

Learning Interpretable Classifiers for PDDL Planning

Arnaud Lequen

Linköping University
arnaud.lequen@liu.se

Abstract

We consider the problem of synthesizing interpretable models that recognize the behaviour of an agent compared to other agents, on a whole set of similar planning tasks expressed in PDDL. Our approach consists in learning logical formulas, from a small set of examples that show how an agent solved small planning instances. These formulas are expressed in a version of First-Order Temporal Logic (FTL) tailored to our planning formalism. Such formulas are human-readable, serve as (partial) descriptions of an agent’s policy, and generalize to unseen instances. We show that learning such formulas is computationally intractable, as it is an NP-hard problem. As such, we propose to learn these behaviour classifiers through a topology-guided compilation to MaxSAT, which allows us to generate a wide range of different formulas. Experiments show that interesting and accurate formulas can be learned in reasonable time.

1 Introduction

One of the main strengths of PDDL planning models is that they are succinct and human-readable, but can nonetheless express general, complex problems, whose state search spaces are exponential in the size of the encoding – as can be the solutions. As a consequence, given a set of examples of the behaviour of an agent (called traces), understanding and recognizing this behaviour can be tedious.

In order to summarize the behaviour of a planning agent in a concise, interpretable way, we propose to learn properties that are specific to the solutions proposed by this agent. Such properties, expressed in a temporal logic tailored to fit PDDL planning models, are not only human-readable, but are also general, and can be evaluated against different instances of the same planning problem. This allows them to recognize the behaviour of an agent on instances that are substantially different from the ones used in the set of examples.

More specifically, the problem we tackle is the one where, given a set of positive example traces (the ones of the agent we seek to recognize) and negative examples traces (the ones of other agents), we wish to learn a model that can discriminate as well as possible between positive and negative traces. A wide variety of techniques and models of different natures have been proposed in the literature. Among these, the learning of finite-state automata (DFA) is a well-studied

problem (Angluin 1987; Shvo et al. 2021; Vaandrager et al. 2022), but DFAs can grow quickly (thus becoming harder to interpret) and do not generalize to instances not in the example set. More recently, neural network-based architectures such as LSTMs (Zhou et al. 2015) have shown very promising results, but lack interpretability, and the rationale for their decision is rarely clear.

In the past decade, significant efforts have been made towards learning logical formulas expressed in (a form of) temporal logic. Such works (Raha et al. 2022; Riener 2019; Gaglione et al. 2021; Neider and Gavran 2018; Camacho and McIlraith 2021; de la Rosa and McIlraith 2011) often leverage symbolic methods to learn Linear Temporal Logic (LTL) formulas (Pnueli 1977) that fit the example traces, and thus share some similarities with our work. Some other authors propose other techniques, such as Latent Dirichlet Allocation (Kim et al. 2019), which stems from the field of natural language processing.

However, in all of these cases, the knowledge extracted from the sets of examples has the major drawback of not generalizing well to unknown instances. This is due to the choice of the language used to express these properties. For instance, since LTL formulas are built over a set of propositional variables, they do not generalize to models that do not share the same variables.

To address this issue, we propose to learn properties in a version of First-Order Temporal Logic (FTL). When tailored to the PDDL planning formalism, FTL can express a wide range of properties that generalize from one planning instance to the other, given that they model similar problems. This was shown in (Bacchus and Kabanza 2000), who proposed to express *search control knowledge* in a language similar to ours, albeit with the aim of guiding the search of a planner designed to use such knowledge. In (de la Rosa and McIlraith 2011), the authors proposed to synthesize such control knowledge automatically, and thus address the problem of learning properties expressed in a fragment of FTL.

In this paper, we show that it is possible to learn richer and more expressive properties, using the whole range of FTL operators and modalities. The properties we wish to learn should describe the behaviour of a given planning agent, without being true for the behaviour of other agents. We show that learning such formulas is computationally intractable, as the associated decision problem is NP-hard.

This is why the core of our approach consists in encoding the learning problem into a MaxSAT instance, which has the added benefit of showing resilience to any potential noise in the set of training examples. To make the search more efficient, we fix the general topology of the target formula before the encoding. In addition to alleviating the load on the MaxSAT solver and rendering the algorithm more parallelizable, this also increases the diversity in the formulas learned by our algorithm, thus providing varied descriptions of the behaviour of the agent of interest.

2 Background

2.1 Planning with PDDL

This section introduces the model that we use to describe planning tasks. Our definition of a PDDL planning task differs from (Helmert 2009), as we require the organization of the objects of our instances into types. The model we use resembles the one defined in (Horčík and Fišer 2021)

Definition 1 (Type tree). *A type tree \mathcal{T} is a non-empty tree where each node is labeled by a symbol, called a type. For any type $\tau \in \mathcal{T}$, we call strict subtype any descendant τ' of τ . τ' is a subtype of τ (denoted $\tau' \preceq \tau$) when τ' is a strict subtype of τ or when $\tau' = \tau$.*

Definition 2 (Object class). *Let \mathcal{O} be a set of elements called objects. We call object class any subset of \mathcal{O} . A class c_i is said to be a subclass of type c_j if $c_i \subseteq c_j$.*

Definition 3 (Type hierarchy). *A type hierarchy \mathcal{H} over type tree \mathcal{T} is a set of object classes such that $\mathcal{O} \in \mathcal{H}$, and such that each object class of \mathcal{H} is mapped to a unique type of \mathcal{T} . This mapping $\tau : \mathcal{H} \rightarrow \mathcal{T}$ is such that for any pair c_i, c_j of object classes:*

- c_i is a subclass of c_j iff $\tau(c_i)$ is a subtype of $\tau(c_j)$;
- $c_i \cap c_j = \emptyset$ iff $\tau(c_i)$ is not a subtype of $\tau(c_j)$ (and conversely).

We say that object $o \in \mathcal{O}$ is of type $\tau(o) := \tau(c)$ where c is the smallest (for inclusion \subseteq) class of \mathcal{H} to which o belongs.

Definition 4 (Predicate, atoms and fluents). *A predicate p is a symbol, with which is associated:*

- An arity $ar(p) \in \mathbb{N}$
- A type for each of its arguments. For $i \in \{1, \dots, ar(p)\}$, the type of its argument at position i is denoted $\tau_p(i) \in \mathcal{T}$

An atom is a predicate for which each argument is associated with a symbol, which can be a variable symbol, or an object of \mathcal{O} . When the i -th argument of the atom is an object $o \in \mathcal{O}$ (associated to type hierarchy \mathcal{H}), then we require that $\tau(o) = \tau_p(i)$. The atom consisting of predicate p and symbols $x_1, \dots, x_{ar(p)}$ is denoted $p(x_1, \dots, x_{ar(p)})$.

A fluent is an atom where each argument corresponds to an object of \mathcal{O} . A state is a set of fluents.

Definition 5 (Action schema and operators). *An action schema is a tuple $a = \langle pre(a), add(a), del(a) \rangle$, such that $pre(a)$, $add(a)$ and $del(a)$ are sets of atoms instantiated with variables only.*

An operator o is akin to an action schema, except that the sets $pre(o)$, $add(o)$ and $del(o)$ are sets of fluents.

Definition 6 (PDDL planning problem). *A PDDL planning problem is a pair $\Pi = \langle \mathcal{D}, \mathcal{I} \rangle$ where $\mathcal{D} = \langle \mathcal{P}, \mathcal{A}, \mathcal{T} \rangle$ is the domain and $\mathcal{I} = \langle \mathcal{O}, \mathcal{H}, I, G \rangle$ is the instance.*

The domain \mathcal{D} consists of a set \mathcal{P} of predicates, a set of actions schemas \mathcal{A} , and a type hierarchy \mathcal{T} .

The instance \mathcal{I} consists of a set of objects \mathcal{O} and an associated type hierarchy \mathcal{H} , as well as two states, I and G , which are respectively the initial state and the goal conditions.

An operator o is applicable in a state s if $pre(o) \subseteq s$. The state that results from the application of o in s is $s[o] = (s \setminus del(o)) \cup add(o)$.

A sequence of operators o_1, \dots, o_n is called a *plan* for Π if there exists a sequence of states s_0, \dots, s_n where $s_0 = I$, and which is such that, for all $i \in \{1, \dots, n\}$, $s_i = s_{i-1}[o_i]$ and o_i is applicable in s_{i-1} . Such a sequence of states (which is unique for each plan) is called a *trace*. A plan is called a *solution-plan* if, in addition to this, $G \subseteq s_n$. We say that a fluent $p(o_1, \dots, o_{ar(p)})$ is true in state s iff $p(o_1, \dots, o_{ar(p)}) \in s$.

2.2 First-Order Temporal Logic (FTL)

Syntax Let \mathcal{X} be a set of variable symbols, \mathcal{P} a set of predicates, and \mathcal{T} a type tree. We define our language \mathcal{L}_{FTL} s. t.

$$\psi := \exists x \in \tau. \psi \mid \forall x \in \tau. \psi \mid \varphi$$

where $\varphi \in \mathcal{L}_{TL}$, and \mathcal{L}_{TL} is such that:

$$\varphi := \top \mid p(x_1, \dots, x_{ar(p)}) \mid \neg \varphi \mid \bigcirc \varphi \mid \diamond \varphi \mid \square \varphi \mid \overline{\bigcirc} \varphi \mid \overline{\diamond} \varphi \mid \overline{\square} \varphi \mid \varphi \mathbf{U} \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \Rightarrow \varphi$$

where $x_1, \dots, x_{ar(p)}$ are variables of \mathcal{X} , p a predicate of \mathcal{P} , and τ a type. In the following, we will denote $\Lambda = \{\wedge, \vee, \Rightarrow, \mathbf{U}, \bigcirc, \diamond, \square\}$ the set of all logical operators. For each operator $\lambda \in \Lambda$, we also note $ar(\lambda) \in \{1, 2\}$ the arity of the operator.

This formulation is akin to Linear Temporal Logic on finite traces (LTL_f) (Pnueli 1977), where propositional variables are replaced with first-order predicates and variables. We only work with formulas in prenex normal form.

Environments The formulas of \mathcal{L}_{TL} are built on atoms whose arguments are variables of \mathcal{X} , while traces contain fluents. We bridge that gap with the notion of *environments*, which are akin to interpretations in first-order logic.

Let us denote $\mathcal{X} = (x_1, \dots, x_q)$. In addition, let \mathcal{I} be an instance, with objects $\mathcal{O} = \{o_1, \dots, o_{|\mathcal{O}|}\}$. We call a *partial environment* any assignment of some of the variables x_1, \dots, x_q to objects of \mathcal{O} . Let us denote $var(e)$ the variables that are assigned an object within the partial environment e . When $var(e) = \mathcal{X}$, we simply say that e is an environment. We denote any (partial) environment $e = \{x_1 := o_{i_1}, \dots, x_q := o_{i_q}\}$, where $i_1, \dots, i_q \in \llbracket 1, |\mathcal{O}| \rrbracket$.

The object to which variable x is associated to in e is denoted $x[e]$. We also denote $p(x, y)[e]$ the grounding of an atom $p(x, y)$ by an environment e such that $x, y \in var(e)$. If $e = \{x := o_1, y := o_2, \dots\}$, then $p(x, y)[e] = p(x[e], y[e]) = p(o_1, o_2)$. By extension, the formula obtained when grounding each atom of φ with e is written $\varphi[e]$.

Semantics Given an environment e , any quantifier-free formula φ of \mathcal{L}_{TL} can be evaluated against a trace $t = \langle s_0, \dots, s_n \rangle$, at any step. When $i \in \llbracket 0, n \rrbracket$, we write $t, e, i \models \varphi$ to denote that formula φ is true at state s_i of trace t with environment e . Temporal modalities, such as \bigcirc , \diamond , \square , etc., are used to reason over the states that follow or precede the current state s_i .

$\bigcirc\varphi$ means that property φ is true in the next state, while $\diamond\varphi$ means that φ is eventually true, in one of the successors of the current state. $\square\varphi$ means that φ is true from this state on, until the end of the trace, and $\varphi_1 \mathbf{U} \varphi_2$ means that φ_2 is true in some successor state, and until then, φ_1 is true. Operators $\overline{\bigcirc}$, $\overline{\diamond}$ and $\overline{\square}$ are the *past* counterparts of the previous connectors: $\overline{\bigcirc}\varphi$ means that φ is true in the previous state, $\overline{\diamond}\varphi$ that φ is true in some previous state, and $\overline{\square}\varphi$ that φ is true in every previous state.

To illustrate the language, we introduce the Childsnack problem, which originates from the International Planning Competition (IPC). It consists in making sandwiches and serving them to a group of children, some of whom are allergic to gluten. Sandwiches can only be prepared in the kitchen, and then have to be put on trays, which is the only way they can be brought to the children for service. Among the following FTL formulas, the first indicates that ‘‘All children will eventually be served’’ (and will be satisfied by any solution-plan). The second formula indicates that every sandwich x will eventually be put on some tray, at a moment $t + 1$. For every moment that precedes moment t , x will not be prepared yet (which indicates that the sandwich is actually put on the tray right after being prepared).

$$\forall x \in \text{Child}. \diamond \text{served}(x) \quad (1)$$

$$\forall x \in \text{Sandwich}. \exists y \in \text{Tray}. \text{notprep}(x) \mathbf{U} \bigcirc \text{on}(x, y) \quad (2)$$

Temporal modalities can be expressed in terms of one another. For any quantifier-free formula φ , we have $\diamond\varphi \equiv \top \mathbf{U} \varphi$, $\square\varphi \equiv \neg \diamond \neg \varphi$ and $\overline{\square}\varphi \equiv \neg \overline{\diamond} \neg \varphi$. This leads us to an inductive definition of the semantics of our language, for quantifier-free formulas of \mathcal{L}_{TL} :

$$\begin{aligned} t, e, i \models p(x, \dots, x) & \text{ iff } p(x, \dots, x)[e] \in s_i \\ t, e, i \models \neg\varphi & \text{ iff } t, e, i \not\models \varphi \\ t, e, i \models \varphi_1 \wedge \varphi_2 & \text{ iff } t, e, i \models \varphi_1 \text{ and } t, e, i \models \varphi_2 \\ t, e, i \models \bigcirc\varphi & \text{ iff } i < n \text{ and } t, e, (i+1) \models \varphi \\ t, e, i \models \overline{\bigcirc}\varphi & \text{ iff } i > 0 \text{ and } t, e, (i-1) \models \varphi \\ t, e, i \models \overline{\diamond}\varphi & \text{ iff } \exists j \in \llbracket 0, i \rrbracket \text{ s.t. } t, e, j \models \varphi \\ t, e, i \models \varphi_1 \mathbf{U} \varphi_2 & \text{ iff } \exists j \in \llbracket i, n \rrbracket \text{ s.t. } t, e, j \models \varphi_2 \\ & \text{ and } \forall k \in \llbracket i, j-1 \rrbracket, t, e, k \models \varphi_1 \end{aligned}$$

We write $t, e \models \varphi$ as a shorthand for $t, e, 0 \models \varphi$, which means that trace t satisfies the formula φ , since it is true in the initial state of t . A formula $\psi \in \mathcal{L}_{\text{FTL}}$ is evaluated against instantiated traces:

Definition 7 (Instantiated trace). *An instantiated trace is a pair $\langle t, \mathcal{I} \rangle$ such that t is a trace where fluents are built on the objects \mathcal{O} of the planning instance \mathcal{I} .*

For any partial environment e , $x \in \mathcal{X}$, $o \in \mathcal{O}$, we denote $e[x := o]$ the environment identical to e , but where variable x is associated o . The semantics of \mathcal{L}_{FTL} is defined as follows:

$$\begin{aligned} \langle t, \mathcal{I} \rangle, e \models \forall x \in \tau. \psi & \text{ iff for all } o \in \mathcal{O} \text{ s.t. } \tau(o) = \tau, \\ & \langle t, \mathcal{I} \rangle, e[x := o] \models \psi \\ \langle t, \mathcal{I} \rangle, e \models \exists x \in \tau. \psi & \text{ iff there exists } o \in \mathcal{O} \text{ s.t. } \tau(o) = \tau, \\ & \langle t, \mathcal{I} \rangle, e[x := o] \models \psi \\ \langle t, \mathcal{I} \rangle, e \models \varphi & \text{ iff } t, e \models \varphi \end{aligned}$$

where x is a variable, and φ is a formula of \mathcal{L}_{TL} (thus quantifier-free). We will often denote $\langle t, \mathcal{I} \rangle \models \psi$ as a shorthand for $\langle t, \mathcal{I} \rangle, \emptyset \models \psi$, where \emptyset is the empty environment.

Note that it is well known that the past modalities do not change the expressivity of LTL. As a consequence, our language could have expressed the same properties without modalities $\overline{\bigcirc}$, $\overline{\diamond}$ or $\overline{\square}$. However, we include these modalities in our language as they may make some properties more succinct to express (Markey 2003).

2.3 The MaxSAT problem

Let Var be a set of propositional variables. The boolean satisfiability problem (SAT) is concerned with finding a valuation that satisfies a propositional formula ϕ . Propositional formulas are defined as follows, where $x \in \text{Var}$ is a propositional variable:

$$\phi := \top \mid x \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi$$

The maximum boolean satisfiability problem (MaxSAT) is a variant of SAT, in which a valuation of the variables Var of a set of formulas $\{\phi_1, \dots, \phi_n\}$ is sought. Each formula ϕ_i is assigned a *weight* $w(\phi_i) \in \mathbb{R} \cup \{\infty\}$. The MaxSAT problem consists in finding a valuation v of Var such that the sum of the weights of the formulas that are not satisfied by v is minimal.

3 The \mathcal{L}_{FTL} learning problem

3.1 Problem definition

Score function Our problem takes in input a *score function*, denoted $\sigma : T \rightarrow \mathbb{R}$, where T is the set of traces. This function allows us to express preferences on which traces are the most important to capture in the output formula, and which traces are the most important to avoid. In the rest of this article, we will say that an instantiated trace $\langle t, \mathcal{I} \rangle$ is *positive* iff $\sigma(\langle t, \mathcal{I} \rangle) \geq 0$. Otherwise, the instantiated trace is said to be *negative*.

In the following, we use $[\langle t, \mathcal{I} \rangle \models \psi]$ as a shorthand for the function equal to 1 if $\langle t, \mathcal{I} \rangle \models \psi$ and equal to 0 otherwise. The score function generalizes to formulas $\psi \in \mathcal{L}_{\text{FTL}}$ as follows:

$$\sigma^T(\psi) = \sum_{\langle t, \mathcal{I} \rangle \in T} \sigma(\langle t, \mathcal{I} \rangle) [\langle t, \mathcal{I} \rangle \models \psi]$$

We define below the main problem that we tackle in this article:

Problem 1. \mathcal{L}_{FTL} learning

Input: \mathcal{D} a domain
 T a set of instantiated traces
 $r \in \mathbb{N}$ the maximum number of logical operators in the output formula
 $q \in \mathbb{N}$ the maximum number of quantifiers
 $\sigma : T \rightarrow \mathbb{R}$ a function called the score function

Output: A formula $\psi \in \mathcal{L}_{FTL}$ such that ψ has at most r logical operators, and q quantifiers, and

$$\sum_{\langle t, \mathcal{I} \rangle \in T} \sigma(\langle t, \mathcal{I} \rangle) [\langle t, \mathcal{I} \rangle \models \psi]$$

is maximal

3.2 Complexity

The decision problem associated to the \mathcal{L}_{FTL} learning problem is the problem for which the output is *Yes* iff there exists a formula $\psi \in \mathcal{L}_{FTL}$ satisfying the requirements above, and with score $\sigma^T(\psi) \geq \ell$, where ℓ is given in input. The proof of intractability consists in a reduction from the NP-hard Set Cover Problem (Karp 1972). The proof is sketched below.

Proposition 1. *The decision problem associated to the \mathcal{L}_{FTL} learning problem is NP-hard*

Problem 2. Set Cover

Input: A set $U = \{1, \dots, n\}$
A set S of subsets: $S = \{S_1, \dots, S_m\} \subseteq 2^U$
 $k \in \mathbb{N}$

Output: *Yes* iff there exists a subset $T \subseteq S$ such that $|T| \leq k$ and $\cup_{s \in T} s = U$
No otherwise

Proof of Proposition 1 (Sketch). Let us consider an instance of Set Cover. We build an instance of \mathcal{L}_{FTL} learning that is positive (i.e. outputs *Yes*) iff the Set Cover instance is positive.

The proof consists in showing that a set of positive traces can be described by a formula \mathcal{L}_{FTL} satisfying the constraints in input iff there exists a set cover of size at most k . Each of the positive traces is associated to an element j of U , and contains a single state (and thus, temporal modalities have no effect). This single state carries the information as to which subsets contain j . The information consists of fluents of the form $\text{in}(S_i)$. For instance, if only sets $S_1, S_3 \in S$ contain element j , then the j -th trace is $\langle \{\text{in}(S_1), \text{in}(S_3)\} \rangle$.

Each subset S_i of S , in the Set Cover instance given in input, is associated a unique type Set_i in the PDDL domain we build. As our \mathcal{L}_{FTL} formulas quantify on these types, and since we restrict the number of quantifiers to $q = k$, any output formula of the \mathcal{L}_{FTL} learning problem can only reason on at most k subsets among those in S . If k quantifiers are enough to distinguish between the positive traces (associated to elements $j \in U$) and a mock trace, then k subsets are enough to cover all elements of U . Otherwise, no set cover of size at most k exists.

More formally, let us define the input of our \mathcal{L}_{FTL} learning instance as follows: the domain is $\mathcal{D} = \langle \mathcal{P}, \mathcal{A}, \mathcal{T} \rangle$,

where types are $\mathcal{T} = \{\text{Set}_1, \dots, \text{Set}_m, \text{Set}\}$, $\mathcal{P} = \{\text{in}\}$ (with $ar(\text{in}) = 1$ and $\tau_{\text{in}}(1) = \text{Set}$), and $\mathcal{A} = \emptyset$ (as all traces have length 1, no action is needed).

Each instantiated trace is associated its own instance, which is in turn associated to a unique element $j \in U$: for $j \leq n$, $\mathcal{I}_j = \langle \mathcal{O}, \mathcal{H}, I_j, G_j \rangle$. The shared set of objects contains the sets of the Set Cover input instance, and is such that $\mathcal{O} = \{S_1, \dots, S_m, d\}$ (where d is a mock object that we use to discriminate traces later). The type hierarchy $\mathcal{H} = \{\{S_i\} \mid S_i \in S\} \cup \mathcal{O}$ associates each set to its own unique type in our PDDL domain: types are such that $\tau(S_i) = \text{Set}_i$ for all $i \leq m$, and $\tau(\mathcal{O}) = \text{Set}$. The information regarding which sets j belongs to is encoded in the initial state and goal of \mathcal{I}_j : for $j \leq n$, $I_j = G_j = \{\text{in}(S_i) \mid j \in S_i\}$.

We also define a mock, negative instance $\mathcal{I}_d = \langle \mathcal{O}, \mathcal{H}, I_d, G_d \rangle$, with $I_d = \{\text{in}(d)\}$, and $G_d = I_d$. We have $T = \{\langle t_j, \mathcal{I}_j \rangle \mid j \leq n\} \cup \{\langle t_d, \mathcal{I}_d \rangle\}$, where $t_j = \langle I_j \rangle$ and $t_d = \langle I_d \rangle$. We set $r = m - 1$, $q = k$, and the score $\sigma(\langle t_j, \mathcal{I}_j \rangle)$ of every trace to 1, except $\langle t_d, \mathcal{I}_d \rangle$ for which $\sigma(\langle t_d, \mathcal{I}_d \rangle) = -1$. The decision problem consists in finding whether there exists a formula $\psi \in \mathcal{L}_{FTL}$ that is satisfied by every trace, except $\langle t_d, \mathcal{I}_d \rangle$. Thus, the score threshold we set is $\ell = m$. Let us now show that this instance is positive iff our Set Cover instance is positive.

Suppose that there exists a set cover $T = \{S_{i_1}, \dots, S_{i_k}\}$. Then there exists a formula $\psi \in \mathcal{L}_{FTL}$ with score exactly m :

$$\psi := \exists x_1 \in \text{Set}_{i_1}, \dots, \exists x_k \in \text{Set}_{i_k}. \bigvee_{j \leq k} \text{in}(x_j)$$

We have that $\langle t_d, \mathcal{I}_d \rangle \not\models \psi$, but all other traces satisfy ψ , since T is a set cover. So the formula has score exactly m , and is a positive instance of \mathcal{L}_{FTL} learning.

Now suppose that there exists no set cover of size at most k for the input instance. By contradiction, suppose that there exists a formula ψ with score m . We denote T' the set of types on which ψ quantifies upon. We have that $|T'| \leq k$. Since there exists no set cover of size at most k , then there exists $w \leq n$ such that $w \notin \cup_{s \in T'} s$. The rest of the proof consists in showing that if $\langle t_w, \mathcal{I}_w \rangle \models \psi$, then $\langle t_d, \mathcal{I}_d \rangle \models \psi$, which contradicts the fact that ψ has score m . This can be proven by structural induction on ψ .

Since the \mathcal{L}_{FTL} learning instance we built is positive iff the set cover instance is positive, the \mathcal{L}_{FTL} learning problem is NP-hard. \square

In (Fijalkow and Lagarde 2021; Mascle, Fijalkow, and Lagarde 2023), the authors tackle the \mathcal{L}_{LTL} learning problem, where \mathcal{L}_{LTL} is the fragment of \mathcal{L}_{TL} without past modalities. Along with other past modalities-free fragments, they showed the corresponding decision problem to be NP-complete.

Given an environment e , a trace t , and a formula $\varphi \in \mathcal{L}_{TL}$, checking that $t, e \models \varphi$ can be done in space polynomial in $|t|$, $|e|$ and $|\varphi|$ (ex. (Fionda and Greco 2016)). The model-checking of $\psi \in \mathcal{L}_{FTL}$ against some $\langle t, \mathcal{I} \rangle$ can be done by enumerating all relevant environments $e \in \mathcal{O}^q$, and checking that $t, e \models \varphi$, where φ is the quantifier-free part of ψ . As a consequence, the \mathcal{L}_{FTL} learning problem is in PSPACE.




Figure 1: A TL chain example, which has been assigned symbols to its nodes. It represents the formula $(q(v, u) \wedge r(z, y)) \cup p(t, x)$

Even though this shows membership, the potential PSPACE-hardness of our problem is still an open problem.

4 Planning problem preprocessing

We present in this section the transformations we bring to the PDDL planning problem before it is passed to our algorithm for learning \mathcal{L}_{FTL} formulas.

Predicate splitting Each predicate is split into several predicates of size 2, in order to curb the number of fluents while conserving the links between pairs of objects. This allows us to synthesize formulas containing predicates of high arity, while keeping the number of quantifiers of the formula low.

Concretely, a predicate of the form $p(x, y, z)$ will be split into newly-created predicates $p_{12}(x, y)$, $p_{13}(x, z)$, and $p_{23}(y, z)$. Predicate splitting leads to significantly fewer fluents than if the task was to be grounded as is: for a predicate of arity $n \geq 2$, to be grounded with instance \mathcal{I} , there are $O(n^2|\mathcal{O}|^2)$ associated fluents, while there would be $O(|\mathcal{O}|^n)$ if the predicate was not split.

Even though the planning model thus obtained is less rich than the original one, we argue that predicate splitting allows us to learn formulas that would be otherwise out of computational reach.

Goal predicates In order to allow the learnt formulas to reason on the goal state, we introduce *goal predicates*. For every predicate $p \in \mathcal{P}$, we introduce the predicate p^G . Then, for each instance \mathcal{I} , we introduce the *latent state* $s_{\mathcal{I}}$, which is intuitively a set of fluents that are true in every state of every trace associated to \mathcal{I} .

For every fluent $p(o_1, \dots, o_{ar(p)})$ of the goal state G of \mathcal{I} , we add the fluent $p^G(o_1, \dots, o_{ar(p)})$ to $s_{\mathcal{I}}$.

5 Topology-based guiding

TL chains An interesting representation for formulas φ of \mathcal{L}_{TL} is a representation as *TL chains*. They are the adaptation to our language of the notion of chain (Knuth 2020; Riener 2019), which is useful for representing formulas of modal or propositional logic.

A *TL chain* is a Directed Acyclic Graph (DAG) which has three types of nodes: logical connector nodes (represented as

\circ in the example of Figure 1), predicate nodes (represented as \diamond) and variable nodes (represented as \square). In order to represent a correct \mathcal{L}_{TL} formula, logical connector nodes can only be children of logical connector nodes, predicate nodes children of logical connector nodes, and variable nodes children of predicate nodes. We also impose that every leaf is a variable node. In addition, to stay consistent with the choices we made in Section 4, we only work with TL chains that are binary trees, whose inner nodes have exactly two children.

By assigning a symbol of the correct type (i.e., a logical connector, a predicate symbol or a variable) to each node, we end up with a representation of a \mathcal{L}_{TL} formula, as illustrated in Figure 1.

For each connector node i of the TL chain, we will denote $\text{succ}_L(i)$ (resp. $\text{succ}_R(i)$) the left (resp. right) child of node i . It is guaranteed to exist, but might sometimes be a predicate node. In the case of connectors $\alpha \in \{\neg, \circ, \overline{\circ}, \diamond, \overline{\diamond}, \square, \overline{\square}\}$ that have arity 1, we will use the convention that the value of the right successor is ignored (and will not appear in the \mathcal{L}_{FTL} formula that ensues), and the left successor will be the root of the formula under the operator α .

In order to alleviate the pressure on the MaxSAT solver, we impose the topology of the output quantifier-free formula before encoding the problem into a propositional formula. This idea was first introduced in (Riener 2019), in the case of a search for LTL formulas. We also fix the quantifiers of the formula before the encoding, as well as the associated types. All that is left to the MaxSAT solver is to “fill in the blanks” in the TL chains that it is given, so that the associated \mathcal{L}_{FTL} formula fits the input as well as possible.

Quantifiers In the rest of the article, for practical reasons, we restrict ourselves to learning formulas of the form $\forall x_1 \dots \forall x_b \exists x_{b+1} \dots \exists x_q \varphi$, where φ is a formula of \mathcal{L}_{TL} , for which every argument of every predicate is a variable x_i . This allows us to curb the size of the MaxSAT encoding.

6 Reduction to MaxSAT

6.1 Learning algorithm

Algorithm 1 summarizes the procedure that we use to learn \mathcal{L}_{FTL} formulas out of our input. The subroutines work as follows: `gen_TLchains(r)` enumerates every TL chain having exactly r connectors. `gen_quantifiers(q)` enumerates sequences of quantifier symbols of size q , such that all universal quantifiers \forall appear before existential quantifiers \exists . `gen_types(\mathcal{D} , q)` enumerates every q -combination of types in the type tree \mathcal{T} of \mathcal{D} . Finally, the main subroutine, `find_formula(\mathcal{D} , \mathcal{T} , ρ , $\{Q_i\}$, $\{\tau_i\}$, σ)`, encodes the problem of finding an \mathcal{L}_{FTL} formula fitting the instantiated traces of \mathcal{T} , with the constraints imposed by the TL chain ρ , the quantifiers $\{Q_i\}$, and the types $\{\tau_i\}$. `find_formula` then returns (one of) the best formula(s) it finds, or the token FAIL if none is found.

6.2 Preliminaries to the encoding

Variables Our MaxSAT encoding is built on the set of variables that follows. When possible, we use the following conventions, as closely as possible: nodes of the FL-chain

Algorithm 1 \mathcal{L}_{FTL} learning

Input: Domain \mathcal{D} , traces \mathcal{T} , parameters r, q , and function σ **Output:** A set of \mathcal{L}_{FTL} formulas

```
1: found_formulas := []
2: for  $\rho \in \text{gen\_TLchains}(r)$  do
3:   for  $Q_1, \dots, Q_q \in \text{gen\_quantifiers}(q)$  do
4:     for  $\tau_1, \dots, \tau_q \in \text{gen\_types}(\mathcal{D}, q)$  do
5:        $\psi \leftarrow \text{find\_formula}(\mathcal{D}, \mathcal{T}, \rho, \{Q_i\}, \{\tau_i\}, \sigma)$ 
6:       if  $\psi \neq \text{FAIL}$  then
7:         found_formulas.add( $\psi$ )
8:       end if
9:     end for
10:  end for
11: end for
12: return found_formulas
```

are denoted by i when they are logical connectors (represented by \bigcirc in Figure 1), by ℓ when they are predicate nodes (represented by \diamond), and by v when they are first-order variable nodes (represented by \square). A trace is denoted by t , and a position in this trace is denoted by k (i.e., the k -th state). Moreover, j is an index for a variable of the quantifiers, and p is a predicate.

This leads us to the following variables, as will be used in the MaxSAT encoding. Greek letters denote decision variables while latin characters are for “technical” variables.

- $y_i^{t,k}[e]$: In position k of trace t , with environment e , the formula rooted at node i is true.
- $\delta_{j,v}^\ell$: The v -th variable of predicate node ℓ is the variable of quantifier j .
- θ_ℓ^p : The predicate of node ℓ is p .
- λ_i^q : The logical connector at node i is q .
- s_t : Trace t is currently satisfied by the first order formula

“Exactly one” constraints In the encoding of a problem into SAT, some situations require that *exactly one* variable, out of a set of variables, is true. Efficient encodings have been widely studied: see for instance (Hölldobler and Nguyen 2013; Nguyen and Mai 2015). In the following, we will denote $\text{ExactlyOne}_{s \in S}(v_s)$ the set of propositional constraints enforcing that exactly one of the variables of $\{v_s \mid s \in S\}$ is true.

6.3 Core constraints

Some of the constraints below are adapted from (Gaglione et al. 2021; Rienr 2019; Neider and Gavran 2018), which are concerned with LTL. Our main contribution is the adaptation of the encoding to our language \mathcal{L}_{FTL} , which differs from LTL by its tighter links with PDDL planning models through first-order components.

In the following, we suppose that an empty TL chain ρ has been computed, and that the associated quantifiers and types have been decided. We will denote n its number of connector nodes, and m its number of predicate nodes. As a consequence, there are $2m$ variable nodes. As previously, the number of quantifiers is denoted q . The first $b \leq q$ quantifiers are universal, while the others are existential.

We also suppose that the types on which the quantifiers range, denoted τ_1, \dots, τ_q , are already chosen. As a consequence, in this section, the set of relevant environments for instance \mathcal{I}_t associated to trace t , denoted $E_{\mathcal{I}_t}$, only consists of environments of the form $\{x_u := o_u\}_{1 \leq u \leq q}$ where, $\tau(o_u) = \tau_u$, for $u \in \llbracket 1, q \rrbracket$.

Syntactic constraints This section describes the constraints that ensure that the formula is syntactically well-formed.

The following constraints respectively ensure that every logical connector node has exactly one logical connector assigned, that every predicate node has exactly one predicate, and that each argument of each predicate is bound to a variable on which the formula quantifies.

$$\bigwedge_{i \leq n} \text{ExactlyOne}_{c \in \Lambda}(\lambda_i^c) \wedge \bigwedge_{\ell \leq m} \text{ExactlyOne}_{p \in \mathcal{P}}(\theta_\ell^p)$$
$$\bigwedge_{\ell \leq m} \bigwedge_{s \in \{1,2\}} \text{ExactlyOne}_{j \leq b}(\delta_{j,s}^\ell)$$

Semantic constraints These constraints ensure that the formula found by the solver is consistent with the traces, and is reminiscent of the model-checking algorithm for modal logic.

The following clauses ensure that the formula ψ that is synthesized is consistent with the traces of \mathcal{T} . This is made in accordance with the environments imposed by the quantifier, which are iterated upon. The variable s_t is true iff for every required environment e , $\varphi[e]$ is satisfied by t (where $\varphi[e]$ is the evaluation of formula φ in environment e , and φ is the quantifier-free part of the formula we synthesize). Thus, for every trace $t \in \mathcal{T}$, we add the following:

$$s_t \Leftrightarrow \bigwedge_{\substack{o_1 \in \mathcal{O}_1 \dots o_k \in \mathcal{O}_k \\ o_{k+1} \in \mathcal{O}_{k+1} \dots o_q \in \mathcal{O}_q}} \bigvee y_1^{t,1}[\{x_u := o_u\}_{1 \leq u \leq q}] \quad (3)$$

The following constraints ensure that formulas that consist of a single literal (i.e., a positive or negative fluent) are consistent with the y variables, that give the truth value of a trace at a certain position in the trace, at each node of the TL chain.

Such constraints appear once for every trace $t \in \mathcal{T}$, for every position $k \leq |t|$ of this trace, for every predicate node $\ell \leq m$ and every predicate $p \in \mathcal{P}$, for every pair of quantifiers (positions) $j_1, j_2 \leq q$, and for each relevant environment $e \in E_{\mathcal{I}_t}$.

$$\theta_\ell^p \wedge \delta_{j_1,1}^\ell \wedge \delta_{j_2,2}^\ell \Rightarrow \begin{cases} y_\ell^{t,k}[e] & \text{if } t[k] \models p(x_{j_1}, x_{j_2})[e] \\ \neg y_\ell^{t,k}[e] & \text{otherwise} \end{cases} \quad (4)$$

Constraints (5) to (8) appear once for each connector node $i \leq n$ of the formula, each position $k \leq |t|$ of each trace $t \in \mathcal{T}$, and for each environment $e \in E_{\mathcal{I}_t}$. They ensure that the logical operators are correctly interpreted.

In the case where the logical connector at node i is a negation \neg , or $\Delta \in \{\wedge, \vee, \Rightarrow\}$, or the next operator \bigcirc , we have:

$$\lambda_i^\neg \Rightarrow \left(y_i^{t,k}[e] \Leftrightarrow \neg y_{\text{succ}_{\mathcal{L}}(i)}^{t,k}[e] \right) \quad (5)$$

$$\lambda_i^\Delta \Rightarrow \left(y_i^{t,k}[e] \Leftrightarrow \left(y_{succ_L(i)}^{t,k}[e] \Delta y_{succ_R(i)}^{t,k}[e] \right) \right) \quad (6)$$

$$\lambda_i^\circ \Rightarrow \left(y_i^{t,k}[e] \Leftrightarrow y_{succ_L(i)}^{t,k+1}[e] \right) \quad (7)$$

with the convention that $y_{succ_L(i)}^{t,|t|+1}[e]$ is replaced by \perp during the encoding itself. In the case of the finally operator \diamond :

$$\lambda_i^\diamond \Rightarrow \left(y_i^{t,k}[e] \Leftrightarrow \bigvee_{\substack{k' \\ k \leq k' \leq |t|}} y_{succ_L(i)}^{t,k'}[e] \right) \quad (8)$$

The case of the temporal operators \square , $\overline{\square}$, $\overline{\diamond}$, $\overline{\square}$ and \mathbf{U} can be encoded in a way that is similar to the constraints above.

Well-formed fluents constraints The following constraints ensure that, in the output formula ψ , there is a consistency between the types of the variables and the arguments of predicates are assigned to. In other words, when a variable x of type τ is chosen to be the v -th argument of a predicate p that occurs in ψ , we require that $\tau = \tau_p(v)$. This can be done through the following constraints:

$$\bigwedge_{j < q} \bigwedge_{\ell \leq m} \bigwedge_{p \in \mathcal{P}} \bigwedge_{j < q} \bigwedge_{\substack{v \leq 2 \\ \tau_p(v) \neq \tau_j}} -\theta_\ell^p \vee -\delta_{j,v}^\ell \quad (9)$$

Weights for the MaxSAT solver Recall that we wish to find a formula ψ that maximizes the function given in Problem 1. The objective of the MaxSAT solver is to minimize the total weight of the falsified soft clauses. As such, for each instantiated trace $\langle t, \mathcal{I} \rangle$, we add the clause s_t , with weight $\sigma(\langle t, \mathcal{I} \rangle)$. This penalizes formulas that falsify traces with a positive score, while rewarding formulas that falsify traces with a negative score.

Pruning non-discriminatory formulas With a given configuration of TL chain, quantifiers and types, it is not guaranteed that there exists a formula ψ that captures (some of) the positive traces while falsifying (some of) the negative traces. To prevent tautology or unsatisfiable formulas from occurring, we enforce the constraint that at least one positive trace is captured and one negative trace is not.

6.4 Formula quality enhancement

The constraints of this section filter the solutions so that less interesting formulas, or formulas that could be computed by a run of our algorithm with smaller parameters, are barred from being output.

Syntactic redundancies prevention These constraints prevent idempotent and involutive modalities and operators from being chained in the output formula. These include the negation \neg , as well as the temporal operators \diamond (for which $\diamond \diamond \varphi \equiv \diamond \varphi$) and \square (which is, likewise, idempotent).

In addition, we prevent redundancies of the form $p(x, y) \Delta p(x, y)$, where $\Delta \in \{\wedge, \vee, \mathbf{U}, \Rightarrow\}$ is a binary operator. In every case, there exists a smaller (sub-)formula that can be found and that expresses the same property, without the redundant atom. For space reasons, we skip the presentation of the constraints.

Variable visibility We wish to ensure that every variable that we quantify upon in the output formula ψ also appears in an atom of ψ . Otherwise, an equivalent formula could be found by running the algorithm with fewer quantifiers. This is why we force each variable to appear at least once in some atom.

7 Experiments

We implemented Algorithm 1 in Python 3.10, using the MaxSAT solver Z3 (de Moura and Bjørner 2008). Our code and data are available online (Lequen 2024). Experiments were conducted on a machine running Rocky Linux 8.5, powered by an Intel Xeon E5-2667 v3 processor, with a 9-hours cutoff and using at most 8GB of memory per run.

Even though we often managed to quickly find a formula that perfectly captures the set of examples, we let the algorithm run to the end, so that all TL chains and combinations of quantifiers and types are enumerated.

Building the data sets To assess the performances of our algorithm, we considered domains from the International Planning Competition (IPC), 2 of which are described in Section 7.1. For each of these domains, we generated 23 instances that model problems with similar goals. Then, for each domain, we designed three domain-specific planners, that solve the tasks in a distinctive way.

We built our training sets by selecting 3 small planning instances of each domain, and the associated traces for each planner – for a total of 9 traces per domain. We then created the tasks of finding a formula recognizing the behaviour of each planner, out of 1, 2 or 3 of the training instances. The 20 remaining instances (and their associated plans) were used in the test set. The traces of our training set have length 5 to 21, with an average of 11.8 states.

7.1 Examples of learnt formulas

Childsnack We designed three different agents that solve Childsnack instances, as introduced in Section 2.2. Agents NGF and NGL compute solution plans of minimal size, and differ in that agent NGF makes sandwiches with *no gluten first*, and agent NGL makes sandwiches with *no gluten last*. Both agents make all sandwiches, put them on a tray, then serve the children. Agent GS greedily serves children: as soon as a sandwich is made, it is put on a tray and brought to a child. It also prioritizes gluten-free sandwiches.

For each of these behaviours, the total computation time and the total number of formulas found are depicted in Table 3 and Table 2, respectively. As shown on Table 1, we manage to learn formulas that perfectly capture the test set, even when the number of examples is very limited, and when few quantifiers and logical operators are allowed. Formula 2, given in Section 2.2, is an example of a concise formula that describes the policy of agent GS perfectly. Other formulas include the following:

$$\forall x \in \text{Kitchen}. \exists y \in \text{Tray}. \diamond(\text{at}(y, x) \wedge \overline{\diamond}\neg\text{at}(y, x)) \quad (10)$$

Formula (10) expresses that agent GS eventually comes back to the (only) kitchen with some tray y , even though the tray was brought out of the kitchen at some point in the past. The

		GS					
# ins.	q	1		2		3	
		1	2	1	2	1	2
$ \varphi =2$	2	67	100	67	100	67	100
	3	100	100	100	100	100	100
	4	.	100	.	100	.	.

		NGF					
# ins.	q	1		2		3	
		1	2	1	2	1	2
$ \varphi =2$	2	67	67	67	67	67	67
	3	100	100	100	100	100	67
	4

		NGL					
# ins.	q	1		2		3	
		1	2	1	2	1	2
$ \varphi =2$	2	100	100	100	100	100	100
	3	100	100	100	100	100	100
	4

Table 1: Proportion of traces correctly classified by the most accurate formula on our test set, expressed in percentage (%). Each table presents a different agent solving a set of Childsnack problems. # ins. indicates the number of instances in the training set, q the number of quantifiers allowed, and $|\varphi|$ the number of logical operators allowed. Dots correspond to configuration where the run did not terminate.

		GS					
# ins.	q	1		2		3	
		1	2	1	2	1	2
$ \varphi =2$	2	22	110	22	112	22	123
	3	92	624	92	624	92	671
	4	.	746	.	757	.	.

		NGF					
# ins.	q	1		2		3	
		1	2	1	2	1	2
$ \varphi =2$	2	23	126	23	126	23	140
	3	92	670	92	679	94	723
	4	.	797	.	858	.	.

		NGL					
# ins.	q	1		2		3	
		1	2	1	2	1	2
$ \varphi =2$	2	7	76	7	78	8	92
	3	79	604	79	681	82	725
	4	.	784	.	845	.	.

Table 2: Total number of formulas found for Childsnack.

formula perfectly captures our test set, but does not perfectly capture our *training* set. Indeed, the smallest instance of our training set contains as many children as there are trays, and thus, no tray has to be brought back to the kitchen. Our use of a reduction to MaxSAT allows us to be resilient to this kind of edge cases, and the formula above is satisfactory despite not perfectly fitting the training set.

Spanner Instances of the Spanner domain involve an operator that has to go from a shed to a gate to tighten some nuts, passing through a sequence of locations where single-use spanners can be picked up. Once a location is left, it can not be returned to. Thus, collecting enough spanners before reaching the gate is essential.

We developed three different behaviours for this domain. Agent ALL picks every possible spanner on its way to the gate, while agent SME picks exactly as many spanners as are needed to tighten the nuts at the gate. Agent SGL takes a single spanner and rushes to the gate, and can then only tighten one nut. Our algorithm learnt the following formulas,

		GS					
# ins.	q	1		2		3	
		1	2	1	2	1	2
2	2	0:00:08	0:05:03	0:00:08	0:16:43	0:01:14	0:44:48
	3	0:00:35	0:19:53	0:00:35	1:07:27	0:06:46	3:35:18
	4	.	0:25:27	.	1:32:07	.	.

		NGF					
# ins.	q	1		2		3	
		1	2	1	2	1	2
2	2	0:00:08	0:05:08	0:00:08	0:17:07	0:01:09	0:47:51
	3	0:00:37	0:20:13	0:00:37	1:11:13	0:07:05	3:44:11
	4	.	0:26:29	.	1:33:27	.	.

		NGL					
# ins.	q	1		2		3	
		1	2	1	2	1	2
2	2	0:00:06	0:04:55	0:00:06	0:14:59	0:00:48	0:40:28
	3	0:00:31	0:19:00	0:00:31	1:07:48	0:05:42	3:14:13
	4	.	0:25:20	.	1:30:16	.	.

Table 3: Total computation times for Childsnack (hh:mm:ss). Rows correspond to the number of allowed logical connectors. Dots correspond to configuration where the run did not terminate.

that perfectly recognize agent ALL:

$$\forall x \in \text{Spanner}. \exists y \in \text{Operator}. \diamond \square \text{carrying}(y, x) \quad (11)$$

$$\forall x \in \text{Spanner}. \exists y \in \text{Location}. \text{at}(x, y) \wedge \diamond \neg \text{at}(x, y) \quad (12)$$

Formula (11) expresses that every spanner will be picked up by the (only) operator and carried for the rest of the plan, and Formula (12) expresses that every spanner will be moved from its initial position.

When searching formulas with a single variable, we split the predicates so that the maximum arity of a fluent is 1. Our algorithm outputs the following very simple formula, which was learnt in a few seconds, while completely characterizing the behaviour of agent ALL:

$$\forall x \in \text{Spanner}. \diamond \text{carrying}_2(x)$$

8 Conclusion

In this paper, we have presented a method to learn temporal logic formulas that recognize agents based on examples of their behaviours. We showed that such formulas can be learned using an algorithm based on a reduction to MaxSAT, and that very few examples are sometimes enough to perfectly capture the behaviour of an agent on instances that can differ from the ones used in the training set. This justifies the cost of resorting to a first-order language, which generalizes to new instances, but is also very concise and easily readable by a human. The formulas that we learn can serve as higher-order descriptions of the behaviour of a planning agent.

Future works include tailoring our algorithm and our datasets so that they can generate domain-specific control knowledge. More specifically, we wish to work on its integration into various systems that can be guided with temporal logic, be them automated planning systems (Bacchus and Kabanza 2000) or reinforcement learning agents (Toro Icarte et al. 2018). This is in line with previous works on generalized planning, which learn logic-based policies fully capable of solving a set of planning problems, out of a set of example instances and plans (Bonet, Frances, and Geffner 2019; Francès, Bonet, and Geffner 2021).

Acknowledgments

The author would like to thank Martin C. Cooper for his insightful advices and his careful proofreading of this manuscript. The author also wishes to thank the anonymous reviewers for their remarks and suggestions.

References

- Angluin, D. 1987. Learning regular sets from queries and counterexamples. *Information and Computation*, 75(2): 87–106.
- Bacchus, F.; and Kabanza, F. 2000. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116(1): 123–191.
- Bonet, B.; Frances, G.; and Geffner, H. 2019. Learning features and abstract actions for computing generalized plans. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2703–2710.
- Camacho, A.; and McIlraith, S. A. 2021. Learning Interpretable Models Expressed in Linear Temporal Logic. *Proceedings of the International Conference on Automated Planning and Scheduling*, 29(1): 621–630.
- de la Rosa, T.; and McIlraith, S. 2011. Learning domain control knowledge for TLPlan and beyond. In *Proceedings of the ICAPS-11 Workshop on Planning and Learning (PAL)*.
- de Moura, L.; and Bjørner, N. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, 337–340. Springer Berlin Heidelberg.
- Fijalkow, N.; and Lagarde, G. 2021. The complexity of learning linear temporal formulas from examples. In Chandlee, J.; Eyraud, R.; Heinz, J.; Jardine, A.; and van Zaanen, M., eds., *Proceedings of the Fifteenth International Conference on Grammatical Inference*, volume 153 of *Proceedings of Machine Learning Research*, 237–250. PMLR.
- Fionda, V.; and Greco, G. 2016. The Complexity of LTL on Finite Traces: Hard and Easy Fragments. In *AAAI Conference on Artificial Intelligence*.
- Francès, G.; Bonet, B.; and Geffner, H. 2021. Learning General Planning Policies from Small Examples Without Supervision. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, 11801–11808. AAAI Press.
- Gaglione, J.-R.; Neider, D.; Roy, R.; Topcu, U.; and Xu, Z. 2021. Learning linear temporal properties from noisy data: a MaxSAT-based approach. In *Automated Technology for Verification and Analysis: 19th International Symposium, ATVA 2021, Gold Coast, QLD, Australia, October 18–22, 2021, Proceedings 19*, 74–90. Springer.
- Helmert, M. 2009. Concise finite-domain representations for PDDL planning tasks. *Artificial Intelligence*, 173(5-6): 503–535.
- Hölldobler, S.; and Nguyen, V.-H. 2013. An Efficient Encoding of the at-most-one Constraint. *Technical Report. Technische Universität Dresden*.
- Horčík, R.; and Fišer, D. 2021. Endomorphisms of lifted planning problems. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 31, 174–183.
- Karp, R. M. 1972. *Reducibility among Combinatorial Problems*, 85–103. Boston, MA: Springer US. ISBN 978-1-4684-2001-2.
- Kim, J.; Muise, C.; Shah, A. J.; Agarwal, S.; and Shah, J. A. 2019. Bayesian inference of linear temporal logic specifications for contrastive explanations. *International Joint Conference on Artificial Intelligence*.
- Knuth, D. E. 2020. *The Art of Computer Programming*, volume 4A. Addison Wesley.
- Lequen, A. 2024. Implementation of an FTL-based classifier learner for PDDL planning. Zenodo. <https://doi.org/10.5281/zenodo.13354031>.
- Markey, N. 2003. Temporal logic with past is exponentially more succinct. *Bull. EATCS*, 79: 122–128.
- Masclé, C.; Fijalkow, N.; and Lagarde, G. 2023. Learning temporal formulas from examples is hard. *arXiv*.
- Neider, D.; and Gavran, I. 2018. Learning linear temporal properties. In *2018 Formal Methods in Computer Aided Design (FMCAD)*, 1–10. IEEE.
- Nguyen, V.-H.; and Mai, S. T. 2015. A New Method to Encode the At-Most-One Constraint into SAT. In *Proceedings of the 6th International Symposium on Information and Communication Technology, SoICT '15*, 46–53. New York, NY, USA: Association for Computing Machinery. ISBN 9781450338431.
- Pnueli, A. 1977. The temporal logic of programs. *18th Annual Symposium on Foundations of Computer Science (SFCS 1977)*, 46–57.
- Raha, R.; Roy, R.; Fijalkow, N.; and Neider, D. 2022. Scalable Anytime Algorithms for Learning Fragments of Linear Temporal Logic. In Fisman, D.; and Rosu, G., eds., *Tools and Algorithms for the Construction and Analysis of Systems*, 263–280. Cham: Springer International Publishing.
- Riener, H. 2019. Exact Synthesis of LTL Properties from Traces. In *2019 Forum for Specification and Design Languages*, 1–6. IEEE.
- Shvo, M.; Li, A. C.; Icarte, R. T.; and McIlraith, S. A. 2021. Interpretable sequence classification via discrete optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 9647–9656.
- Toro Icarte, R.; Klassen, T. Q.; Valenzano, R.; and McIlraith, S. A. 2018. Teaching Multiple Tasks to an RL Agent using LTL. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, 452–461. International Foundation for Autonomous Agents and Multiagent Systems.
- Vaandrager, F.; Garhewal, B.; Rot, J.; and Wißmann, T. 2022. A new approach for active automata learning based on apartness. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 223–243. Springer.
- Zhou, C.; Sun, C.; Liu, Z.; and Lau, F. 2015. A C-LSTM Neural Network for Text Classification. *arXiv*.