega}^{\varphi^?}$.

For the other direction, we construct a satisfying assignment v using the substitution function s and tables $T_{uv^\omega}^{f_s(\varphi)}$ for $uv^\omega \in P \cup N$. First, we assign $v(x_{i,\lambda}) = 1$ if and only if $s(?) = \lambda$ for a Node i labeled with a Type-1 or Type-2 placeholder ?. Second, we assign $v(y_{i,t}^{u,v}) = T_{uv^\omega}^{f_s(\varphi)}[f_s(\varphi^?)[i],t]$ for each $uv^\omega \in P \cup N$ and $0 \le t \le |uv|$. This assignment v satisfies $\Phi_?^{1,2}$, since we obtain v from the syntax DAG of a valid LTL formula. Further, this assignment satisfies Φ_{sem}, Φ_{con} and Φ_{suf} because the tables satisfy Semantic, Consistency and Suffix properties, respectively, on which the constraints are based. Overall, v is a model for $\Phi^{\varphi^?,\mathcal{S}}$. $\qquad \square$

Finally, we have the following remark to assess the size of the encoding $\Phi^{\varphi^?,\mathcal{S}}$.

Remark 2. Let $n = |\varphi^?|$ and $m = \sum_{uv^\omega \in P \cup N} |uv|$. The formula $\Phi^{\varphi^?,\mathcal{S}}$ ranges over $\mathcal{O}(n + nm)$ variables and is of size $\mathcal{O}(n + nm^3 + m^2)$.

5   Algorithms to complete an LTL sketch

In this section, we describe two novel algorithms for solving the *LTL sketching* problem, which aim at searching for concise LTL formulas from sketches. Our rationale behind finding concise specifications is motivated by two considerations: first, they are more human understandable and thus, easier for the engineers to interpret; second, potentially all verification and synthesis algorithms perform better with smaller specifications. To this end, our first algorithm relies on existing techniques for learning *minimal* LTL formulas. Our second algorithm, alternatively, searches for formulas of increasing size using constraint solving.

5.1   Algorithm based on LTL learning

Our first algorithm for finding concise LTL formulas builds upon the decision procedure for checking the existence of a complete substitution presented in Section 4.2. In particular, we exploit the formula $\Phi^{\varphi^?,\mathcal{S}}$ to reduce the *LTL sketching* problem to a number of instances of the *LTL learning* problem, one for each Type-0 placeholder. One can then exploit recent advancements [27,31] in solving the *LTL learning* problem to find concise ways of filling out the sketch.

Algorithm 1 outlines the details of the algorithm. The first step is to construct $\Phi^{\varphi^?,\mathcal{S}}$ from the given sample and sketch, as described in Section 4.2. If $\Phi^{\varphi^?,\mathcal{S}}$ is unsatisfiable, the algorithm straight-away returns that no solution exists (Line 12), as established by Theorem 4. If it is satisfiable, we use a model (say v) of $\Phi^{\varphi^?,\mathcal{S}}$, which can be obtained from any off-the-shelf SAT solver, to fill out $\varphi^?$. We now describe in detail how to fill out a sketch.

Given a model v of $\Phi^{\varphi^?,\mathcal{S}}$, in a rather straightforward manner, one can substitute the Type-1 and Type-2 placeholders in $\varphi^?$ (Line 4). For each Node i where $\ell(i)$ is a Type-1 and Type-2 placeholders, we assign $s(\ell(i)) = \lambda$ where, λ is the unique operator for which $v(x_{i,\lambda}) = 1$.

The Type-0 placeholders, however, are more challenging to substitute. This is because, they represent entire LTL formulas. Towards substituting Type-0 placeholders, for every Node i where $\ell(i)$ is a Type-0 placeholder (i.e., $\ell(i) \in \Pi^0$), we first construct a sample $\mathcal{S}_i = (P_i, N_i)$ (Line 6) as follows:

$$P_i = \{uv^\omega[t, \infty) \in \mathit{suffix}(\mathcal{S}) \mid v(y_{1,t}^{u,v}) = 1\}, \text{ and}$$
$$N_i = \{uv^\omega[t, \infty) \in \mathit{suffix}(\mathcal{S}) \mid v(y_{1,t}^{u,v}) = 0\}.$$

Now, we learn a concise LTL formula φ_i consistent with the sample $\mathcal{S}_i$ (using e.g., algorithms by Neider and Gavran [27]) for substituting $\ell(i)$ (Line 7).

The correctness of the algorithm is established by the following theorem:

Theorem 3. *Given sketch $\varphi^?$ and sample $\mathcal{S}$, Algorithm 1 terminates and completes $\varphi^?$ to output an LTL formula that is consistent with $\mathcal{S}$, if such a formula exists.*

Algorithm 1 Algorithm based on LTL learning

 1: **Input:** Sketch $\varphi^?$, Sample $\mathcal{S}$
 2: Construct $\Phi^{\varphi^?,\mathcal{S}} = \Phi^{1,2}_? \wedge \Phi_{sem} \wedge \Phi_{con} \wedge \Phi_{suf}$
 3: **if** $\Phi^{\varphi^?,\mathcal{S}}$ is satisfiable (say with model v) **then**
 4:      Substitute Type-1 and Type-2 placeholders in $\varphi^?$ using v
 5:      **for** every i such that $\ell(i) \in \Pi^0$ **do**
 6:          Construct $\mathcal{S}_i = (P_i, N_i)$
 7:          $\Phi_i \leftarrow Learn(\mathcal{S}_i)$
 8:          Substitute Node i with Φ_i in t
 9:      **end for**
10:      **return** $\varphi^?$
11: **else**
12:      **return** LTL formula does not exist
13: **end if**

Observe that Algorithm 1 constructs new samples for each Type-0 placeholder. Each of these samples have size $\mathcal{O}(|suffix(\mathcal{S})|) = \mathcal{O}(|\mathcal{S}|^2)$, which poses a challenge to the scalability of this algorithm. Furthermore, the new samples are not optimized to produce the minimal possible substitutions. Our next algorithm improves both on the runtime and the size of the inferred specification.

5.2   Algorithm based on incremental SAT solving

We, now, introduce an algorithm that is based on a series of SAT solving problems, similar to the SAT-based algorithm by Neider and Gavran [27]. Given a sample $\mathcal{S}$ and a natural number $n \in \mathbb{N}\backslash\{0\}$, we construct a propositional formula $\Psi_n^{\varphi^?,\mathcal{S}}$ of size $\mathsf{poly}(n, |\mathcal{S}|)$ that has the following two properties:

- $\Psi_n^{\varphi^?,\mathcal{S}}$ is satisfiable if and only if one can complete $\varphi^?$ to obtain an LTL formula of size n that is consistent with $\mathcal{S}$; and
- if v is a model of $\Psi_n^{\varphi^?,\mathcal{S}}$, then v contains sufficient information to complete $\varphi^?$ to construct an LTL formula φ of size n that is consistent with $\mathcal{S}$.

However, in contrast to the algorithms by Neider and Gavran, we first solve $\Phi^{\varphi^?,\mathcal{S}}$ (discussed in Section 4.1) to determine the existence of a complete substitution. If and only if $\Phi^{\varphi^?,\mathcal{S}}$ is satisfiable, our algorithm checks the satisfiability of $\Psi_n^{\varphi^?,\mathcal{S}}$ for increasing values of n (starting from 1) to search for an LTL formula of size at most n that has the same syntactic structure as $\varphi^?$. We construct the resulting LTL formula by substituting the placeholders in $\varphi^?$ based on a model v of the satisfiable formula $\Psi_n^{\varphi^?,\mathcal{S}}$. This idea is illustrated in Algorithm 2.

On a technical level, the formula $\Psi_n^{\varphi^?,\mathcal{S}}$ is obtained by modifying certain parts of the formula $\Phi^{\varphi^?,\mathcal{S}}$. Precisely, $\Psi_n^{\varphi^?,\mathcal{S}} = \Phi^{1,2}_? \wedge \Phi'_{sem} \wedge \Phi_{con} \wedge \Phi^0_{?,n}$ employs a two modifications: a new formula $\Phi^0_{?,n}$ replaces Φ_{suf}; and Φ'_{sem} is obtained by adding more constraints to Φ_{sem}. The formula $\Phi^0_{?,n}$ encodes the structure LTL formulas

---

Algorithm 2 Algorithm based on incremental SAT solving

1: **Input:** Sketch $\varphi^?$, Sample $\mathcal{S}$
2: Construct $\Phi^{\varphi^?,\mathcal{S}} = \Phi^{1,2}_? \wedge \Phi_{sem} \wedge \Phi_{con} \wedge \Phi_{suf}$
3: **if** $\Phi^{\varphi^?,\mathcal{S}}$ is satisfiable **then**
4:     $n \leftarrow |\varphi^?| - 1$
5:     **repeat**
6:         $n \leftarrow n + 1$
7:         Construct $\Psi^{\varphi^?,\mathcal{S}}_n = \Phi^{1,2}_? \wedge \Phi'_{sem} \wedge \Phi_{con} \wedge \Phi^0_{?,n}$
8:     **until** $\Psi^{\varphi^?,\mathcal{S}}_n$ is satisfiable (say with model v)
9:     Substitute placeholders in t using v
10:      **return** $\varphi^?$
11: **else**
12:     **return** LTL formula does not exist
13: **end if**

---

that substitute the Type-0 placeholders. Φ'_{sem}, again as in Φ_{sem}, ensures that the variables $y^{u,v}_{i,t}$ encode table entries $T^{\varphi}_{uv^{\omega}}[\varphi[i],]$ that satisfy Equations 2 to 6. We now briefly describe the constraints for the newly introduced formulas.

Constraints for $\Phi^0_{?,n}$. The constraints here rely on an additional set of variables: (i) $x_{i,\lambda}$ for each Node i labeled with a Type-0 placeholder, and for each $i \in \{|\varphi^?| + 1, \cdots, n\}$ such that $\ell(i)$ is a Type-0 placeholder, and each $\lambda \in \Lambda$; and (ii) $l_{i,j}$ and $r_{i,j}$ for each Node i labeled with a Type-0 placeholder, and each $i \in \{|\varphi^?| + 1, \cdots, n\}$ and $j \in \{\max(i, |t|), \cdots, n\}$. The variable $x_{i,\lambda}$, again, encode that Node i is labeled with λ. The variables $l_{i,j}$ (and $r_{i,j}$) encode that the left child (and the right child) of Node i is Node j. Together the new variables encode the structure of the prospective LTL formulas for Type-0 placeholders.

We now impose constraints, similar to Constraint 8, on the variables $x_{i,\lambda}$ to ensure each node is labeled by a unique LTL operator from Λ. Additionally, we ensure that each Node i has a unique left child using the following constraint:

$$\left[\bigvee_{\max(i,|t|) \leq j \leq n} l_{i,j} \right] \wedge \left[\bigwedge_{\max(i,|t|) \leq j \neq j' \leq n} \neg l_{i,j} \vee \neg l_{i,j'} \right], \tag{13}$$

We have a similar constraint to ensure the uniqueness of right child of a node. Now, we construct Ψ_{str} as the conjunction of all such structural constraints.

Constraints for Φ'_{sem}. We rely on new variables $y^{u,v}_{i,t}$ for each Node i labeled with a Type-0 variables and each $i \in \{|\varphi^?| + 1, n\}$, each $t \in \{0, \cdots, |uv| - 1\}$ and each uv^{ω} in $\mathcal{S}$. On these variables, we impose the following constraint, which is similar to Constraint 10:

$$\left[x_{i,\mathbf{X}} \wedge l_{i,j} \right] \rightarrow \bigwedge_{0 \leq t \leq |uv| - 1} \left[y^{u,v}_{i,t} \leftrightarrow y^{u,v}_{j,t+1} \right] \wedge \left[y^{u,v}_{i,|uv|-1} \leftrightarrow y^{u,v}_{j,|u|} \right], \tag{14}$$

that ensures that the $y_{i,t}^{u;v}$ variables encode entries of table that satisfy Equation 3. We construct Φ'_{sem} as the conjunction of Φ_{sem} and the new semantic constraints.
    We, now, establish the correctness of Algorithm 2 using the following theorem:

Theorem 4. *Given sketch $\varphi^?$ and sample $\mathcal{S}$, Algorithm 2 terminates and completes $\varphi^?$ to output an LTL formula that is consistent with $\mathcal{S}$, if such a formula exists.*

While Algorithm 2 optimizes for the size of the inferred specification, it may not always return an minimal LTL formula. Example 2 demonstrates one such situation where Algorithm 2 could produce a sub-optimal result.

Example 2. Consider the sketch $\varphi^? =?_0 \vee \mathbf{X}\,\mathbf{X}\,p$ (Figure 5a) and sample $\mathcal{S}$ consisting of positive examples $\{\}\{p\}\{\}^\omega$ and $\{\}\{\}\{p\}\{\}^\omega$ and a negative example $\{\}^\omega$. For this input, a possible output by Algorithm 2 is the formula $\mathbf{F}\,p \vee \mathbf{X}\,\mathbf{X}\,p$ (Figure 5b). The minimal consistent formula $\mathbf{X}\,p \vee \mathbf{X}\,\mathbf{X}\,p$ (Figure 5c), however, is smaller.

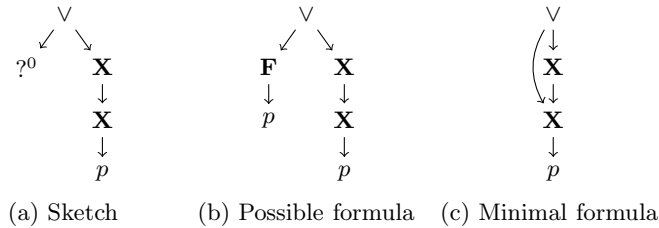(a) Sketch   (b) Possible formula   (c) Minimal formula

Fig. 5: Minimal formulas require sharing of nodes

We leave the challenging problem of devising algorithms for searching for minimal LTL formulas from a given sketch as a part of future work.

6   Experimental evaluation

In this section, we design experiments to answer the following research question: how do the presented algorithms perform in terms of their running times and the size of inferred specification? Can the algorithms recover the specification intended by the engineer?
    To answer these questions, we perform a comparative study of our algorithms (presented in Section 5). Note that, to the best of our knowledge, none of the existing tools can solve the *LTL sketch existence* problem the newly generated sample from which they learn is not ensured to produce the minimal formula.in
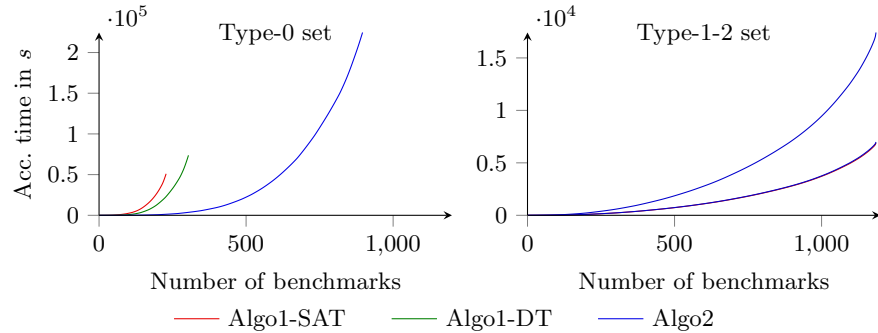
Fig. 6: Accumulated runtime of the algorithms on the two benchmark sets

its full generality. Thus, we restrict ourselves to comparisons of the presented algorithms.

We have implemented a prototype of our algorithms in a tool called `LTL-Sketcher` which is publicly available[3]. While Algorithm 1 can exploit any LTL learning algorithm, we have chosen the algorithms presented by Neider and Gavran [27] in their state-of-the-art LTL learning tool—`Flie`—one based on SAT solving and one on decision tree learning. We refer to Algorithm 1 relying on the SAT-based algorithm as Algo1-SAT and the one relying on decision tree learning as Algo1-DT. We refer to Algorithm 2 as Algo2.

For generating our benchmarks, we used the synthetic benchmark set presented by Neider and Gavran. Their benchmark set consists of 1196 samples generated from 12 different LTL properties, which originate from a study by Dwyer et al. [11]. Based on these samples, we generated two benchmark sets, which we refer to as the Type-0 set and the Type-1-2 set. Each benchmark in these two sets is simply a pair $(\varphi^?, \mathcal{S})$ consisting of a sample and a sketch, which forms the input to our presented algorithms.

For generating a benchmark $(\varphi^?, \mathcal{S})$ in Type-0 set, we choose a sample $\mathcal{S}$ generated from LTL formula φ from the 1196 chosen samples; then construct a sketch $\varphi^?$ from φ by substituting an aritrarily chosen subformula of size at least $\lfloor |\varphi|/2 \rfloor$ by a Type-0 placeholder. In this manner, we obtain 1196 pairs $(\varphi^?, \mathcal{S})$ for the Type-0 set. We repeat the same process for the Type-1-2 set, except that we construct a sketch $\varphi^?$ from φ by substituting one aritrarily chosen operator by a Type-1 placeholder, if unary, otherwise with a Type-2 placeholder. Here again we generate 1196 pairs $(\varphi^?, \mathcal{S})$ for the Type-1-2 set.

All the experiments are conducted on a single core of a Debian machine with Intel Xeon E7-8857 CPU (at 3 GHz) using up to 6 GB of RAM. The timeout was set to be 1200 s for the run of each algorithm on each benchmark.

Comparison of runtime First, we compare the three algorithms based on their runtimes on the Type-0 and the Type-1-2 benchmark sets. Figure 6 presents the accumulated runtimes of the algorithms on the two benchmark sets. From the

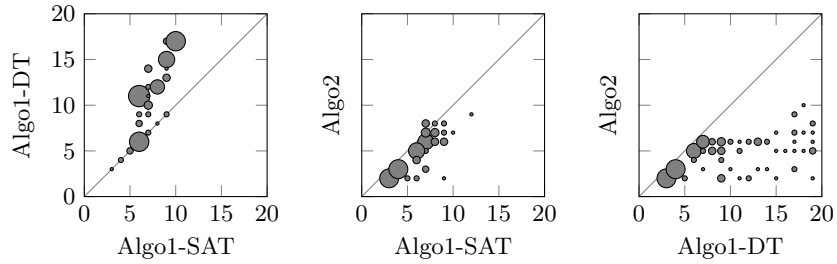---

[3] `https://github.com/rajarshi008/LTL-sketcher`

Fig. 7: Comparison of sizes of resulting LTL formula on Type-0 set.

figure, we observe that the runtimes of Algo1-SAT and Algo1-DT is higher than that of Algo2 on the Type-0 set. In fact, Algo1-SAT and Algo1-DT experiences 957 and 881 time-outs, respectively, while, Algo2 only 289 time-outs. The superior performance of Algo2 can be attributed to the fact that Algo1-SAT and Algo1-DT perform LTL learning on a newly generated sample $\mathcal{S}'$ (as discussed in Section 5.1) of size $\mathcal{O}(|\mathcal{S}|^2)$, which turns out to be an expensive procedure in most cases.

On the Type-1-2 set, however, Algo1-SAT and Algo1-DT display better runtime performance. This is because for substituting Type-1 or Type-2 placeholders only, Algo1-SAT and Algo1-DT do not generate new samples. Rather, they only solve $\Phi^{\varphi^?,\mathcal{S}}$ using a SAT-solver. Algo2, on the other hand, additionally performs at least one iteration of its incremental SAT solving procedure. This observation presents a possible future optimization of Algo2 to adapt based on the type of placeholder present in the sketch.

Comparison of formula sizes We now compare the three algorithms based on the size of the resulting LTL formulas on the Type-0 set. Figure 7 presents a pair-wise comparison of the three algorithms in terms of size. We notice that Algo2 produces the smallest formulas among all the algorithms. This is because, Algo2 iteratively searches for consistent formulas of increasing size. While Algo1-SAT and Algo1-DT also try to optimize the size, the newly generated sample from which they learn is not ensured to produce the minimal formula.

Comparison of formula recovery Finally, we compare the three algorithms based on whether they were able to recover the intended formula. In our experimental setup, we consider the intended formula to be the one that was used to generate the input sample. Table 1 illustrates the percentage of runs in which the intended formula was recovered by the three algorithms on Type-0 and Type-1-2 benchmark sets. On Type-0 benchmark set, both Algo1-SAT and Algo1-DT almost never produced the intended formula, while Algo2 produced the original formula in several runs. A possible explanation can be that Algo2 searches for simple completions first, finding the intended formula often. For Algo1, however, the intermediate samples generated will not be optimized to obtain the intended formulas. On Type-1-2 benchmark set, on the other hand, all the algorithms

recover the intended formulas in most runs. This is because, the intermediate samples do not play a role in this benchmark set and thus, all the algorithms have a similar performance.

	Type-0	Type-1-2
Algo1-SAT	0.1	97.4
Algo1-DT	0	97.4
Algo2	54.9	97.9

Table 1: Percentage of runs in which intended formula is recovered

Overall, we conclude that Algorithm 2 performs better, in terms of runtime, size and recovery of specifications, even when Algorithm 1 uses state-of-the-art LTL learning techniques.

7   Conclusion and Future Work

We have introduced LTL sketching, a novel way of writing formal specifications in LTL. The key idea is that a user can write a partial specification, i.e., a sketch, which is then completed based on given examples of desired and undesired system behavior. We have shown that the sketching problem is in NP and presented two SAT-based sketching algorithms. Our experimental evaluation has shown that our algorithms can effectively complete sketches consisting of different types of missing information.

A natural direction for future work is to lift the idea of specification sketching to other specification languages, such as Signal Temporal Logic (STL) [26], the Property Specification Language (PSL) [12], or Computation Tree Logic (CTL) [8]. We also plan to investigate how specification sketching can be applied to visual specifications, including UML (high-level) message sequence charts [17]. Moreover, we intend to extend the notion of sketching beyond the use of examples to fill out placeholders (e.g., by allowing the engineer to constraint placeholders using simple logical formulas or regular expressions).

References

1. Arif, M.F., Larraz, D., Echeverria, M., Reynolds, A., Chowdhury, O., Tinelli, C.: SYSLITE: syntax-guided synthesis of PLTL formulas from finite traces. In: FM-CAD. pp. 93–103. IEEE (2020)
2. Bacherini, S., Fantechi, A., Tempestini, M., Zingoni, N.: A story about formal methods adoption by a railway signaling manufacturer. In: FM. Lecture Notes in Computer Science, vol. 4085, pp. 179–189. Springer (2006)

3. Badeau, F., Amelot, A.: Using B as a high level programming language in an industrial project: Roissy VAL. In: ZB. Lecture Notes in Computer Science, vol. 3455, pp. 334–354. Springer (2005)
4. Balyo, T., Heule, M.J.H., Järvisalo, M.: SAT competition 2016: Recent developments. In: 31st AAAI Conference on Artificial Intelligence, AAAI '17. pp. 5061–5063. AAAI Press (2017)
5. Bowen, J.P.: Gerard o'regan: Concise guide to formal methods: Theory, fundamentals and industry applications. Formal Aspects Comput. **32**(1), 147–148 (2020)
6. Camacho, A., Icarte, R.T., Klassen, T.Q., Valenzano, R.A., McIlraith, S.A.: LTL and beyond: Formal languages for reward function specification in reinforcement learning. In: Kraus, S. (ed.) Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019. pp. 6065–6073. ijcai.org (2019). https://doi.org/10.24963/ijcai.2019/840, `https://doi.org/10.24963/ijcai.2019/840`
7. Camacho, A., McIlraith, S.A.: Learning interpretable models expressed in linear temporal logic. In: ICAPS. pp. 621–630. AAAI Press (2019)
8. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching-time temporal logic. In: Logics of Programs, Workshop, Yorktown Heights, New York, USA, May 1981. Lecture Notes in Computer Science, vol. 131, pp. 52–71. Springer (1981). https://doi.org/10.1007/BFb0025774
9. Cofer, D.D., Miller, S.P.: DO-333 certification case studies. In: NASA Formal Methods. Lecture Notes in Computer Science, vol. 8430, pp. 1–15. Springer (2014)
10. Courtois, P.J., Seidel, F., Gallardo, F., Bowell, M.: Licensing of safety critical software for nuclear reactors. common position of international nuclear regulators and authorised technical support organisations. (12 2015). https://doi.org/10.13140/RG.2.1.2789.8968
11. Dwyer, M.B., Avrunin, G.S., Corbett, J.C.: Property specification patterns for finite-state verification. In: FMSP. pp. 7–15. ACM (1998)
12. Eisner, C., Fisman, D.: A Practical Introduction to PSL. Series on Integrated Circuits and Systems, Springer (2006). https://doi.org/10.1007/978-0-387-36123-9
13. Fecko, M.A., Uyar, M.Ü., Amer, P.D., Sethi, A.S., Dzik, T., Menell, R., McMahon, M.: A success story of formal description techniques: Estelle specification and test generation for MIL-STD 188-220. Comput. Commun. **23**(12), 1196–1213 (2000)
14. Fijalkow, N., Lagarde, G.: The complexity of learning linear temporal formulas from examples. CoRR **abs/2102.00876** (2021)
15. Fix, L.: Fifteen years of formal property verification in intel. In: 25 Years of Model Checking. Lecture Notes in Computer Science, vol. 5000, pp. 139–144. Springer (2008)
16. Gario, M., Cimatti, A., Mattarei, C., Tonetta, S., Rozier, K.Y.: Model checking at scale: Automated air traffic control design space exploration. In: CAV (2). Lecture Notes in Computer Science, vol. 9780, pp. 3–22. Springer (2016)
17. Harel, D., Thiagarajan, P.S.: Message sequence charts. In: UML for Real - Design of Embedded Real-Time Systems, pp. 77–105. Kluwer (2003). https://doi.org/10.1007/0-306-48738-1_4
18. Holzmann, G.J.: The model checker SPIN. IEEE Trans. Software Eng. **23**(5), 279–295 (1997)
19. Holzmann, G.J.: The logic of bugs. In: SIGSOFT FSE. pp. 81–87. ACM (2002)
20. Kim, J., Muise, C., Shah, A., Agarwal, S., Shah, J.: Bayesian inference of linear temporal logic specifications for contrastive explanations. In: IJCAI. pp. 5591–5598. ijcai.org (2019)

21. Klein, G., Andronick, J., Elphinstone, K., Heiser, G., Cock, D., Derrin, P., Elka-duwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T., Tuch, H., Winwood, S.: sel4: formal verification of an operating-system kernel. Commun. ACM **53**(6), 107–115 (2010)
22. Lemieux, C., Beschastnikh, I.: Investigating program behavior using the texada LTL specifications miner. In: ASE. pp. 870–875. IEEE Computer Society (2015)
23. Lemieux, C., Park, D., Beschastnikh, I.: General LTL specification mining (T). In: ASE. pp. 81–92. IEEE Computer Society (2015)
24. Li, W., Dworkin, L., Seshia, S.A.: Mining assumptions for synthesis. In: MEM-OCODE. pp. 43–50. IEEE (2011)
25. Lowe, G.: Breaking and fixing the needham-schroeder public-key protocol using FDR. Softw. Concepts Tools **17**(3), 93–102 (1996)
26. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems, Joint International Conferences on Formal Modelling and Analysis of Timed Systems, FORMATS 2004 and Formal Techniques in Real-Time and Fault-Tolerant Systems, FTRTFT 2004, Grenoble, France, September 22-24, 2004, Proceedings. Lecture Notes in Computer Science, vol. 3253, pp. 152–166. Springer (2004). https://doi.org/10.1007/978-3-540-30206-3_12
27. Neider, D., Gavran, I.: Learning linear temporal properties. In: FMCAD. pp. 1–10. IEEE (2018)
28. Pakonen, A., Pang, C., Buzhinsky, I., Vyatkin, V.: User-friendly formal specification languages - conclusions drawn from industrial experience on model checking. In: ETFA. pp. 1–8. IEEE (2016)
29. Pnueli, A.: The temporal logic of programs. In: FOCS. pp. 46–57. IEEE Computer Society (1977)
30. Raha, R., Roy, R., Fijalkow, N., Neider, D.: Scalable anytime algorithms for learning formulas in linear temporal logic. CoRR **abs/2110.06726** (2021)
31. Riener, H.: Exact synthesis of LTL properties from traces. In: FDL. pp. 1–6. IEEE (2019)
32. Roy, R., Fisman, D., Neider, D.: Learning interpretable models in the property specification language. In: IJCAI. pp. 2213–2219. ijcai.org (2020)
33. Schlör, R., Josko, B., Werth, D.: Using a visual formalism for design verification in industrial environments. In: Services and Visualization: Towards User-Friendly Design. Lecture Notes in Computer Science, vol. 1385, pp. 208–221. Springer (1998)
34. Shah, A., Kamath, P., Shah, J.A., Li, S.: Bayesian inference of temporal task specifications from demonstrations. In: NeurIPS. pp. 3808–3817 (2018)
35. Solar-Lezama, A.: Program sketching. Int. J. Softw. Tools Technol. Transf. **15**(5-6), 475–495 (2013)
36. Solar-Lezama, A., Rabbah, R.M., Bodík, R., Ebcioglu, K.: Programming by sketching for bit-streaming programs. In: PLDI. pp. 281–294. ACM (2005)
37. Verhulst, E., de Jong, G.G.: Opencomrtos: An ultra-small network centric embedded RTOS designed using formal modeling. In: SDL Forum. Lecture Notes in Computer Science, vol. 4745, pp. 258–271. Springer (2007)
38. Wasylkowski, A., Zeller, A.: Mining temporal specifications from object usage. Autom. Softw. Eng. **18**(3-4), 263–292 (2011)