

Learning temporal formulas from examples is hard*

Corto Mascle

LaBRI, University of Bordeaux, France

CORTO.MASCLE@LABRI.FR

Nathanaël Fijalkow

CNRS, LaBRI, Université de Bordeaux, France

NATHANAEL.FIJALKOW@GMAIL.COM

Guillaume Lagarde

LaBRI, University of Bordeaux, France

GUILLAUME.LAGARDE@GMAIL.COM

Abstract

We study the problem of learning linear temporal logic (LTL) formulas from examples, as a first step towards expressing a property separating positive and negative instances in a way that is comprehensible for humans. In this paper we initiate the study of the computational complexity of the problem. Our main results are hardness results: we show that for almost all fragments of LTL, the learning problem is NP-complete. This motivates the search for efficient heuristics, and highlights the complexity of expressing separating properties in concise natural language.

1. Introduction

We are interested in the complexity of learning formulas of Linear Temporal Logic (**LTL**) from examples, in a passive scenario: from a set of positive and negative words, the objective is to construct a formula, as small as possible, which satisfies the positive words and does not satisfy the negative words. Passive learning of languages has a long history paved with negative results. Learning automata is notoriously difficult from a theoretical perspective, as witnessed by the original NP-hardness result of learning a Deterministic Finite Automaton (DFA) from examples (Gold, 1978). This line of hardness results culminates with the inapproximability result of (Pitt & Warmuth, 1993) stating that there is no polynomial time algorithm for learning a DFA from examples even up to a polynomial approximation of their size.

One approach to cope with such hardness results is to change representation, for instance replacing automata by logical formulas; their syntactic structures make them more amenable to principled search algorithms. There is a range of potential logical formalisms to choose from depending on the application domain. Linear Temporal Logic (Pnueli, 1977), which we abbreviate as **LTL**, is a prominent logic for specifying temporal properties over words, it has become a de facto standard in many fields such as model checking, program analysis, and motion planning for robotics. Verification of **LTL** specifications is routinely employed in industrial settings and marks one of the most successful applications of formal methods to real-life problems. A key property making **LTL** a strong candidate as a concept class is

*. This article is a long version of the article presented in the proceedings of the International Conference on Grammatical Inference (ICGI) in 2021 (Fijalkow & Lagarde, 2021). It includes much stronger and more general results than the extended abstract.

that its syntax does not include variables, contributing to the fact that **LTL** formulas are typically easy to interpret and therefore useful as explanations.

Over the past five to ten years learning temporal logics (of which **LTL** is the core) has become an active research area. There are many applications, let us only mention a few: program specification (Lemieux, Park, & Beschastnikh, 2015) and fault detections (Bombara, Vasile, Penedo Alvarez, Yasuoka, & Belta, 2016). We refer to (Camacho & McIlraith, 2019) for a longer discussion on the potential and actual applications of learning temporal logics.

Since learning temporal logics is a computationally hard problem, a number of different approaches have been explored. One of the first and probably most natural is leveraging SAT solvers (Neider & Gavran, 2018), which can then accommodate noisy data (Gaglione, Neider, Roy, Topcu, & Xu, 2022). Another line of work relies on their connections to automata (Camacho & McIlraith, 2019), and a third completely different idea approaches it from the lens of Bayesian inference (Kim, Muise, Shah, Agarwal, & Shah, 2019). Learning specifically tailored fragments of **LTL** yields the best results in practice (Raha, Roy, Fijalkow, & Neider, 2022). There are a number of temporal logics, and the ideas mentioned above have been extended to more expressive logics such as Property Specification Language (PSL) (Roy, Fisman, & Neider, 2020; Baharisangari, Gaglione, Neider, Topcu, & Xu, 2021), Computational Tree Logic (CTL) (Ehlers, Gavran, & Neider, 2020), and Metric Temporal Logic (MTL) (Raha, Roy, Fijalkow, Neider, & Pérez, 2024). Other paradigms have been explored to make learning temporal logics useful, such as sketching (Lutz, Neider, & Roy, 2023).

Despite this growing interest, very little is known about the computational complexity of the underlying problem; indeed the works cited above focused on constructing efficient algorithms for practical applications. The goal of this paper is to initiate the study of the complexity of learning **LTL** formulas from examples.

Our contributions. We present a set of results for several fragments of **LTL**, showing in almost all cases that the learning problem is NP-complete. Section 2 gives definitions.

- Our first NP-hardness result is presented in Section 3, it states that the learning problem for full **LTL** is NP-hard when the alphabet size is part of the input.
- To obtain membership in NP for the learning problem, we show in Section 4 that all fragments of **LTL** have the short formula property. We then study some (degenerate) fragments in Section 5 and show that for these fragments the **LTL** learning problem is in polynomial time.
- We construct in Section 6 a polynomial-time approximation algorithm for **LTL** with only the next operator and conjunctions, and show that assuming $P \neq NP$, the approximation ratio of this algorithm is optimal.
- Our most technical results are presented in Section 7 and Section 8: in the first section we show that almost all fragments with the next operator are hard to approximate, and in the next a similar result for almost all fragments without the next operator.

We conclude in Section 9.

2. Preliminaries

Let us fix a (finite) alphabet Σ . We index words from position 1 (not 0) and the letter at position i in the word w is written $w(i)$, so $w = w(1)\dots w(\ell)$ where ℓ is the length of w , written $|w| = \ell$. We write $w[k\dots] = w(k)\dots w(\ell)$. To avoid unnecessary technical complications we only consider non-empty words, and let Σ^+ denote the set of (non-empty) words.

The syntax of Linear Temporal Logic (**LTL**) includes atomic formulas $c \in \Sigma$ and their negations, as well as \top and \perp , the boolean operators \wedge and \vee , and the temporal operators **X**, **F**, **G**, and **U**. Note that as usually done, we work with **LTL** in negation normal form, meaning that negation is only used on atomic formulas. The semantic of **LTL** over finite words is defined inductively over formulas, through the notation $w \models \varphi$ where $w \in \Sigma^+$ is a non-empty word, and φ is an **LTL** formula. The definition is given below for the atomic formulas and temporal operators **U**, **F**, **G**, and **X**, with boolean operators interpreted as usual.

- We have $w \models \top$ and $w \not\models \perp$.
- $w \models c$ if $w(1) = c$.
- $w \models \mathbf{X}\varphi$ if $|w| > 1$ and $w[2\dots] \models \varphi$. It is called the next operator.
- $w \models \mathbf{F}\varphi$ if $w[i\dots] \models \varphi$ for some $i \in [1, |w|]$. It is called the eventually operator.
- $w \models \mathbf{G}\varphi$ if $w[i\dots] \models \varphi$ for all $i \in [1, |w|]$. It is called the globally operator.
- $w \models \varphi\mathbf{U}\psi$ if there exists $i \in [1, |w|]$ such that for $j \in [1, i - 1]$ we have $w[j\dots] \models \varphi$, and $w[i\dots] \models \psi$. It is called the until operator.

Note that $\mathbf{F}\varphi$ is syntactic sugar for $\top\mathbf{U}\varphi$. We say that w satisfies φ when $w \models \varphi$ is true. It is sometimes useful to write $w, i \models \varphi$ to mean $w[i\dots] \models \varphi$. We consider fragments of **LTL** by specifying which boolean connectives and temporal operators are allowed. For instance **LTL**(**X**, \wedge) is the set of all **LTL** formulas using only atomic formulas, conjunctions, and the next operator. The full logic is **LTL** = **LTL**(**U**, **F**, **G**, **X**, \wedge , \vee). More generally, for a set of operators $Op \subseteq \{\mathbf{U}, \mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge, \vee\}$, we write **LTL**(Op) for the logic using operators from Op . The size of a formula is the size of its syntactic tree. We say that two formulas are equivalent if they have the same semantics, and we write $\varphi \equiv \psi$ to say that φ and ψ are equivalent.

The LTL learning problem. A sample is a pair (P, N) where $P = \{u_1, \dots, u_n\}$ is a set of positive words and $N = \{v_1, \dots, v_m\}$ a set of negative words. Without loss of generality we can assume that $n = m$ (adding duplicate identical words to have an equal number of positive and negative words). The **LTL** learning decision problem is:

INPUT: a sample (P, N) and $k \in \mathbb{N}$,
QUESTION: does there exist an **LTL** formula φ of size at most k such that for all $u \in P$, we have $u \models \varphi$, and for all $v \in N$, we have $v \not\models \varphi$?

In that case we say that φ separates P from N , or simply that φ is a separating formula if the sample is clear from the context. The **LTL** learning problem is analogously defined for any fragment of **LTL**.

Parameters for complexity analysis. The three important parameters for the complexity of the **LTL** learning problem are: n the number of words, ℓ the maximum length of the words, and k the desired size for the formula. As we will see, another important parameter is the size of the alphabet. We will consider two settings: either the alphabet Σ is fixed, or it is part of the input. The size of the alphabet is $|\Sigma|$, the number of letters.

Representation. The words given as input are represented in a natural way: we work with the RAM model with word size $\log(|\Sigma| + n + \ell)$, which allows us to write a letter in each cell and to manipulate words and positions in a natural way. We write $|P|$ for $\sum_{j \in [1, n]} |u_j|$ and similarly $|N| = \sum_{j \in [1, n]} |v_j|$. The size of a sample is $|P| + |N|$.

We emphasise a subtlety on the representation of k : it can be given in binary (a standard assumption) or in unary. In the first case, the input size is $O(n \cdot \ell + \log(k))$, so the formula φ we are looking for may be exponential in the input size! This means that it is not clear a priori that the **LTL** learning problem is in NP. Opting for a unary encoding, the input size becomes $O(n \cdot \ell + k)$, and in that case an easy argument shows that the **LTL** learning problem is in NP. We follow the standard representation: k is given in binary, and therefore it is not immediate that the **LTL** learning problem is in NP.

A naive algorithm. Let us start our complexity analysis of the learning **LTL** problem by constructing a naive algorithm for the whole logic.

Theorem 1. *There exists an algorithm for solving the **LTL** learning problem in time and space $O(|\Sigma| + \exp(k) \cdot n \cdot \ell)$, where $\exp(k)$ is exponential in k .*

Notice that the dependence of the algorithm presented in Theorem 1 is linear in n and ℓ , and it is exponential only in k , but since k is represented in binary this is a doubly-exponential algorithm.

Proof. For a formula $\varphi \in \mathbf{LTL}$, we write $\langle \varphi \rangle : P \cup N \rightarrow \{0, 1\}^\ell$ for the function defined by

$$\langle \varphi \rangle(w)(i) = \begin{cases} 1 & \text{if } w[i \dots] \models \varphi, \\ 0 & \text{if } w[i \dots] \not\models \varphi, \end{cases}$$

for $w \in P \cup N$.

Note that φ is separating if and only if $\langle \varphi \rangle(u)(1) = 1$ and $\langle \varphi \rangle(v)(1) = 0$ for all $u \in P, v \in N$. The algorithm simply consists in enumerating all formulas φ of **LTL** of size at most k inductively, constructing $\langle \varphi \rangle$, and checking whether φ is separating. Initially, we

construct $\langle a \rangle$ for all $a \in \Sigma$, and then once we have computed $\langle \varphi \rangle$ and $\langle \psi \rangle$, we can compute $\langle \varphi \mathbf{U} \psi \rangle$, $\langle \mathbf{F} \varphi \rangle$, $\langle \mathbf{G} \varphi \rangle$, $\langle \mathbf{X} \varphi \rangle$, $\langle \varphi \wedge \psi \rangle$, and $\langle \varphi \vee \psi \rangle$ in time $O(n \cdot \ell)$. To conclude, we note that the number of formulas of **LTL** of size at most k is exponential in k . \square

Subwords

Let $u = u(1) \dots u(\ell)$ and $w = w(1) \dots w(\ell')$. We say that

- u is a *subword* of w if there exist $p_1 < \dots < p_\ell$ such that $u(i) = w(p_i)$ for all $i \in [1, \ell]$.
- u is a *weak subword* of w if there exist $p_1 \leq \dots \leq p_\ell$ such that $u(i) = w(p_i)$ for all $i \in [1, \ell]$.

For instance, *abba* is a subword of *babaaaaba*, and *bba* is a weak subword of *abaa* (using the *b* twice). We say that a word is non-repeating if every two consecutive letters are different. If u is non-repeating, then u is a weak subword of w if and only if it is a subword of w .

As a warm-up, let us construct simple **LTL** formulas related to subwords.

- We consider the word $u = u(1) \dots u(\ell)$ and $p_1 < p_2 < \dots < p_\ell$. They induce the following **LTL**(\mathbf{X}, \wedge) formula, called a *pattern*:

$$\mathcal{P} = \mathbf{X}^{p_1-1}(u(1) \wedge \mathbf{X}^{p_2-p_1}(\dots \wedge \mathbf{X}^{p_\ell-p_{\ell-1}}u(\ell)) \dots).$$

It is equivalent to the (larger in size) formula $\bigwedge_{i \in [1, \ell]} \mathbf{X}^{p_i-1}u(i)$, which states that for each $i \in [1, \ell]$, the letter in position p_i is $u(i)$.

- We consider the word $u = u(1) \dots u(\ell)$. It induces the following **LTL**(\mathbf{F}, \wedge) formula, called a *fattern* (pattern with an \mathbf{F}):

$$\mathcal{F} = \mathbf{F}(u(1) \wedge \mathbf{F}(\dots \wedge \mathbf{F}u(\ell)) \dots).$$

A word w satisfies \mathcal{F} if and only if u is a weak subword of w .

3. NP-hardness when the alphabet is part of the input

We prove our first hardness result: the learning problem for **LTL** is NP-hard for non-constant alphabets.

Theorem 2. *For all $\{\mathbf{F}, \vee\} \subseteq Op$, the **LTL**(Op) learning problem is NP-hard when the alphabet is part of the input.*

Recall that the hitting set problem takes as input C_1, \dots, C_n subsets of $[1, \ell]$ and $k \in \mathbb{N}$, and asks whether there exists $H \subseteq [1, \ell]$ of size at most k such that for every $j \in [1, n]$ we have $H \cap C_j \neq \emptyset$. It is known to be NP-complete.

Proof. We construct a reduction from the hitting set problem. Let $C_1, \dots, C_n \subseteq [1, \ell]$ and $k \in \mathbb{N}$ an instance of the hitting set problem. Let us define the alphabet $\Sigma = \{a_j, b_j : j \in [1, n]\}$, it has size $2n$. For $j \in [1, n]$, we define u_j of length ℓ by

$$u_j(i) = \begin{cases} a_j & \text{if } i \in C_j, \\ b_j & \text{otherwise.} \end{cases}$$

Let $v = b_1 \cdots b_n$. We claim that there exists a hitting set H of size at most k if and only if there exists a separating formula of **LTL** of size at most $2k$.

Given a hitting set H of size k , we construct the formula $\mathbf{F}(\bigvee_{j \in H} a_j)$, it is separating by definition of H being a hitting set, and it has size $2k$. Conversely, let us consider a separating formula φ of size $2k$. Since all operators have arity one or two, the syntactic tree of φ contains at most k leaves, so the set $H = \{j \in [1, \ell] : a_j \text{ appears in } \varphi\}$ has size at most k . Suppose H is not a hitting set of C_1, \dots, C_n , there exists $j \in [1, n]$ such that $H \cap C_j = \emptyset$. We prove that for all subformulas ψ of φ and for all $i \in [1, \ell]$, if $u_j[i \dots] \models \psi$ then $v[j \dots] \models \psi$. We proceed by induction.

- If ψ is a letter. If $u_j[i \dots] \models \psi$, this implies that $u_j(i) = b_j$ because a_j cannot appear in both ψ and u_j since $H \cap C_j = \emptyset$. Remark that $v(i) = b_i$, so $v[i \dots] \models \psi$.
- The other cases are easily proved by applying directly the induction hypothesis.

In particular, since u_j satisfies φ , then so does v , a contradiction with φ being separating. \square

This could be the end of our study! However the assumption that the alphabet is part of the input is very unusual, and in the rest of the paper we will therefore fix the alphabet size.

In the remainder of the paper, we will only consider fragments of **LTL** without the until operator. In other words, when consider Op a set of operators, we assume that $Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge, \vee\}$.

4. Membership in NP: the short formula property

Definition 1. Let Op a set of operator, we say that Op has the short formula property if there exists a polynomial P such that for all samples (P, N) , if there exists an **LTL**(Op) separating formula, then there exists one of size at most $P(|P| + |N|)$.

Note that if Op has the short formula property then the **LTL**(Op) learning problem is in NP.

Theorem 3. All $Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge, \vee\}$ have the short formula property, and therefore the **LTL**(Op) learning problem is in NP.

This theorem is the consequence of the following propositions, which establish it for various subsets of operators. We start with some technical lemmas. The following lemma is easily proved by induction, it shows that **LTL** cannot detect repetitions in the last letter of a word.

Lemma 1. Let u whose last letter is a . For all formulas $\varphi \in \mathbf{LTL}$, for all $k \in \mathbb{N}$, if u satisfies φ then ua^k satisfies φ .

Lemma 2. If a formula of the form $\varphi = \mathbf{F}(\psi_1 \wedge \mathbf{F}(\psi_2 \wedge (\cdots \wedge \mathbf{F}\psi_r) \dots))$ is not satisfied by some word v then there exist $i_1 < \cdots < i_k$ with $k \leq |v| + 1$ such that v does not satisfy $\psi = \mathbf{F}(\psi_{i_1} \wedge \mathbf{F}(\psi_{i_2} \wedge (\cdots \wedge \mathbf{F}\psi_{i_k}) \dots))$ but all words satisfying φ also satisfy ψ .

Proof. We set $v_1 = v$. For all $i \in [2, r]$ we set v_i as the largest suffix of v_{i-1} satisfying ψ_i . If this suffix does not exist, we set $v_i = \varepsilon$. As v does not satisfy φ , we easily obtain that for all $i < r$ v_i does not satisfy $\mathbf{F}(\psi_i \wedge \mathbf{F}(\psi_{i+1} \wedge \dots \wedge \mathbf{F}\psi_r) \dots)$. Hence v_{r-1} does not satisfy $\mathbf{F}\psi_r$ and thus $v_r = \varepsilon$. We can now extract a decreasing subsequence of suffixes v_{i_1}, \dots, v_{i_k} with $i_1 = 1$ and for all $j > 1$, i_j is the smallest index larger than i_{j-1} such that $v_{i_j} \neq v_{i_{j-1}}$ if it exists. Note that as this sequence is decreasing k cannot be larger than $|v| + 1$. We set $\psi = \mathbf{F}(\psi_{i_1} \wedge \mathbf{F}(\psi_{i_2} \wedge \dots \wedge \mathbf{F}\psi_{i_k}) \dots)$. We have that for all $j > 1$, v_{i_j} is the largest suffix of $v_{i_{j-1}}$ satisfying ψ_{i_j} , and $v_{i_k} = \varepsilon$. It is then easy to infer by reverse induction on j that for all j , v_{i_j} does not satisfy $\mathbf{F}(\psi_{i_j} \wedge \mathbf{F}(\psi_{i_{j+1}} \wedge \dots \wedge \mathbf{F}\psi_{i_k}) \dots)$ and thus that in particular $v_{i_1} = v$ does not satisfy ψ .

Let u be a word satisfying φ , there is a non-increasing sequence of suffixes u_i of u such that $u_1 = u$ and each u_i satisfies $\mathbf{F}(\psi_i \wedge \mathbf{F}(\psi_{i+1} \wedge \dots \wedge \mathbf{F}\psi_r) \dots)$. It is then easy to see, by reverse induction on j , that for each $j \in [1, k]$ u_{i_j} satisfies $\mathbf{F}(\psi_{i_j} \wedge \mathbf{F}(\psi_{i_{j+1}} \wedge \dots \wedge \mathbf{F}\psi_{i_k}) \dots)$. In particular $u_1 = u$ satisfies ψ . \square

We state and prove a dual version.

Lemma 3. *If a formula of the form $\varphi = \mathbf{G}(\psi_1 \vee \mathbf{G}(\psi_2 \vee \dots \vee \mathbf{G}\psi_r) \dots)$ is satisfied by some word u then there exist $i_1 < \dots < i_k$ with $k \leq |u| + 1$ such that u satisfies $\psi = \mathbf{G}(\psi_{i_1} \vee \mathbf{G}(\psi_{i_2} \vee \dots \vee \mathbf{G}\psi_{i_k}) \dots)$ and all words satisfying ψ also satisfy φ .*

Proof. Let us momentarily use negation. Note that the proof of Lemma 2 still holds when we allow negations. The word u does not satisfy $\varphi' = \mathbf{F}(\neg\psi_1 \wedge \mathbf{F}(\neg\psi_2 \vee \dots \wedge \mathbf{F}\neg\psi_r) \dots)$ as it is equivalent to the negation of φ , thus there exist $i_1 < \dots < i_k$ with $k \leq |u| + 1$ such that u does not satisfy $\psi' = \mathbf{F}(\psi_{i_1} \wedge \mathbf{F}(\psi_{i_2} \wedge \dots \wedge \mathbf{F}\psi_{i_k}) \dots)$ but all words satisfying φ' satisfy ψ' . The formula $\psi = \mathbf{G}(\psi_{i_1} \vee \mathbf{G}(\psi_{i_2} \vee \dots \vee \mathbf{G}\psi_{i_k}) \dots)$ is equivalent to the negation of ψ' , thus it is satisfied by u , and all words satisfying ψ satisfy φ . \square

The following fact will be useful for obtaining weak normal forms.

Fact 1. *For all formulas $\varphi, \psi, \varphi_1, \dots, \varphi_n \in \mathbf{LTL}$, the following equivalences hold:*

1. $\mathbf{GG}\varphi \equiv \mathbf{G}\varphi$ and $\mathbf{FF}\varphi \equiv \mathbf{F}\varphi$.
2. $\mathbf{X}(\varphi \wedge \psi) \equiv \mathbf{X}\varphi \wedge \mathbf{X}\psi$ and $\mathbf{X}(\varphi \vee \psi) \equiv \mathbf{X}\varphi \vee \mathbf{X}\psi$.
3. $\mathbf{GX}\varphi \equiv \perp$ and $\mathbf{FX}\varphi \equiv \mathbf{XF}\varphi$.
4. $\mathbf{GF}\varphi \equiv \mathbf{FG}\varphi$ (both state that ψ is satisfied on the last position)
5. $\mathbf{G}(\varphi \wedge \psi) \equiv \mathbf{G}\varphi \wedge \mathbf{G}\psi$ and $\mathbf{F}(\varphi \vee \psi) \equiv \mathbf{F}\varphi \vee \mathbf{F}\psi$.
6. $\mathbf{F}(\psi \wedge \bigwedge_{i=1}^n \mathbf{F}\varphi_i) \equiv \bigwedge_{i=1}^n \mathbf{F}(\psi \wedge \mathbf{F}\varphi_i)$ and $\mathbf{G}(\psi \vee \bigvee_{i=1}^n \mathbf{G}\varphi_i) \equiv \bigvee_{i=1}^n \mathbf{G}(\psi \vee \mathbf{G}\varphi_i)$.

Proposition 1. *All $Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge\}$ have the short formula property.*

Proof. Let us consider a formula φ separating P from N in $\mathbf{LTL}(Op)$, we apply a series of transformations to φ to construct another separating formula of polynomial size.

First of all we push the \mathbf{G} in front of the letters, which is possible by using repeatedly the equivalences 1, 3, 4 and 5 of Fact 1. Then we push the \mathbf{X} to the bottom of the formula

as well, as they commute with \mathbf{F} and \wedge . We obtain a $\mathbf{LTL}(\mathbf{F}, \wedge)$ formula with atoms of the form either $\mathbf{X}^k \mathbf{G}a$ or $\mathbf{X}^k a$. Once that is done, we make the formula follow a normal form by repeatedly using equivalence 6 of Fact 1. While φ has a subformula of the form $\mathbf{F}(\psi \wedge \bigwedge_{i=1}^p \mathbf{F}\varphi_i)$, we turn it into $\bigwedge_{i=1}^p \mathbf{F}(\psi \wedge \mathbf{F}\varphi_i)$. In the end φ is of the form $\psi \wedge \bigwedge_{i=1}^p \varphi_i$ with each φ_i of the form $\mathbf{F}(\psi_1 \wedge \mathbf{F}(\psi_2 \wedge \dots \wedge \mathbf{F}\psi_r) \dots)$ where ψ and the ψ_i are conjunctions of formulas of the form $\mathbf{X}^k \mathbf{G}a$ or $\mathbf{X}^k a$. Note that p may be exponential in $|\varphi|$.

Observe that given two formulas α and β of the form either a or $\mathbf{G}a$ for some $a \in \Sigma$, $\mathbf{X}^k \alpha \wedge \mathbf{X}^k \beta$ is equivalent to either $\mathbf{X}^k \alpha$, $\mathbf{X}^k \beta$, or \perp . It is then easy to see that every conjunction of formulas of the form $\mathbf{X}^k \mathbf{G}a$ or $\mathbf{X}^k a$ is equivalent over non-empty words of length at most ℓ to either \perp or a conjunction of at most ℓ formulas of the form $\mathbf{X}^k a$ or $\mathbf{X}^k \mathbf{G}a$ with $k < \ell$. We can thus assume ψ and all ψ_i to be of that form, and thus to be of size polynomial in ℓ .

This formula being equivalent to φ , all u_i satisfy all φ_r and ψ , and for all v_j either v_j does not satisfy ψ or it does not satisfy some φ_r . We focus on the second case. Say v_j does not satisfy $\varphi_r = \mathbf{F}(\psi_1 \wedge \mathbf{F}(\psi_2 \wedge \dots \wedge \mathbf{F}\psi_s) \dots)$, by Lemma 2 we can turn φ_r into a short formula φ'_j that is still satisfied by all u_i but not by v_j . We set φ'_j to be as described above for all v_j that satisfy ψ and to be ψ for the others. As a result, the formula $\bigwedge_{i=1}^m \varphi'_i$ is a formula of polynomial size in m and ℓ separating the u_i and v_j .

Furthermore, we did not add any operators to the formula, hence the final formula uses the same set of operators as φ . \square

Proposition 2. *All $Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \vee\}$ have the short formula property.*

Proof. The proof is nearly identical to the one of Proposition 1. We start by pushing the \mathbf{F} and \mathbf{X} to the bottom of the formula to obtain a formula of $\mathbf{LTL}(\mathbf{G}, \vee)$ with atoms of the form $\mathbf{X}^k a$ or $\mathbf{X}^k \mathbf{F}a$. We also use repeatedly equivalence 6 of Fact 1 to obtain a formula of the form $\psi \vee \bigvee_{i=1}^p \varphi_i$ with each φ_i of the form $\mathbf{G}(\psi_1 \vee \mathbf{G}(\psi_2 \vee \dots \vee \mathbf{G}\psi_r) \dots)$ where ψ and the ψ_i are disjunctions of formulas of the form $\mathbf{X}^k \mathbf{F}a$ or $\mathbf{X}^k a$. Then we note that every disjunction of formulas of the form $\mathbf{X}^k a$ and $\mathbf{X}^k \mathbf{F}a$ can be replaced by a disjunction of at most $\ell|\Sigma|$ such formulas, equivalent on non-empty words of length at most ℓ (as $\mathbf{X}^k \mathbf{F}a \vee \mathbf{X}^k a$ is equivalent to $\mathbf{X}^k \mathbf{F}a$). We can thus assume that all ψ_i and ψ mentioned above are of polynomial size in ℓ and $|\Sigma|$. No v_j satisfies either ψ or any ψ_r . For each u_i either there exists φ_r that is satisfied by u_i or u_i satisfies ψ . In the first case we use Lemma 3 to turn that φ_r into a formula φ'_i of polynomial size in ℓ and $|\Sigma|$ satisfied by u_i but not by any v_j . If u_i does not satisfy any φ_r we set $\varphi'_i = \psi$. In the end the formula $\bigvee_{i=1}^n \varphi'_i$ is satisfied by all u_i but no v_j , and is of polynomial size in ℓ , n and $|\Sigma|$. \square

Proposition 3. *All $\{\mathbf{G}, \mathbf{X}, \wedge, \vee\} \subseteq Op$ have the short formula property.*

Proof. Let (P, N) a sample. Consider the formula $\varphi = \bigvee_{i=1}^n \varphi_{u_i}$ where $\varphi_{u_i} = \bigwedge_{j=1}^{m-1} \mathbf{X}^{j-1} a_j \wedge \mathbf{X}^{m-1} \mathbf{G}a_m$ with $u_i = a_1 \dots a_m$. If there exist u_i, v_j such that $v_j \in u_i a^*$ with a the last letter of u_i , then there is no separating formula: By Lemma 1 every formula satisfied by u_i is also satisfied by v_j . Otherwise, every u_i satisfies the associated formula φ_{u_i} while no v_j satisfies any of them. Hence φ is a separating formula of polynomial size. \square

Proposition 4. *All Op such that $\{\mathbf{X}, \wedge, \vee\} \subseteq Op \subseteq \{\mathbf{F}, \mathbf{X}, \wedge, \vee\}$ have the short formula property.*

Proof. Let (P, N) a sample. Consider the formula $\varphi = \bigvee_{i=1}^n \varphi_{u_i}$ where $\varphi_{u_i} = \bigwedge_{j=1}^m \mathbf{X}^{j-1} a_j$ with $u_i = a_1 \cdots a_m$. We distinguish two cases.

- If there exist u_i, v_j such that u_i is a prefix of v_j , then there is no separating formula: An easy induction shows that every formula satisfied by u_i is also satisfied by v_j .
- Otherwise, every u_i satisfies the associated formula φ_{u_i} while no v_j satisfies any of them. Hence φ is a separating formula of polynomial size.

□

Fact 2. For all $\varphi, \varphi_1, \varphi_2 \in \mathbf{LTL}(\mathbf{F}, \mathbf{G}, \wedge, \vee, \neg)$, the following equivalences hold:

- $\neg(\varphi_1 \wedge \varphi_2) \equiv \neg\varphi_1 \vee \neg\varphi_2$
- $\neg(\varphi_1 \vee \varphi_2) \equiv \neg\varphi_1 \wedge \neg\varphi_2$
- $\neg\mathbf{F}\varphi \equiv \mathbf{G}\neg\varphi$
- $\neg\mathbf{G}\varphi \equiv \mathbf{F}\neg\varphi$

Proposition 5. $\{\mathbf{F}, \wedge, \vee\}$ and $\{\mathbf{G}, \wedge, \vee\}$ have the short formula property.

Proof. Let (P, N) a sample. Let $j \in [1, n]$, if there exists w_j a weak subword of all u_i that is not a weak subword of v_j then we set $\psi(w_j) = \mathbf{F}(a_1 \wedge \mathbf{F}(a_2 \wedge \cdots \mathbf{F}a_k))$ with $w_j = a_1 \cdots a_k$. This formula is satisfied by exactly the words which have w_j as a weak subword, hence by all u_i but not v_j . We set φ_j as $\psi(w_j)$ if it exists and a if all u_i start with an a and v_j does not for some $a \in \Sigma$.

If φ_j is defined for all j then the formula $\bigwedge_j \varphi_j$ is a separating formula. Otherwise there exists a v_j such that there are no weak subwords of all the u_i that are not weak subwords of v_j and either the u_i do not start with the same letter or v_j also starts with that letter. In that case we can easily prove by induction that all $\mathbf{LTL}(\mathbf{F}, \wedge, \vee)$ formulas of the form $\mathbf{F}\psi$ that are satisfied by all u_i are also satisfied by v_j , and as all formulas of the form a with $a \in \Sigma$ satisfied by all u_i are also satisfied by v_j , the same can be said about boolean combinations of those two types of formulas, and thus there are no separating formulas in $\mathbf{LTL}(\mathbf{F}, \wedge, \vee)$.

Concerning $\{\mathbf{G}, \wedge, \vee\}$, one can easily infer from Fact 2 that we can turn any formula of $\mathbf{LTL}(\mathbf{F}, \wedge, \vee)$ into one of $\mathbf{LTL}(\mathbf{G}, \wedge, \vee)$ equivalent to its negation, and vice-versa. Hence the short formula property of $\{\mathbf{G}, \wedge, \vee\}$ is a consequence of the one of $\{\mathbf{F}, \wedge, \vee\}$. □

Proposition 6. $\{\mathbf{F}, \mathbf{G}, \wedge, \vee\}$ has the short formula property.

Proof. Let (P, N) a sample. For all u_i we set $\varphi(u_i) = \mathbf{F}(a_1 \wedge \mathbf{F}(a_2 \wedge \cdots \mathbf{F}a_k))$ with $u_i = a_1 \cdots a_k$. This formula is satisfied by exactly the words with u_i as a weak subword. For all v_j we set $\varphi(v_j) = \mathbf{G}(\bar{a}_1 \vee \mathbf{G}(\bar{a}_2 \vee \cdots \mathbf{G}\bar{a}_k))$ with $v_j = a_1 \cdots a_k$ and $\bar{a} = \bigvee_{b \in \Sigma \setminus \{a\}} b$ for all $a \in \Sigma$. This formula is satisfied by exactly the words which do not have v_j as a weak subword. For all pairs (u_i, v_j) we set $\psi_{i,j}$ as $\varphi(u_i) \wedge \varphi(v_j)$, and finally we define ψ as $\bigvee_i \bigwedge_j \psi_{i,j}$. If this formula does not separate the u_i and v_j then there must exist u_i and v_j which are weak subwords of each other. An easy induction on the formula then shows that u_i and v_j satisfy the same formulas in $\mathbf{LTL}(\mathbf{F}, \mathbf{G}, \wedge, \vee)$, and thus there does not exist any separating formula. □

5. Degenerate cases: the shortest formula property

Definition 2. Let Op a set of operators, we say that Op has the shortest formula property if there is a polynomial time algorithm solving the $\mathbf{LTL}(Op)$ learning problem and outputting the minimal separating formula if it exists.

Theorem 4. Let Op such that either $Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}\}$, $Op \subseteq \{\vee, \wedge\}$, $Op = \{\mathbf{G}, \wedge\}$. Then Op has the shortest formula property, and therefore the $\mathbf{LTL}(Op)$ learning problem is in \mathbf{P} . The same holds for $Op = \{\mathbf{F}, \vee\}$ if the size of the alphabet is fixed.

This theorem is the consequence of the following propositions, which establish it for various subsets Op of operators.

Proposition 7. All $Op \subseteq \{\vee, \wedge\}$ have the shortest formula property.

Proof. Let (P, N) a sample. Let $A \subseteq \Sigma$ be the set of first letters of the u_i . Consider the formula $\psi = \bigvee_{a \in A} a$. We distinguish two cases.

- Either ψ does not separate the u_i and v_i , meaning that some u_i and v_j share the same first letter. In that case u_i and v_j satisfy the same formulas of $\mathbf{LTL}(\vee, \wedge)$ and thus there is no separating formula.
- Or ψ does separate the u_i and v_j . In that case it is also of minimal size: say some element a of A does not appear in a formula $\varphi \in \mathbf{LTL}(\vee, \wedge)$, then as it is satisfied by some u_i starting with a , ψ is a tautology and is also satisfied by the v_j . Hence a separating formula has to contain all letters of A , and thus also at least $|A| - 1$ boolean operators. As a result, ψ is of minimal size.

□

Proposition 8. All $Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}\}$ have the shortest formula property.

Proof. Let $\ell \in \mathbb{N}$, every formula $\varphi \in \mathbf{LTL}(\mathbf{F}, \mathbf{G}, \mathbf{X})$ is equivalent over words of length at most ℓ to a formula of smaller or equal size in

$$\left\{ \top, \perp, \mathbf{X}^k a, \mathbf{X}^k \mathbf{F} a, \mathbf{X}^k \mathbf{G} a, \mathbf{X}^k \mathbf{F} \mathbf{G} a \mid a \in \Sigma, 0 \leq k < \ell \right\}.$$

This can be shown by observing that over those words, for all ψ , $\mathbf{F}\mathbf{X}\psi$ and $\mathbf{X}\mathbf{F}\psi$ are equivalent, as well as $\mathbf{G}\mathbf{X}\psi$ and \perp . This allows to push all \mathbf{X} operators at the top of the formula. Finally, every formula of the form $\mathbf{X}^k \psi$ with $k \geq N$ is equivalent to \perp .

As a result we can compute a minimal separating formula by enumerating formulas from this set (of polynomial size) and checking which ones separate the positive and negative words. □

Proposition 9. $\{\mathbf{G}, \wedge\}$ has the shortest formula property.

Proof. An easy induction on ψ shows that all $\mathbf{LTL}(\mathbf{G}, \wedge)$ formulas are equivalent over finite words to a formula of the form \top , \perp , a or $\mathbf{G}a$ with $a \in \Sigma$, i.e., to a formula of size at most 2. As those can be enumerated in polynomial time, we can compute a separating formula of minimal size or conclude that it does not exist in polynomial time. □

A corollary of our results is that the classification between P and NP mostly does not depend upon whether we consider the alphabet as part of the input or not. The only affected subcase is $Op = \{\mathbf{F}, \vee\}$, which is NP-hard when the alphabet is not fixed (Theorem 2).

Proposition 10. $\{\mathbf{F}, \vee\}$ has the shortest formula property when the alphabet is fixed.

Proof. Since \mathbf{F} commutes with \vee , every formula can be turned into an equivalent disjunction of formulas of the form a or $\mathbf{F}a$ with $a \in \Sigma$, of polynomial size. There are only $2^{2^{|\Sigma|}}$ such formulas, hence when given a sample (P, N) one can simply select the formulas that are not satisfied by any v_j , take their disjunction, and check that all u_i satisfy the resulting formula. If they do we have a separating formula, otherwise there cannot exist any. \square

6. An approximation algorithm for $\mathbf{LTL}(\mathbf{X}, \wedge)$

An α -approximation algorithm for learning a fragment of \mathbf{LTL} does the following: the algorithm either determines that there are no separating formulas, or constructs a separating formula φ which has size at most $\alpha \cdot m$ with m the size of a minimal separating formula.

Theorem 5. There exists a $O(n \cdot \ell^2)$ time $\log(n)$ -approximation algorithm for learning $\mathbf{LTL}(\mathbf{X}, \wedge)$.

Recall the definition of patterns: a word $u = u(1) \dots u(\ell)$ and $p_1 < p_2 < \dots < p_\ell$ induce the following $\mathbf{LTL}(\mathbf{X}, \wedge)$ formula, called a *pattern*:

$$\mathcal{P} = \mathbf{X}^{p_1-1}(u(1) \wedge \mathbf{X}^{p_2-i_1}(\dots \wedge \mathbf{X}^{p_\ell-p_{\ell-1}}u(\ell)) \dots).$$

It is equivalent to the (larger in size) formula $\bigwedge_{i \in [1, \ell]} \mathbf{X}^{p_i-1}u(i)$, which states that for each $i \in [1, \ell]$, the letter in position p_i is $u(i)$. To determine the size of a pattern \mathcal{P} we look at two parameters: its last position $\mathbf{last}(\mathcal{P}) = i_p$ and its width $\mathbf{width}(\mathcal{P}) = p$. The size of \mathcal{P} is $\mathbf{last}(\mathcal{P}) + 2(\mathbf{width}(\mathcal{P}) - 1)$. The two parameters of a pattern, last position and width, hint at the key trade-off we will have to face in learning $\mathbf{LTL}(\mathbf{X}, \wedge)$ formulas: do we increase the last position, to reach further letters in the words, or the width, to further restrict the set of satisfying words?

Lemma 4. For every formula $\varphi \in \mathbf{LTL}(\mathbf{X}, \wedge)$ there exists an equivalent pattern of size smaller than or equal to φ .

Proof. We proceed by induction on φ .

- Atomic formulas are already a special case of patterns.
- If $\varphi = \mathbf{X}\varphi'$, by induction hypothesis we get a pattern \mathcal{P} equivalent to φ' , then $\mathbf{X}\mathcal{P}$ is a pattern and equivalent to φ .
- If $\varphi = \varphi_1 \wedge \varphi_2$, by induction hypothesis we get two patterns \mathcal{P}_1 and \mathcal{P}_2 equivalent to φ_1 and φ_2 . We use the inductive definition for patterns to show that $\mathcal{P}_1 \wedge \mathcal{P}_2$ is equivalent to another pattern. We focus on the case $\mathcal{P}_1 = \mathbf{X}^{i_1}(c_1 \wedge \mathcal{P}'_1)$ and $\mathcal{P}_2 = \mathbf{X}^{i_2}(c_2 \wedge \mathcal{P}'_2)$, the other cases are simpler instances of this one.

There are two cases: $i_1 = i_2$ or $i_1 \neq i_2$.

If $i_1 = i_2$, either $c_1 \neq c_2$ and then $\mathcal{P}_1 \wedge \mathcal{P}_2$ is equivalent to false, which is the pattern $c_1 \wedge c_2$, or $c_1 = c_2$, and then $\mathcal{P}_1 \wedge \mathcal{P}_2$ is equivalent to $\mathbf{X}^{i_1}(c_1 \wedge \mathcal{P}'_1 \wedge \mathcal{P}'_2)$. By induction hypothesis $\mathcal{P}'_1 \wedge \mathcal{P}'_2$ is equivalent to a pattern \mathcal{P}' , so the pattern $\mathbf{X}^{i_1}(c_1 \wedge \mathcal{P}')$ is equivalent to $\mathcal{P}_1 \wedge \mathcal{P}_2$, hence to φ .

If $i_1 \neq i_2$, without loss of generality $i_1 < i_2$, then $\mathcal{P}_1 \wedge \mathcal{P}_2$ is equivalent to $\mathbf{X}^{i_1}(c_1 \wedge \mathcal{P}'_1 \wedge \mathbf{X}^{i_2-i_1}(c_2 \wedge \mathcal{P}'_2))$. By induction hypothesis $\mathcal{P}'_1 \wedge \mathbf{X}^{i_2-i_1}(c_2 \wedge \mathcal{P}'_2)$ is equivalent to a pattern \mathcal{P}' , so the pattern $\mathbf{X}^{i_1}(c_1 \wedge \mathcal{P}')$ is equivalent to $\mathcal{P}_1 \wedge \mathcal{P}_2$, hence to φ .

□

Proof. Let $u_1, \dots, u_n, v_1, \dots, v_n$ a set of $2n$ words of length at most ℓ . Thanks to Lemma 4 we are looking for a pattern. For a pattern \mathcal{P} we define $I(\mathcal{P}) = \{i_q \in [1, \ell] : q \in [1, p]\}$. Note that $\mathbf{last}(\mathcal{P}) = \max I(\mathcal{P})$ and $\mathbf{width}(\mathcal{P}) = \text{Card}(I(\mathcal{P}))$.

We define the set $X = \{i \in [1, \ell] : \exists c \in \Sigma, \forall j \in [1, n], u_j(i) = c\}$. Note that \mathcal{P} satisfies u_1, \dots, u_n if and only if $I(\mathcal{P}) \subseteq X$. Further, given $I \subseteq X$, we can construct a pattern \mathcal{P} such that $I(\mathcal{P}) = I$ and \mathcal{P} satisfies u_1, \dots, u_n : we simply choose $c_q = u_1(i_q) = \dots = u_n(i_q)$ for $q \in [1, p]$. We call \mathcal{P} the pattern corresponding to I .

Recall that the size of the pattern \mathcal{P} is $\mathbf{last}(\mathcal{P}) + 2(\mathbf{width}(\mathcal{P}) - 1)$. This makes the task of minimising it difficult: there is a trade-off between minimising the last position $\mathbf{last}(\mathcal{P})$ and the width $\mathbf{width}(\mathcal{P})$.

Let us consider the following easier problem: construct a $\log(n)$ -approximation of a minimal separating pattern with fixed last position. Assuming we have such an algorithm, we obtain a $\log(n)$ -approximation of a minimal separating pattern by running the previous algorithm on prefixes of length ℓ' for each $\ell' \in [1, \ell]$.

We now focus on the question of constructing a $\log(n)$ -approximation of a minimal separating pattern with fixed last position. For a set I , we write $C_I = \bigcup \{Y_i : i \in I\}$: the pattern corresponding to I does not satisfy v_j if and only if $j \in C_I$. In particular, the pattern corresponding to I is separating if and only if $C_I = [1, n]$.

The algorithm constructs a set I incrementally through the sequence $(I_x)_{x \geq 0}$, with the following easy invariant: for $x \geq 0$, we have $C_x = C_{I_x}$. The algorithm is greedy: I_x is augmented with $i \in X \setminus I_x$ maximising the number of words added to C_x by adding i , which is the cardinality of $Y_i \setminus C_x$.

We now prove that this yields a $\log(n)$ -approximation algorithm. Let \mathcal{P}_{opt} a minimal separating pattern with last position ℓ , inducing $I_{\text{opt}} = I(\mathcal{P}_{\text{opt}}) \subseteq [1, \ell]$ of cardinal m . Note that $C_{I_{\text{opt}}} = [1, n]$.

We let $n_x = n - |C_x|$ and show the following by induction on $x \geq 0$:

$$n_{x+1} \leq n_x \cdot \left(1 - \frac{1}{m}\right) = n_x \cdot \frac{m-1}{m}.$$

We claim that there exists $i \in X \setminus I_x$ such that $\text{Card}(Y_i \setminus C_x) \geq \frac{n_x}{m}$. Indeed, assume towards contradiction that for all $i \in X \setminus I_x$ we have $\text{Card}(Y_i \setminus C_x) < \frac{n_x}{m}$, then there are no sets I of cardinal m such that $C_I \supseteq [1, n] \setminus C_x$, contradicting the existence of I_{opt} . Thus there exists $i \in X \setminus I_x$ such that $\text{Card}(Y_i \setminus C_x) \geq \frac{n_x}{m}$, implying that the algorithm chooses such an i and $n_{x+1} \leq n_x - \frac{n_x}{m} = n_x \cdot \left(1 - \frac{1}{m}\right)$.

The proved inequality implies $n_x \leq n \cdot \left(1 - \frac{1}{m}\right)^x$. This quantity is less than 1 for $x \geq \log(n) \cdot m$, implying that the algorithm stops after at most $\log(n) \cdot m$ steps. Consequently, the pattern corresponding to I has size at most $\log(n) \cdot |\mathcal{P}_{\text{opt}}|$, completing the claim on approximation.

A naive complexity analysis yields an implementation of the greedy algorithm running in time $O(n \cdot \ell)$, leading to an overall complexity of $O(n \cdot \ell^2)$ by running the greedy algorithm on the prefixes of length ℓ' of $u_1, \dots, u_n, v_1, \dots, v_n$ for each $\ell' \in [1, \ell]$. \square

7. Almost all fragments with the next operator are hard to approximate

Theorem 6. *The $LTL(\mathbf{X}, \wedge)$ learning problem is NP-hard, and there are no $(1 - o(1)) \cdot \log(n)$ polynomial time approximation algorithms unless $\mathbf{P} = \mathbf{NP}$, even for a single positive word.*

Note that Theorem 5 and Theorem 6 yield matching upper and lower bounds on approximation algorithms for learning $LTL(\mathbf{X}, \wedge)$.

The hardness result stated in Theorem 6 follows from a reduction to the set cover problem, that we define now. The set cover decision problem is: given S_1, \dots, S_ℓ subsets of $[1, n]$ and $k \in \mathbb{N}$, does there exist $I \subseteq [1, \ell]$ of size at most k such that $\bigcup_{i \in I} S_i = [1, n]$? In that case we say that I is a cover. An α -approximation algorithm returns a cover of size at most $\alpha \cdot k$ where k is the size of a minimal cover. The following results form the state of the art for solving exact and approximate variants of the set cover problem.

Theorem 7 ((Dinur & Steurer, 2014)). *The set cover problem is NP-complete, and there are no $(1 - o(1)) \cdot \log(n)$ polynomial time approximation algorithms unless $\mathbf{P} = \mathbf{NP}$.*

Proof. We construct a reduction from set cover. Let S_1, \dots, S_ℓ subsets of $[1, n]$ and $k \in \mathbb{N}$.

Let us consider the word $u = a^{\ell+1}$, and for each $j \in [1, n]$ and $i \in [1, \ell]$, writing $v_j(i)$ for the i^{th} letter of v_j :

$$v_j(i) = \begin{cases} b & \text{if } j \in S_i, \\ a & \text{if } j \notin S_i, \end{cases}$$

and we set $v_j(\ell + 1) = a$ for any $j \in [1, n]$. We also add $v_{n+1} = a^\ell b$.

We claim that there is a cover of size k if and only if there is a formula of size $\ell + 2k - 1$ separating u from v_1, \dots, v_{n+1} .

Thanks to Lemma 4 we can restrict our attention to patterns, *i.e.* formulas of the form (we adjust the indexing for technical convenience)

$$\varphi = \mathbf{X}^{i_1-1}(c_1 \wedge \mathbf{X}^{i_2-i_1}(\dots \wedge \mathbf{X}^{i_{p+1}-i_p} c_{p+1}) \dots),$$

for some positions $i_1 \leq \dots \leq i_{p+1}$ and letters $c_1, \dots, c_{p+1} \in \Sigma$. If φ satisfies u , then necessarily $c_1 = \dots = c_{p+1} = a$. This implies that if φ does not satisfy v_{n+1} , then necessarily $i_{p+1} = \ell + 1$.

We associate to φ the set $I = \{i_1 \leq \dots \leq i_p\}$. Note that φ is equivalent to $\bigwedge_{q \in [1, p]} \mathbf{X}^{i_q-1} a \wedge \mathbf{X}^\ell a$, and the size of φ is $\ell + 1 + 2(|I| - 1)$.

By construction, φ separates u from v_1, \dots, v_{n+1} if and only if I is a cover. Indeed, I is a cover if and only if for every $j \in [1, n]$ there exists $i \in I$ such that $j \in S_i$, which is equivalent to for every $j \in [1, n]$ we have $v_j \not\models \varphi$. \square

We can extend the previous hardness result to all sets of operators Op such that $\{\mathbf{X}, \wedge\} \subseteq Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge, \vee\}$, by reducing their learning problems to the previous one. The reduction consists in transforming the input words so that the \mathbf{F} and \mathbf{G} operators are essentially useless. The \vee operator is also useless since we only have one positive word, implying that the minimal separating formulas in $\mathbf{LTL}(Op)$ and $\mathbf{LTL}(\mathbf{X}, \wedge)$ are in fact the same.

The first step is a reduction lemma for disjunctions.

Lemma 5. *For all $\varphi \in \mathbf{LTL}(\mathbf{X}, \wedge, \vee)$, for all u, v_1, \dots, v_n , if φ separates u from v_1, \dots, v_n , then there exists $\psi \in \mathbf{LTL}(\mathbf{X}, \wedge)$ such that $|\psi| \leq |\varphi|$ which separates u from v_1, \dots, v_n .*

Proof. We define $D(\varphi) \subseteq \mathbf{LTL}(\mathbf{X}, \wedge)$ by induction:

- If $\varphi = c$ then $D(\varphi) = \{c\}$.
- If $\varphi = \varphi_1 \wedge \varphi_2$ then $D(\varphi) = \{\psi_1 \wedge \psi_2 : \psi_1 \in D(\varphi_1), \psi_2 \in D(\varphi_2)\}$.
- If $\varphi = \varphi_1 \vee \varphi_2$ then $D(\varphi) = D(\varphi_1) \cup D(\varphi_2)$.
- If $\varphi = \mathbf{X}\varphi'$ then $D(\varphi) = \{\mathbf{X}\psi : \psi \in D(\varphi')\}$.
- If $\varphi = \mathbf{F}\varphi'$ then $D(\varphi) = \{\mathbf{F}\psi : \psi \in D(\varphi')\}$.

Observe that all formulas of $D(\varphi)$ are of size at most $|\varphi|$, and are in $\mathbf{LTL}(\mathbf{X}, \wedge)$. We now show that for all φ , for all u, v_1, \dots, v_n , if φ separates u from v_1, \dots, v_n then there exists $\psi \in D(\varphi)$ separating them, which proves the lemma.

We proceed by induction on φ .

- If $\varphi = c$ this is clear.
- If $\varphi = \varphi_1 \wedge \varphi_2$ then $D(\varphi) = \{\psi_1 \wedge \psi_2 : \psi_1 \in D(\varphi_1), \psi_2 \in D(\varphi_2)\}$. Since φ separates u from v_1, \dots, v_n , there exists $I_1, I_2 \subseteq [1, n]$ such that $I_1 \cup I_2 = [1, n]$, φ_1 separates u from $\{v_i : i \in I_1\}$, and φ_2 separates u from $\{v_i : i \in I_2\}$. By induction hypothesis applied to both φ_1 and φ_2 there exists $\psi_1 \in D(\varphi_1)$ separating u from $\{v_i : i \in I_1\}$ and $\psi_2 \in D(\varphi_2)$ separating u from $\{v_i : i \in I_2\}$. It follows that $\psi_1 \wedge \psi_2$ separates u from v_1, \dots, v_n , and $\psi_1 \wedge \psi_2 \in D(\varphi)$.
- If $\varphi = \varphi_1 \vee \varphi_2$ then $D(\varphi) = D(\varphi_1) \cup D(\varphi_2)$. Since φ separates u from v_1, \dots, v_n , either φ_1 or φ_2 does as well; without loss of generality let us say that φ_1 separates u from v_1, \dots, v_n . The induction hypothesis implies that $\psi_1 \in D(\varphi_1)$ separates u from v_1, \dots, v_n , and $\psi_1 \in D(\varphi)$.
- The case $\varphi = \mathbf{X}\varphi'$ follows directly by induction hypothesis.

□

Proposition 11. *For all $\{\mathbf{X}, \wedge\} \subseteq Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge, \vee\}$, there is a polynomial-time reduction from the $\mathbf{LTL}(\mathbf{X}, \wedge)$ learning problem for a single positive word to the $\mathbf{LTL}(Op)$ one over the same alphabet.*

We recall two (folklore) facts about **LTL**. They can easily be proven by induction on the formula.

Fact 3. For all $w_1 \in \Sigma^*$, $w_2 \in \Sigma^+$, $N \in \mathbb{N}$, and $\varphi \in \mathbf{LTL}$ with $|\varphi| \leq N$, then $w_1 w_2^N \models \varphi$ if and only if $w_1 w_2^{N+1} \models \varphi$.

Fact 4. For all $w_1 \in \Sigma^*$, $N \in \mathbb{N}$, and $\varphi \in \mathbf{LTL}(\mathbf{X}, \wedge, \vee)$ with at most $N + 1$ operators \mathbf{X} in φ , then $w_1 \models \varphi$ if and only if $w_1^{\leq N} \models \varphi$.

We now proceed to the proof of Proposition 11.

Proof. Let $u, v_1, \dots, v_n \in \Sigma^*$ be words, all of length ℓ . Let $a \in \Sigma$, we set M as the size of the formula

$$\psi_u = \bigwedge_{i=0}^{|u|-1} X^i u_i,$$

$u' = ua^M$, for all i , we define $v'_i = v_i a^M$,

$$\mathbf{u} = (u' v'_1 \dots v'_n)^{M+1} \quad ; \quad \mathbf{v}_i = v'_i \dots v'_n (u' v'_1 \dots v'_n)^M.$$

The formula ψ_u separates u from $\{v_1, \dots, v_n\}$ unless one of the v_i is equal to u . This can be checked in polynomial time, and in that case we can answer no to the learning problem immediately. In all that follows we assume that ψ_u separates u from the v_i .

Let $\varphi \in \mathbf{LTL}(Op)$ be a formula of minimal size separating \mathbf{u} from the $(\mathbf{v}_i)_{1 \leq i \leq n}$. Note that ψ_u is satisfied by \mathbf{u} but not by any \mathbf{v}_i , thus since Op contains \mathbf{X} and \wedge , φ exists and $|\varphi| \leq |\varphi_u| = M$.

We first show that φ contains no \mathbf{F} or \mathbf{G} . Let $C \in \mathbf{LTL}(\mathbf{X}, \wedge, \vee)$ be a context with free variables x_1, \dots, x_m (each appearing exactly once in C) such that $\varphi = C[\psi_1 \rightarrow x_1, \dots, \psi_m \rightarrow x_m]$ for some $\psi_i \in \mathbf{LTL}(Op)$ of the form either $\mathbf{F}\psi'$ or $\mathbf{G}\psi'$.

As the \mathbf{X} operator commutes with all the others, we can push the \mathbf{X} to the variables in C . Hence there exist a boolean context B with free variables y_1, \dots, y_m (each appearing exactly once in B) and $j_1, \dots, j_m \in \mathbb{N}$ such that $B[X^{j_1} x_1 \rightarrow y_1, \dots, X^{j_m} x_m \rightarrow y_m]$ is equivalent to C . Furthermore, the formulas B and C have the same depth, thus as C is of size at most M , we have $j_p \leq M$ for all p .

Note that B does not contain any negation, thus if \mathbf{u} did not satisfy some $X^{j_p} \psi_p$, we would have that \mathbf{u} satisfies $B[X^{j_1} \psi_1 \rightarrow y_1, \dots, X^{j_m} \perp \rightarrow y_p, \dots, X^{j_m} \psi_m \rightarrow y_m]$, while no \mathbf{v}_i satisfies it, since they do not satisfy $B[X^{j_1} \psi_1 \rightarrow y_1, \dots, X^{j_m} \psi_m \rightarrow y_m]$. As a result, $C[\psi_1 \rightarrow x_1, \dots, \perp \rightarrow x_p, \dots, \psi_m \rightarrow x_m]$ would be satisfied by \mathbf{u} but no \mathbf{v}_i , contradicting the minimality of φ . Hence \mathbf{u} satisfies all $X^{j_p} \psi_p$.

Let $1 \leq p \leq m$, let $1 \leq i \leq n$, we show that \mathbf{v}_i satisfies $\mathbf{X}^{j_p} \psi_p$. We distinguish two cases:

- $\psi_p = \mathbf{G}\psi'$. Since \mathbf{v}_i is a suffix of \mathbf{u} of length greater than M , $\mathbf{v}_i[j_p \dots]$ is not empty and is a suffix of $\mathbf{u}[j_p \dots]$. As the latter satisfies $\mathbf{G}\psi'$, so does the former. Hence $\mathbf{v}_i \models \mathbf{X}^{j_p} \psi_p$.

- $\psi_p = \mathbf{F}\psi'$. We have that $\mathbf{u}[j_p\dots] = (u'[j_p\dots]v'_1 \cdots v'_n)(u'v'_1 \cdots v'_n)^M$ satisfies ψ_p . By Fact 3, as ψ_p is of size at most $|\psi| \leq M$, $(u'[j_p\dots]v'_1 \cdots v'_n)(u^M v_1 \cdots u^M v_n)^{M-1}$ satisfies ψ_p as well. As the latter is a suffix of $\mathbf{v}_i[j_p\dots]$, $\mathbf{v}_i[j_p\dots]$ also satisfies the formula. Hence $\mathbf{v}_i \models \mathbf{X}^{j_p}\psi_p$.

Now consider the formula $\varphi' = C[\top \rightarrow x_1, \dots, \top \rightarrow x_m]$, which is equivalent to $B[\mathbf{X}^{j_1}\top \rightarrow y_1, \dots, \mathbf{X}^{j_m}\top \rightarrow y_m]$. As all \mathbf{v}_i satisfy all ψ_p but not φ , no \mathbf{v}_i satisfies φ' . Further, as \mathbf{u} satisfies φ and all ψ_p , it also satisfies φ' . This contradicts the minimality of φ , unless $m = 0$. As a result, φ does not contain any \mathbf{F} or \mathbf{G} . Thus the minimal $\mathbf{LTL}(Op)$ formula separating \mathbf{u} and the \mathbf{v}_i is the same as the minimal $\mathbf{LTL}(Op \cap \{\mathbf{X}, \wedge, \vee\})$ separating them.

By Fact 4, we have that this formula is also the minimal formula separating u' and the v'_i . Clearly any $\mathbf{LTL}(Op \cap \{\mathbf{X}, \wedge, \vee\})$ formula separating u and the v_i also separates those. By Lemma 5, there exists $\varphi' \in \mathbf{LTL}(\mathbf{X}, \wedge)$ with $|\varphi'| \leq |\varphi|$ separating u' and the $v_i u^{M-1}$, of the form $X^{i_1-1}(c_1 \wedge \mathbf{X}^{i_2-i_1}(\dots \wedge \mathbf{X}^{i_p-i_{p-1}}c_p)\dots)$ with $0 < i_1 < \dots < i_p$ and $c-1, \dots, c_p \in \Sigma$. As u' and the v'_i are equal after the first $|u|$ letters, by minimality of $|\varphi|$, we have $i_p < |u|$. As a result, by Fact 4, φ' separates u and the v_i .

We conclude that the minimal size of a formula of $\mathbf{LTL}(Op)$ separating \mathbf{u} and the \mathbf{v}_i is the same as the minimal size of a formula of $\mathbf{LTL}(\mathbf{X}, \wedge)$ separating u and the v_i . This completes the reduction. \square

Proposition 12. *For all $\{\mathbf{X}, \vee\} \subseteq Op \subseteq \{\mathbf{F}, \mathbf{G}, \mathbf{X}, \wedge, \vee\}$, there is a polynomial-time reduction from the $\mathbf{LTL}(\mathbf{X}, \vee)$ learning problem for a single negative word to the $\mathbf{LTL}(Op)$ one over the same alphabet.*

Proof. The proof is identical to the one of Proposition 11, with the roles of positive and negative words reversed and disjunctions and conjunctions reversed. \square

8. Almost all fragments without the next operator are hard to approximate

A study of $\mathbf{LTL}(\mathbf{F}, \wedge)$

As we will see, $\mathbf{LTL}(\mathbf{F}, \wedge)$ over an alphabet of size 2 is very weak. This degeneracy vanishes when considering alphabets of size at least 3. Instead of defining a normal form as we did for $\mathbf{LTL}(\mathbf{X}, \wedge)$ we characterise the expressive power of $\mathbf{LTL}(\mathbf{F}, \wedge)$ and construct for each property expressible in this logic a minimal formula.

Lemma 6. *For every formula $\varphi \in \mathbf{LTL}(\mathbf{F}, \wedge)$, either it is equivalent to false or there exists a finite set of non-repeating words w_1, \dots, w_p and $c \in \Sigma \cup \{\varepsilon\}$ such that for every word z ,*

$$z \models \varphi \text{ if and only if } \begin{cases} \text{for all } q \in [1, p], w_q \text{ is a subword of } z, \\ \text{and } z \text{ starts with } c. \end{cases}$$

Proof. We proceed by induction over φ .

- For the atomic formula $c \in \Sigma$, the property is satisfied using the empty set of words and c .

- If $\varphi = \mathbf{F}\varphi'$, by induction hypothesis we get w_1, \dots, w_p and c for φ' . We let $w'_i = cw_i$ if $w_i(1) \neq c$ and $w'_i = w_i$ otherwise, then $z \models \varphi$ if and only if for all $q \in [1, p]$, w'_q is a subword of z and z starts with ε (the latter condition is always satisfied).
- If $\varphi = \varphi_1 \wedge \varphi_2$, by induction hypothesis we get $w_1^1, \dots, w_{p_1}^1, c_1$ for φ_1 and $w_1^2, \dots, w_{p_2}^2, c_2$ for φ_2 . There are two cases. If c_1 and c_2 are non-empty and $c_1 \neq c_2$ then φ is equivalent to false. Otherwise, either both are non-empty and equal or at least one is ε , say c_2 . In both cases, $u \models \varphi$ if and only if for all $(e, q) \in (1, [1, p_1]) \cup (2, [1, p_2])$, w_q^e is a subword of u and u starts with c_1 .

□

Lemma 6 gives a characterisation of the properties expressible in $\mathbf{LTL}(\mathbf{F}, \wedge)$. It implies that over an alphabet of size 2 the fragment $\mathbf{LTL}(\mathbf{F}, \wedge)$ is very weak. Indeed, there are very few non-repeating words over the alphabet $\Sigma = \{a, b\}$: only prefixes of $abab\dots$ and $baba\dots$. This implies that formulas in $\mathbf{LTL}(\mathbf{F}, \wedge)$ over $\Sigma = \{a, b\}$ can only place lower bounds on the number of alternations between a and b (starting from a or from b) and check whether the word starts with a or b . In particular, the $\mathbf{LTL}(\mathbf{F}, \wedge)$ learning problem over this alphabet is (almost) trivial and thus not interesting. Hence we now assume that Σ has size at least 3.

We move back from semantics to syntax, and show how to construct minimal formulas. Let w_1, \dots, w_p a finite set of non-repeating words and $c \in \Sigma \cup \{\varepsilon\}$, we define a formula φ as follows.

The set of prefixes of w_1, \dots, w_p are organised in a forest (set of trees): a node is labelled by a prefix w of some w_1, \dots, w_p , and its children are the words wc which are prefixes of some w_1, \dots, w_p . The leaves are labelled by w_1, \dots, w_p . We interpret each tree t as a formula φ_t in $\mathbf{LTL}(\mathbf{F}, \wedge)$ as follows, in an inductive fashion: for $c \in \Sigma$, if t is labelled wa with subtrees t_1, \dots, t_q , then

$$\varphi_t = \mathbf{F}(c \wedge \bigwedge_i \varphi_{t_i}).$$

If $c = \varepsilon$, the formula associated to w_1, \dots, w_p and c is the conjunction of the formulas for each tree of the forest, and if $c \in \Sigma$, then the formula additionally has a conjunct c .

As an example, consider the set of words ab, ac, bab , and the letter a . The forest corresponding to ab, ac, bab contains two trees: one contains the nodes b, ba, bab , and the other one the nodes a, ab, ac . The two corresponding formulas are

$$\mathbf{F}(b \wedge \mathbf{F}(a \wedge \mathbf{F}b)) \quad ; \quad \mathbf{F}(a \wedge \mathbf{F}b \wedge \mathbf{F}c).$$

And the formula corresponding to the set of words ab, ac, bab , and the letter a is

$$a \wedge \mathbf{F}(b \wedge \mathbf{F}(a \wedge \mathbf{F}b)) \wedge \mathbf{F}(a \wedge \mathbf{F}b \wedge \mathbf{F}c).$$

Lemma 7. *For every non-repeating words w_1, \dots, w_p and $c \in \Sigma \cup \{\varepsilon\}$, the formula φ constructed above is minimal, meaning there are no smaller equivalent formulas.*

Applying the construction above to a single non-repeating word $w = c_1 \dots c_p$ we obtain what we call a “fattern” (pattern with an F):

$$F = \mathbf{F}(c_1 \wedge \mathbf{F}(\dots \wedge \mathbf{F}c_p) \dots),$$

We say that the non-repeating word w induces the fattern F above, and conversely that the fattern F induces the word w . The size of a fattern F is $3|w| - 1$. Adding the initial letter we obtain a grounded fattern $c \wedge F$, in that case the letter c is added at the beginning of w and the size is $3|w| - 2$.

Lemma 8. *Let $u_1, \dots, u_n, v_1, \dots, v_n$. If there exists $\varphi \in \mathbf{LTL}(\mathbf{F}, \wedge)$ separating u_1, \dots, u_n from v_1, \dots, v_n , then there exists a conjunction of at most n fatterns separating u_1, \dots, u_n from v_1, \dots, v_n .*

Proof. Thanks to Lemma 6, to the separating formula φ we can associate a finite set of non-repeating words w_1, \dots, w_p and $c \in \Sigma \cup \{\varepsilon\}$ such that for every word z ,

$$z \models \varphi \text{ if and only if } \begin{cases} \text{for all } q \in [1, p], w_q \text{ is a subword of } z, \\ \text{and } z \text{ starts with } c. \end{cases}$$

Let $j \in [1, n]$, since v_j does not satisfy φ either v_j does not start with c or for some $q \in [1, p]$ the word w_q is not a subword of u . For each $j \in [1, n]$ such that v_j starts with c , we pick one $q_j \in [1, p]$ for which w_{q_j} is not a subword of v_j , and consider the set $\{w_{q_j} : j \in [1, n]\}$ together with $c \in \Sigma \cup \{\varepsilon\}$. The formula induced by the construction above is a conjunction of at most n fatterns and it separates u_1, \dots, u_n from v_1, \dots, v_n . \square

Hardness result when the alphabet is part of the input

Theorem 8. *The $\mathbf{LTL}(\mathbf{F}, \wedge)$ learning problem is NP-hard when the alphabet is part of the input, and there are no $(1 - o(1)) \cdot \log(n)$ polynomial time approximation algorithms unless $\mathbf{P} = \mathbf{NP}$, even with a single positive word.*

The result follows from a reduction from the hitting set problem. For proving the correctness of the reduction we will need a normalisation lemma specialised to the case of a single positive word.

Lemma 9. *Let u, v_1, \dots, v_n . If there exists $\varphi \in \mathbf{LTL}(\mathbf{F}, \wedge)$ separating u from v_1, \dots, v_n , then there exists a fattern of size smaller than or equal to φ separating u from v_1, \dots, v_n .*

Proof. Thanks to Lemma 6, to the separating formula φ we can associate a finite set of non-repeating words w_1, \dots, w_p and $c \in \Sigma \cup \{\varepsilon\}$ such that for every word z ,

$$z \models \varphi \text{ if and only if } \begin{cases} \text{for all } q \in [1, p], w_q \text{ is a subword of } z, \\ \text{and } z \text{ starts with } c. \end{cases}$$

Since u satisfies φ , it starts with c and for all $q \in [1, p]$, w_q is a subword of u . For each $q \in [1, p]$ there exists φ_q mapping the positions of w_q to u . Let us write w for the word obtained by considering all positions mapped by φ_q for $q \in [1, p]$. By definition w is a subword of u , and for all $q \in [1, p]$ w_q is a subword of w . It follows that the fattern induced by w separates u from v_1, \dots, v_n . The size of w is at most the sum of the sizes of the w_q for $q \in [1, p]$, hence the fattern induced by w is smaller than the original formula φ . \square

We can now prove Theorem 8.

Proof. We construct a reduction from the hitting set problem. Let C_1, \dots, C_n subsets of $[1, m]$ and $k \in \mathbb{N}$. Let us consider the alphabet $[0, m]$, we define the word $u = 012 \dots m$. For each $j \in [1, n]$ we let $[1, m] \setminus C_j = \{a_{j,1} < \dots < a_{j,m_j}\}$, and define $v_j = 0a_{j,1} \dots a_{j,m_j}$.

We claim that there exists a hitting set of size at most k if and only if there exists a formula in $\mathbf{LTL}(\mathbf{F}, \wedge)$ of size at most $3k - 1$ separating u from v_1, \dots, v_n .

Let $H = \{c_1, \dots, c_k\}$ a hitting set of size k with $c_1 < c_2 < \dots < c_k$, we construct the (non-grounded) fattern induced by $w = c_1 \dots c_k$, it separates u from v_1, \dots, v_n and has size $3k - 1$.

Conversely, let φ a formula in $\mathbf{LTL}(\mathbf{F}, \wedge)$ of size $3k - 1$ separating u from v_1, \dots, v_n . Thanks to Lemma 9 we can assume that φ is a fattern, let $w = c_1 \dots c_k$ the non-repeating word it induces. Necessarily $c_1 < c_2 < \dots < c_k$. If φ is grounded then $c_1 = 0$, but then the (non-grounded) fattern induced by $c_2 \dots c_k$ is also separating, so we can assume that φ is not grounded. We let $H = \{c_1, \dots, c_k\}$, and argue that H is a hitting set. Indeed, H is a hitting set if and only if for every $j \in [1, n]$ we have $H \cap C_j \neq \emptyset$, which is equivalent to for every $j \in [1, n]$ we have $v_j \not\models \varphi$; indeed for $c_i \in H \cap C_j$ by definition c_i does not appear in v_j so $v_j \not\models \mathbf{F}c_i$. \square

Hardness result for fixed alphabet

The most technical result in this paper is the following theorem, which strengthens Theorem 8 by restricting to an alphabet of size 3.

Theorem 9. *For all $\{\mathbf{F}, \wedge\} \subseteq \text{Op} \subseteq \{\mathbf{F}, \mathbf{G}, \wedge, \vee, \neg\}$, the learning problem for $\mathbf{LTL}(\text{Op})$ is NP-hard even for an alphabet of size 3.*

Its main difficulty stems from the use of a fixed alphabet, combined with the absence of \mathbf{X} , which forbids us from pointing at a specific position. To circumvent this problem we will adapt the previous reductions from hitting set: we again use a single positive example u against several negative ones v_1, \dots, v_n . Each of those words is a sequence of factors $abab \dots$ (one for each subset) separated by a letter c . We introduce small differences between the factors in the v_i and in u . We ensure that detecting those differences is costly in terms of operators, and thus that a minimal formula has to “talk about” as few of them as possible. The main difficulty in the proof is showing that we do not have more efficient ways to separate the words. Much like in the previous hardness proofs, the minimal formula has to select a minimal amount of positions (here factors) in the words that are enough to distinguish u from the v_i . Those positions answer the hitting set problem. We temporarily allow the negation operator as it does not make the proof more difficult and allows us to obtain a dual version of Theorem 9 easily.

We consider the alphabet $\Sigma = \{a, b, c\}$. Let $S = \{1, \dots, m\}$, let $T_1, \dots, T_n \subseteq S$, let $k' \in \mathbb{N}$. We set $M = 3m + 2$, $u = ((ab)^{M+1}c)^m$ and for all i , $v_i = cw_{i,1}cw_{i,2}c \dots w_{i,m}c$ with

$$w_{i,j} = \begin{cases} (ab)^M & \text{if } j \in T_i, \\ (ab)^{M+1} & \text{otherwise.} \end{cases}$$

Let k be size of a minimal hitting set.

Claim 1. *There exists a separating formula in $LTL(\mathbf{F}, \wedge)$ of size at most $6kM + 9m + 2$.*

Proof. Let H be a hitting set of size k . We define for $i \in [1, m]$,

$$z_j = \begin{cases} (ab)^{M+1} & \text{if } i \in H, \\ ab & \text{otherwise.} \end{cases}$$

and $w = c z_1 c z_2 c \cdots z_m c$. Let us write $w = x_1 \cdots x_p$ where x_1, \dots, x_p are letters, and define the formula

$$\psi = \mathbf{F}(x_1 \wedge \mathbf{F}(x_2 \wedge \mathbf{F}(x_3 \wedge \mathbf{F}(\cdots \wedge \mathbf{F}x_p)))).$$

This formula has size $3p - 1 = 3(2kM + 3m + 1) - 1 \leq 6kM + 9m + 2$. We claim that φ is a separating formula. First, note that since all letters in w are non-repeating, ψ is satisfied by exactly the words which have w as a subword.

- Since w is a subword of u , we have that u satisfies w .
- Let $j \in [1, n]$. Since v_j and w contain the same number of c 's, for w to be a subword of v_j , z_i needs to be a subword of $w_{i,j}$ for all i , i.e., we need to have $y_{i,j} = (ab)^{M+1}$ for all $j \in H$. However, as H is a hitting set of $(C_j)_{1 \leq j \leq n}$, there exists $j \in H$ such that $i \in C_j$ and thus $w_{i,j} = (ab)^M$.

□

The rest of the proof is about proving the converse implication:

Proposition 13. *Let φ be an $LTL(\mathbf{F}, \mathbf{G}, \wedge, \vee, \neg)$ formula separating u and the v_i , then we have $k(6M - 3) \leq \|\varphi\|$, and thus, as $\|\varphi\| \leq |\varphi|$, this implies $k(6M - 3) \leq |\varphi|$.*

We define a *separation tree* of φ as a finite tree with each node x labelled by a subformula ψ^x of φ , a set of indices $I^x \subseteq \{1, \dots, n\}$, a family of numbers $(J_i^x)_{i \in I^x}$ with $J_i^x \in \{0, \dots, m\}$ and two families of words $(P_i^x)_{i \in I^x}, (V_i^x)_{i \in I^x}$ with $P_i^x \in \{\varepsilon, a, b, ab, ba\}$ and V_i^x a suffix of w_{i, J_i^x} for all i (with $w_{i,0} = \varepsilon$ for all i).

This tree and labelling must respect the following rules. Those rules should be understood as follows: imagine two players, one trying to prove that u satisfies the formula φ but the v_i do not and the other trying to prove the contrary. The game is defined by induction on the formula. For instance, if φ is of the form $\mathbf{F}\psi$, then the first player will choose a suffix of u on which she claims ψ to be satisfied, and the second player will do the same for each v_i . A separation tree describes a play of this game in which the second player follows a specific strategy: he tries to copy everything the first player does, and to take a suffix that is as similar as possible to the current one of the first player.

Here are the rules:

- If $\psi^x = \psi_1 \vee \psi_2$ or $\psi^x = \psi_1 \wedge \psi_2$ then x has two children x_1 and x_2 and $\psi^{x_1} = \psi_1$, $\psi^{x_2} = \psi_2$ and I^x is the disjoint union of I^{x_1} and I^{x_2} , and $(J_i^{x_1}, P_i^{x_1}, V_i^{x_1}) = (J_i^x, P_i^x, V_i^x)$ for all $i \in I^{x_1}$ and $(J_i^{x_2}, P_i^{x_2}, V_i^{x_2}) = (J_i^x, P_i^x, V_i^x)$ for all $i \in I^{x_2}$.
- If $\psi^x = \mathbf{G}\psi$ then x has one child y , $\psi^y = \psi$, $I^y = I^x$ and either $J_i^y = J_i^x$ and V_i^y is a suffix of V_i^x or $0 \leq J_i^y < J_i^x$ and V_i^y is a suffix of w_{i, J_i^y} . In both cases $P_i^y = \varepsilon$.

- If ψ_x is a letter then x is a leaf and all $P_i^x V_i^x c$ start with that letter but none of the $V_i^x c$ do.
- If ψ_x is the negation of a letter then x is a leaf and all $V_i^x c$ start with that letter but none of the $P_i^x V_i^x c$ do.
- If $\psi_x = \mathbf{F}\psi$ with ψ a boolean combination of formulas **all of the form $\mathbf{F}\psi'$ or $\mathbf{G}\psi'$** then x has one child y , $\psi^y = \psi$ and $I^y = I^x$. Furthermore, there exists (J, U) such that for all $i \in I^x$ either $J_i^x = J$ and U is a suffix of $P_i^x V_i^x$ or $J_i^x < J$ and U is a suffix of $(ab)^{M+1}$. For all i we have $J_i^y = J$.

We have V_i^y equal to the shortest of U and V_i^x if $J_i^x = J_i^y$ and to the shortest of U and w_{i, C_i^y} otherwise. In the first case U is a suffix of $P_i^x V_i^x$ and in the second it is a suffix of $abw_{i, J}$. Thus in all cases we can select $P_i^y \in \{\varepsilon, a, b, ab, ba\}$ such that $U = P_i^x V_i^x$.

- If $\psi_x = \mathbf{G}\psi$ with ψ a boolean combination of formulas **all of the form $\mathbf{F}\psi'$ or $\mathbf{G}\psi'$** then x has one child y , $\psi^y = \psi$ and $I^y = I^x$. Furthermore, there exists (J, U) such that for all $i \in I^x$ either $J_i^x = J$ and U is a suffix of $P_i^x V_i^x$ or $J_i^x < J$ and U is a suffix of $(ab)^{M+1}$. For all i we have $J_i^y = J$.

For the values of P_i^y and V_i^y we have several cases:

- if $J_i^y > J_i^x$ then $V_i^y = U$ and $P_i^y = \varepsilon$ if U is a suffix of $w_{i, J}$. Otherwise we have $w_{i, J} = (ab)^M$ and either $U = (ab)^{M+1}$, in which case $P_i^y = ab$ and $V_i^y = (ab)^M$, or $U = b(ab)^M$, and then $P_i^y = ba$ and $V_i^y = b(ab)^{M-1}$.
- if $J_i^y = J_i^x$ and U is a suffix of V_i^x then $V_i^y = U$ and $P_i^y = \varepsilon$.
- if $J_i^y = J_i^x$ and $U = P_i^x V_i^x$ then $P_i^y = P_i^x$ and $V_i^y = V_i^x$.
- if $J_i^y = J_i^x$ and $U = aV_i^x$ and $|V_i^x| > 1$ then $P_i^y = ab$. Similarly if $U = bV_i^x$ and $|V_i^x| > 0$ then $P_i^y = ba$. In both cases V_i^y is such that $U = P_i^y V_i^y$.
- if $J_i^y = J_i^x$ and $U = aV_i^x$ and $|V_i^x| \leq 1$ or if $U = bV_i^x$ and $|V_i^x| = 0$ then $i \notin I^y$.

This also defines I^y .

Lemma 10. *If φ is satisfied by u but none of the v_i then there exists a separation proof tree for φ with a root r such that $I^r = \{1, \dots, n\}$ and for all $i \in I^r$, $J_i^r = 0$ and $P_i^r = V_i^r = \varepsilon$.*

Proof. We prove the following statement:

Let $\psi \in \mathbf{LTL}(\mathbf{F}, \mathbf{G}, \wedge, \vee, \neg)$, let $I \subseteq \{1, \dots, n\}$, let $(J_i)_{i \in I}$ be a family of indices in $\{0, \dots, m\}$, and let $(P_i)_{i \in I}$ and $(V_i)_{i \in I}$ be families of words such that for all $i \in I$, we have $P_i \in \{\varepsilon, a, b, ab, ba\}$, V_i is a suffix of w_{i, J_i} and $P_i V_i$ is a suffix of $(ab)^{M+1}$. If ψ is satisfied by all $u_i = P_i V_i c ((ab)^{M+1} c)^{m-J_i}$ but none of the $v_i = V_i c w_{i, J_i+1} c \dots w_{i, m} c$, then there exists a separation tree with a root r labelled by $I^r = I$, $\psi^r = \psi$ and $(J_i, P_i, V_i)_{i \in I}$.

This implies the lemma.

We proceed by induction on ψ .

Case 1: If ψ is a letter then all u_i start with that letter but none of the v_i do, thus the same can be said of the P_iV_ic and V_ic .

Case 2: Similarly if ψ is the negation of a letter then all v_i start with that letter but none of the u_i do, thus the same can be said of the V_ic and P_iV_ic .

Case 3: If $\psi = \psi_1 \vee \psi_2$ then we set $I_1 = \{i \in I \mid u_i \models \psi_1\}$ and $I_2 = I \setminus I_1$. As all u_i satisfy ψ , all u_i with $i \in I_2$ satisfy ψ_2 . As no v_i satisfies ψ , they satisfy neither ψ_1 nor ψ_2 . By induction hypothesis there exists a separation tree with a root r_1 (resp. r_2) labelled by I_1 (resp. I_2), ψ_1 (resp. ψ_2) and $(J_i, P_i, V_i)_{i \in I_1}$ (resp. $(J_i, P_i, V_i)_{i \in I_2}$). The tree with a root r labelled by I , ψ and $(J_i, P_i, V_i)_{i \in I}$ and with those two subtrees as children is a separation tree. We can proceed similarly when $\psi^x = \psi_1 \wedge \psi_2$ by setting $I_1 = \{i \in I \mid v_i \not\models \psi_1\}$.

Case 4: If $\psi = \mathbf{G}\psi'$ then for all i there exists a suffix v'_i of v_i not respecting ψ' . We set $J'_i = m - C'_i$ with C'_i the number of c in v' and V'_i the largest prefix of v'_i without c . We also set $P'_i = \varepsilon$. As u_i satisfies $\mathbf{G}\psi'$ and as $u'_i = V_i c ((ab)^{M+1} c)^{m-J_i}$ is a suffix of u_i , u'_i satisfies ψ' . By induction hypothesis, there exists a separation tree t with a root labelled by I , ψ' and $(J'_i, P'_i, V'_i)_{i \in I}$. As a result, there exists a separation tree for ψ , obtained by taking a root labelled by I , ψ and $(J'_i, P'_i, V'_i)_{i \in I}$ and giving it one child whose subtree is t .

Case 5: If $\psi = \mathbf{F}\psi'$ then let u' be the shortest suffix of u satisfying ψ' . As all u_i are suffixes of u satisfying ψ , u' is a suffix of all u_i . Let $J' = m - C'$ with C' the number of c in u' and U' the largest prefix of u' without c . For all $i \in I$ we set $J'_i = J'$. We will define a set $I' \subseteq I$ along the way. For all $i \in I$, if $J' = J_i$ and U' is a suffix of V_i , or if $J' > J_i$ and U' is a suffix of $w_{i,J'}$, then we can set $V'_i = U'$ and $P'_i = \varepsilon$. Then $v'_i = V'_i c w_{i,J'+1} \cdots c w_{i,m} c$ is a suffix of v_i and thus does not satisfy ψ' (as v_i does not satisfy $\mathbf{F}\psi'$). Otherwise, we have to distinguish cases according to the shape of ψ' .

Case 5.1: Suppose ψ' is a boolean combination of formulas of the form $\mathbf{F}\psi''$ or $\mathbf{G}\psi''$.

If $J_i = J'$ then we set $V'_i = V_i$. In that case U' is a suffix of $P_i V_i$ and V_i is a suffix of U_i , hence there exists $P'_i \in \{\varepsilon, a, b, ab, ba\}$ such that $U' = P'_i V'_i$. As $v_i = v'_i$ does not satisfy $\mathbf{F}\psi'$, it does not satisfy ψ' either.

If $J_i > J'$ then we set $V'_i = w_{i,J'}$. In that case U' is a suffix of $(ab)^{M+1}$ but not of $w_{i,J'}$, hence there exists $P'_i \in \{\varepsilon, a, b, ab, ba\}$ such that $U' = P'_i V'_i$. As v_i does not satisfy $\mathbf{F}\psi'$, and as $v'_i = V'_i c w_{i,J'+1} \cdots c w_{i,m} c$ is a suffix of v_i , v'_i does not satisfy ψ' .

Case 5.2: Now suppose ψ' is a boolean combination of formulas of which at least one is a letter or its negation.

If $J_i = J'$ and U' starts with aba then we set $P'_i = ab$ and V'_i such that $U' = abV'_i$. As U' is a suffix of $P_i V_i$ and $|P_i| \leq 2$, V'_i is a suffix of V_i . As a result, $v'_i = V'_i c w_{i,J'+1} \cdots c w_{i,m} c$ is a suffix of v_i , hence it does not satisfy ψ' .

Similarly, if $J_i = J'$ and U' starts with bab then we set $P'_i = ba$ and V'_i such that $U' = baV'_i$. Again, as U' is a suffix of $P_i V_i$ and $|P_i| \leq 2$, V'_i is a suffix of V_i . As a result, $v'_i = V'_i c w_{i,J'+1} \cdots c w_{i,m} c$ is a suffix of v_i , hence it does not satisfy ψ' .

If $J_i = J'$ and $|U'| \leq 2$ then $i \notin I'$.

If $J_i < J'$ then as U' is not a suffix of $w_{i,J'}$ we must have $w_{i,J'} = (ab)^M$ and U' is either $(ab)^{M+1}$ or $b(ab)^M$. If $U' = (ab)^{M+1}$ then we set $P'_i = ab$ and $V'_i = (ab)^M$. As $v'_i = V'_i c w_{i,J'+1} \cdots c w_{i,m} c$ is a suffix of v_i , it does not satisfy ψ' .

Similarly if $U' = b(ab)^M$ then we set $P'_i = ba$ and $V'_i = b(ab)^{M-1}$. Since $v'_i = V'_i cw_{i,J'+1} \cdots cw_{i,m}c$ is a suffix of v_i , it does not satisfy ψ' . We can then apply the induction hypothesis to obtain a separation tree t whose root is labelled by I' , ψ' and $(J'_i, P'_i, V'_i)_{i \in I'}$. We obtain the result by taking a tree whose root is labelled by I , ψ and $(J_i, P_i, V_i)_{i \in I}$ and giving it one child whose subtree is t .

This concludes our induction. \square

Lemma 11. *If there exists a separation proof tree for φ then $\|\varphi\| \geq 6k(M-1)$.*

Proof. We make some key observations. Let x be a node labelled by some letter or its negation. Then:

- either x only has ancestors labelled by formulas of the form $\varphi_1 \vee \varphi_2$ or $\varphi_1 \wedge \varphi_2$, in which case $P_i^x V_i^x c$ and $V_i^x c$ both start with c .
- or x has an ancestor labelled by a formula of the form $\mathbf{F}\varphi$ or $\mathbf{G}\varphi$. Let y be its closest such ancestor, and z the only child of y . Then ψ^y is of the form $\mathbf{F}\psi^z$ or $\mathbf{G}\psi^z$ with ψ^z a boolean combination of formulas including ψ^x (which is a letter or its negation).

Hence for all $i \in I^z$ we have that $P_i^z V_i^z c$ and $V_i^z c$ start with the same letter. Moreover as all nodes from z to x are labelled by formulas of the form $\psi_1 \vee \psi_2$ or $\psi_1 \wedge \psi_2$, we have $I^x \subseteq I^z$ and $P_i^x V_i^x c = P_i^z V_i^z c$ and $V_i^x c = V_i^z c$ for all $i \in I^x$. As a result, by definition of a separation tree, we must have $I^x = \emptyset$.

Thus for all node x , if x is a leaf then $I^x = \emptyset$. Furthermore, if x has two children y and z then I^x is the disjoint union of I^y and I^z , and if x has a single child y then $I^y \subseteq I^x$. Those facts are direct consequences of the definition of separation tree.

The second key observation is that if y is a child of x and $J_i^x = J_i^y$ and $P_i^y \neq \varepsilon$ for some i then $P_i^x \neq \varepsilon$ and $|V_i^y| \geq |V_i^x| - 1$. If furthermore $|V_i^y| = |V_i^x| - 1$ then ψ^x is of the form $\mathbf{F}\psi^y$ with ψ^y a boolean combination of formulas including at least one letter or its negation.

The third observation is that if y is a child of x and $J_i^x < J_i^y$ and $P_i^y \neq \varepsilon$ for some i then $|V_i^y| \geq 2M - 1$ and $w_{i,J_i^y} = (ab)^M$.

Let $i \in \{1, \dots, n\}$. In light of the previous observations, there is a branch of the tree x_1, \dots, x_{k+1} (x_{j+1} being a child of x_j for all j) such that i is in I^{x_j} for all $j \leq k$ but not in $I^{x_{k+1}}$. As $i \in I^{x_k}$ but $i \notin I^{x_{k+1}}$, we have $J_i^{x_{k+1}} = J_i^{x_k}$ and $|V_i^{x_k}| \leq 1$ and $P_i^{x_k} \neq \varepsilon$. Let j be the minimal index such that $J_i^{x_j} = J_i^{x_{k+1}}$. As the sequence of $J_i^{x_\ell}$ is non-increasing, all $J_i^{x_\ell}$ are equal for $j \leq \ell \leq k+1$. Thanks to the previous remarks, we infer that $P_i^{x_\ell} \neq \varepsilon$ for all $j \leq \ell \leq k+1$. Therefore x_j cannot be the root, hence $j > 1$. As a consequence, by our third remark, we have $|V_i^{x_j}| \geq 2M - 1$ and $w_{i,J_i^{x_j}} = (ab)^M$. Furthermore $|V_i^{x_k}| \leq 1$. Our second remark allows us to conclude that there exist at least $2M - 2$ nodes x_ℓ such that ψ^{x_ℓ} is of the form $\mathbf{F}\psi^{x_{\ell+1}}$ with $\psi^{x_{\ell+1}}$ a boolean combination of formulas including at least one letter or its negation. For each such ℓ , $x_{\ell+1}$ and the leaf corresponding to that letter (or its negation) are all labelled by $J_i^{x_\ell}$. Hence for all i there exists J_i such that $w_{i,J_i} = (ab)^M$ and $J_i^x = J_i$ for at least $6M - 6$ distinct nodes in the separation tree. As $w_{i,J_i} = (ab)^M$, we have $i \in T_{J_i}$ for all i , hence the set of J_i is a solution to the set cover problem. As a

result, there are at least k distinct J_i , and thus there are at least $6k(M-1)$ nodes in the separation tree. As the size of the separation tree is exactly $|\varphi|$, we obtain the result. \square

We may now finish the proof of Theorem 9. Recall that we considered an input $S = \{1, \dots, m\}$, $T_1, \dots, T_n \subseteq S$ and $k' \in \mathbb{N}$ (the encoding is irrelevant). We set k to be the size of a minimal hitting set of $H \subseteq S$ hitting all T_i , and $M = 3m + 2$.

Let φ be a formula of minimal size satisfied by u but not by any v_i . By Proposition 13 and Claim 1, we have $6kM - 3k \leq |\varphi| \leq 6kM + 9m + 2$. We set $K = 6k'M + 9m + 2$ and show that $k' \geq k$ if and only if $|\varphi| \leq K$.

- Suppose $k' \geq k$, then $K \geq 6kM + 9m + 2 \geq |\varphi|$.
- Suppose $k' \leq k - 1$, then $K \leq 6(k-1)M + 9m + 2 = 18km - 9m + 12k - 10 \leq (6kM - 3k) - 9m + 3k - 10 < 6kM - 3k \leq |\varphi|$ as $k \leq m$.

As a result, the hitting set problem has a positive answer on instance S, T_1, \dots, T_n, k' if and only if so does the **LTL**(Op) learning problem on instance u, v_1, \dots, v_n, K . As the latter instance is constructible in polynomial time from the former, Theorem 9 is proven.

Dual hardness result

We now show hardness for fragments with operators **G** and \vee . As we allowed negation in the previous result, we can infer that one almost directly.

Theorem 10. *For all $\{\mathbf{G}, \vee\} \subseteq Op \subseteq \{\mathbf{F}, \mathbf{G}, \wedge, \vee, \neg\}$, the learning problem for **LTL**(Op) is NP-hard even for an alphabet of size 3.*

We take the same instance of the hitting problem, but this time we consider the **LTL**(Op) learning problem with the v_i as positive words and u as the only negative one.

Corollary 1. *Let φ be an **LTL**($\wedge, \vee, \neg, \mathbf{F}, \mathbf{G}$) formula separating the u_i and v , then we have: $k(6M - 3) \leq |\varphi|$.*

Proof. Let φ be such a formula, and let $\bar{\varphi}$ be its negation, with the negations pushed to the bottom. Formally, $\bar{\varphi}$ is defined by induction on φ :

- $\bar{a} = \neg a$
- $\overline{\varphi_1 \wedge \varphi_2} = \bar{\varphi}_1 \vee \bar{\varphi}_2$
- $\overline{\varphi_1 \vee \varphi_2} = \bar{\varphi}_1 \wedge \bar{\varphi}_2$
- $\overline{\mathbf{F}\varphi_1} = \mathbf{G}\bar{\varphi}_1$
- $\overline{\mathbf{G}\varphi_1} = \mathbf{F}\bar{\varphi}_1$

A clear induction shows that a word satisfies one if and only if it does not satisfy the other (note that this would not be true if we allowed the operator **X**). Hence $\bar{\varphi}$ is satisfied by u but not by the v_i , so by Theorem 9 we have $k(6M - 3) \leq |\bar{\varphi}| = |\varphi|$. \square

Lemma 12. *There exists a formula φ of $\mathbf{LTL}(\mathbf{G}, \vee)$ separating the u_i and v with $|\varphi| \leq 6kM + 11m + 4$.*

Proof. Let J be a cover of S of size k . We set, for all $1 \leq i \leq n$,

$$z_j = \begin{cases} (ab)^{M+1} & \text{if } j \in J, \\ ab & \text{otherwise.} \end{cases}$$

and we set $w = cz_1cz_2c \cdots z_m c$. We use the formula $\psi = \mathbf{G}(\bar{x}_1 \vee \mathbf{G}(\bar{x}_2 \vee \mathbf{G}(\bar{x}_3 \vee \mathbf{G}(\dots \vee \mathbf{G}(\bar{x}_p))))$, where x_1, \dots, x_p are letters such that $x_1 \cdots x_p = w$, $\bar{a} = b$, $\bar{b} = a$ and $\bar{c} = a \vee b$. This formula has size $3p - 1 + 2m + 2 = 3(2kM + 3m + 1) + 2m + 1 \leq 6kM + 11m + 4$.

Observe that ψ is satisfied by exactly the words which do not have any w' as a weak subword, with $w' \in (2^\Sigma)^*$ obtained by replacing each c by $\{c\}$, a by $\{a, c\}$ and b by $\{b, c\}$. As w is a weak subword of v , so is w' , hence v does not satisfy φ . Let $i \in S$, as u_i and w contain the same number of c , for u_i to not satisfy the formula, z_j needs to be a subword of $y_{i,j}$ for all j , i.e., we need to have $y_{i,j} = (ab)^{M+1}$ for all $j \in J$. However, as J is a cover of S , there exists $j \in J$ such that $i \in C_j$ and thus $y_{i,j} = (ab)^M$. Therefore all u_i satisfy the formula. \square

Let φ be a formula of minimal size satisfied by all v_i but not by u . By Proposition 1 and Lemma 12, we have $6kM - 3k \leq |\varphi| \leq 6kM + 11m + 4$.

We set $K = 6k'M + 11m + 4$ and show that $k' \geq k$ if and only if $|\varphi| \leq K$.

- Suppose $k' \geq k$, then $K \geq 6kM + 11m + 4 \geq |\varphi|$.
- Suppose $k' \leq k - 1$, then $K \leq 6(k - 1)M + 11m + 4 = 18km - 7m + 12k - 8 \leq (6kM - 3k) - 7m + 3k - 8 < 6kM - 3k \leq |\varphi|$ as $k \leq m$.

As a result, the hitting set problem has a positive answer on instance S, T_1, \dots, T_n, k' if and only if so does the $\mathbf{LTL}(Op)$ learning problem on instance u, v_1, \dots, v_n, K . As the latter instance is constructible in polynomial time from the former, Theorem 10 is proven.

9. Perspectives and open problems

In this paper, we showed NP-completeness of the \mathbf{LTL} learning problem for all fragments which do not include the until operator. The same holds adding until for non-constant size alphabets. Hence the main open question is the following:

Open problem 1. *Is the learning problem NP-complete for full \mathbf{LTL} with constant size alphabet?*

The negative results in this paper suggest looking for approximation algorithms. Beyond $\mathbf{LTL}(\mathbf{X}, \wedge)$, nothing is known in terms of upper and lower bounds on polynomial time approximation algorithms, leaving an open field of exciting research directions. In the same vein, one could wonder about the parameterized complexity of the \mathbf{LTL} learning problem, in particular when fixing the number of words. We leave this question open and hope to inspire further studies!

References

- Baharisangari, N., Gaglione, J., Neider, D., Topcu, U., & Xu, Z. (2021). Uncertainty-aware signal temporal logic inference. In Bloem, R., Dimitrova, R., Fan, C., & Sharygina, N. (Eds.), *International Conference on Software Verification, VSTTE*, Vol. 13124 of *Lecture Notes in Computer Science*, pp. 61–85. Springer.
- Bombara, G., Vasile, C. I., Penedo Alvarez, F., Yasuoka, H., & Belta, C. (2016). A Decision Tree Approach to Data Classification using Signal Temporal Logic. In *Hybrid Systems: Computation and Control, HSCC*.
- Camacho, A., & McIlraith, S. A. (2019). Learning interpretable models expressed in linear temporal logic. *International Conference on Automated Planning and Scheduling, ICAPS*, 29.
- Dinur, I., & Steurer, D. (2014). Analytical approach to parallel repetition. In *Symposium on Theory of Computing, STOC*, pp. 624–633.
- Ehlers, R., Gavran, I., & Neider, D. (2020). Learning properties in $LTL \cap ACTL$ from positive examples only. In *Formal Methods in Computer Aided Design, FMCAD*.
- Fijalkow, N., & Lagarde, G. (2021). The complexity of learning linear temporal formulas from examples. In Chandlee, J., Eyraud, R., Heinz, J., Jardine, A., & van Zaanen, M. (Eds.), *International Conference on Grammatical Inference, ICGI*, Vol. 153 of *Proceedings of Machine Learning Research*, pp. 237–250. PMLR.
- Gaglione, J., Neider, D., Roy, R., Topcu, U., & Xu, Z. (2022). Maxsat-based temporal logic inference from noisy data. *Innovations in Systems and Software Engineering*, 18(3), 427–442.
- Gold, E. M. (1978). Complexity of automaton identification from given data. *Information and Control*, 37(3), 302–320.
- Kim, J., Muise, C., Shah, A., Agarwal, S., & Shah, J. (2019). Bayesian inference of linear temporal logic specifications for contrastive explanations. In *International Joint Conference on Artificial Intelligence, IJCAI*.
- Lemieux, C., Park, D., & Beschastnikh, I. (2015). General LTL specification mining. In *International Conference on Automated Software Engineering, (ASE)*.
- Lutz, S., Neider, D., & Roy, R. (2023). Specification sketching for linear temporal logic. In André, É., & Sun, J. (Eds.), *International Symposium on Automated Technology for Verification and Analysis, ATVA*, Vol. 14216 of *Lecture Notes in Computer Science*, pp. 26–48. Springer.
- Neider, D., & Gavran, I. (2018). Learning linear temporal properties. In *Formal Methods in Computer Aided Design, FMCAD*, pp. 1–10.
- Pitt, L., & Warmuth, M. K. (1993). The minimum consistent DFA problem cannot be approximated within any polynomial. *Journal of the ACM*, 40(1), 95–142.
- Pnueli, A. (1977). The temporal logic of programs. In *Symposium on Foundations of Computer Science, SFCS*.

- Raha, R., Roy, R., Fijalkow, N., & Neider, D. (2022). Scalable anytime algorithms for learning fragments of linear temporal logic. In Fisman, D., & Rosu, G. (Eds.), *International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS*, Vol. 13243 of *Lecture Notes in Computer Science*, pp. 263–280. Springer.
- Raha, R., Roy, R., Fijalkow, N., Neider, D., & Pérez, G. A. (2024). Synthesizing efficiently monitorable formulas in metric temporal logic. In *International Conference on Verification, Model Checking, and Abstract Interpretation, VMCAI*.
- Roy, R., Fisman, D., & Neider, D. (2020). Learning interpretable models in the property specification language. In *International Joint Conference on Artificial Intelligence, IJCAI*.