

INFERRING SYMBOLIC AUTOMATA

DANA FISMAN ^a, HADAR FRENKEL ^b, AND SANDRA ZILLES ^c

^a Ben-Gurion University, Be'er Sheva, Israel

^b CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

^c University of Regina, Regina, Canada

ABSTRACT. We study the learnability of *symbolic finite state automata* (SFA), a model shown useful in many applications in software verification. The state-of-the-art literature on this topic follows the *query learning* paradigm, and so far all obtained results are positive. We provide a necessary condition for efficient learnability of SFAs in this paradigm, from which we obtain the first negative result.

The main focus of our work lies in the learnability of SFAs under the paradigm of *identification in the limit using polynomial time and data*, and its strengthening *efficient identifiability*, which are concerned with the existence of a systematic set of *characteristic samples* from which a learner can correctly infer the target language. We provide a necessary condition for identification of SFAs in the limit using polynomial time and data, and a sufficient condition for efficient learnability of SFAs. From these conditions we derive a positive and a negative result.

The performance of a learning algorithm is typically bounded as a function of the size of the representation of the target language. Since SFAs, in general, do not have a canonical form, and there are trade-offs between the complexity of the predicates on the transitions and the number of transitions, we start by defining size measures for SFAs. We revisit the complexity of procedures on SFAs and analyze them according to these measures, paying attention to the special forms of SFAs: *normalized SFAs* and *neat SFAs*, as well as to SFAs over a *monotonic* effective Boolean algebra.

This is an extended version of the paper with the same title published in CSL'22 [FFZ22].

1. INTRODUCTION

Symbolic finite state automata, SFAs for short, are an automata model in which transitions between states correspond to predicates over a domain of concrete alphabet letters. Their purpose is to cope with situations where the domain of concrete alphabet letters is large or infinite. As an example for automata over finite large alphabets consider automata over the alphabet 2^{AP} where AP is a set of atomic propositions; these are used in model checking [CGP01, BK08]. Another example, used in string sanitizer algorithms [HLM⁺11], are automata over predicates on the Unicode alphabet which consists of over a million symbols. An infinite alphabet is used for example in *event recording automata*, a determinizable class of timed automata [AFH99] in which an alphabet letter consists of both a symbol from a

finite alphabet, and a non-negative real number. Formally, the transition predicates in an SFA are defined with respect to an effective Boolean algebra as defined in section 2.

SFAs have proven useful in many applications [DVLM14, PGLM15, ASJ⁺16, HD17, SV17, MRA⁺17] and consequently have been studied as a theoretical model of automata. Many algorithms for natural questions over these automata already exist in the literature, in particular, Boolean operations, determinization, and emptiness [VdHT10]; minimization [DV16]; and language inclusion [KT14]. Recently the subject of learning automata in verification has also attracted attention, as it has been shown useful in many applications, see Vaandrager’s survey [Vaa17].

There already exists substantial literature on learning restricted forms of SFAs [GJL10, MM14, ASKK16, MM17, CDYS17], as well as general SFAs [DD17, AD18], and even non-deterministic residual SFAs [CHYS19]. For other types of automata over infinite alphabets, [HSM11] suggests learning abstractions, and [She19] presents a learning algorithm for deterministic variable automata. All these works consider the query learning paradigm, and provide extensions to Angluin’s \mathbf{L}^* algorithm for learning DFAs using membership and equivalence queries [Ang87a]. Unique to these works is the work [AD18] which studies the learnability of SFAs taking as a parameter the learnability of the underlying algebras, providing positive results regarding specific Boolean algebras.

One of our contributions is to demonstrate that these positive learnability results are far from trivial. In particular, we show that there are limitations to the power of membership and equivalence queries when it comes to learning SFAs. To do so, we provide a necessary condition for efficient learnability of SFAs in the query learning paradigm, from which we obtain a negative result regarding query learning of SFAs over the propositional algebra. This is, to the best of our knowledge, the first negative result on learning SFAs with membership and equivalence queries and thus gives useful insights into the limitations of the \mathbf{L}^* framework in this context.

The main focus of our work lies on the learning paradigm of *identification in the limit using polynomial time and data*. We are interested in providing sufficient or necessary conditions for a class of SFAs to be learnable under this paradigm. To this end, we show that the type of the algebra, in particular whether it is monotonic or not, largely influences the learnability of the class.

Learnability of a class of languages in a certain paradigm greatly depends on the representation chosen for the language. For instance, regular languages are efficiently learnable (both in the paradigm of identification in the limit using polynomial time and data, and in the query learning paradigm using membership and equivalence queries) when represented as DFAs but not when represented as NFAs. While we are interested in SFAs as the representations, there are various types of SFAs (with the same expressive power), and the learnability results for them may vary.

The literature on SFAs has mainly focused on a special type of SFA, termed *normalized*, in which there is at most one transition between every pair of states. This minimization of the number of transitions comes at the cost of obtaining more complex predicates. We consider, in addition to normalized SFAs, another special type of SFAs that we term *neat SFAs*, which by contrast, allows several transitions between the same pair of states, but restricts the predicates to be *basic*, as formally defined in subsection 2.1.

To get on the right track, we first take a global look at the complexity of the standard operations on SFAs, and how they vary according to the special form. We revisit the results in the literature and analyze them along the measures we find adequate for the size of an SFA:

the number of states (n), the number of transitions (m) and the size of the most complex predicate (l).¹ The results show that most procedures are more efficient on neat SFAs.

We then turn to study identification of SFAs in the limit using polynomial time and data. We provide a necessary condition a class of SFAs \mathbb{M} should meet in order to be identified in the limit using polynomial time and data, and a sufficient condition a class of SFAs \mathbb{M} should meet in order to be efficiently identifiable. These conditions are expressed in terms of the existence of certain efficiently computable functions, which we call $\text{Generalize}_{\mathbb{M}}$, $\text{Concretize}_{\mathbb{M}}$, and $\text{Decontaminate}_{\mathbb{M}}$. We then provide positive and negative results regarding the learnability of specific classes of SFAs in this paradigm. In particular, we show that the class of SFAs over *any* monotonic algebra is efficiently identifiable.

Comparison to the conference version Preliminary results of this work appear in [FFZ22]. This paper extends the results of [FFZ22] by adding a thorough discussion of the different SFA types and their effect on the complexity on different automata procedures; as well as a new theorem regarding efficient learnability, and additional examples for learning SFAs. In particular, sections 3, 4, and 5, are all new, as well as Theorem 7.3 and Examples 10.6 and 10.7.

Outline The rest of the paper is organized as follows. In section 2 we provide the necessary definitions on effective Boolean algebras and SFAs. Section 3 introduces the special forms of SFAs. In section 4 we discuss transformations between the special forms. Section 5 then reviews the complexity of standard automata procedures along the mentioned parameters.

We then turn to discuss the learnability of symbolic automata. Section 6 provides a short overview and definitions regarding learnability of SFAs. In section 7 we discuss the paradigm of learnability in the limit using polynomial time and data, and provide an overview of learning DFAs in this paradigm. Sections 8 and 9 present a necessary condition and a sufficient condition for the efficient learnability of SFAs, and sections 10 and 11 use these conditions to prove a positive result on the learnability of SFAs over monotonic algebras, and a negative result on the learnability of SFAs over the propositional algebra. Section 12 discusses query learning of SFAs and provides a negative result. We conclude in section 13 with a short discussion.

2. PRELIMINARIES

2.1. Effective Boolean Algebra. A *Boolean Algebra* \mathcal{A} can be represented as a tuple $(\mathbb{D}, \mathbb{P}, \llbracket \cdot \rrbracket, \perp, \top, \vee, \wedge, \neg)$ where \mathbb{D} is a set of domain elements; \mathbb{P} is a set of predicates closed under the Boolean connectives, where $\perp, \top \in \mathbb{P}$; the component $\llbracket \cdot \rrbracket : \mathbb{P} \rightarrow 2^{\mathbb{D}}$ is the so-called *semantics function*. It satisfies the following three requirements: (i) $\llbracket \perp \rrbracket = \emptyset$, (ii) $\llbracket \top \rrbracket = \mathbb{D}$, and (iii) for all $\varphi, \psi \in \mathbb{P}$, $\llbracket \varphi \vee \psi \rrbracket = \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket$, $\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$, and $\llbracket \neg \varphi \rrbracket = \mathbb{D} \setminus \llbracket \varphi \rrbracket$. A Boolean Algebra is *effective* if all the operations above, as well as satisfiability, are decidable. Henceforth, we implicitly assume Boolean algebras to be effective.

One way to define a Boolean algebra is by defining a set \mathbb{P}_0 of *atomic formulas* that includes \top and \perp and obtaining \mathbb{P} by closing \mathbb{P}_0 for conjunction, disjunction and negation. For a predicate $\psi \in \mathbb{P}$ we say that ψ is *atomic* if $\psi \in \mathbb{P}_0$. We say that ψ is *basic* if ψ is a conjunction of atomic formulas.

We now introduce two Boolean algebras that are discussed extensively in the paper.

¹Previous results have concentrated mainly on the number of states.

The Interval Algebra is the Boolean algebra in which the domain \mathbb{D} is the set $\mathbb{Z} \cup \{-\infty, \infty\}$ of integers augmented with two special symbols with their standard semantics, and the set of atomic formulas \mathbb{P}_0 consists of intervals of the form $[a, b)$ where $a, b \in \mathbb{D}$. The semantics associated with intervals is the natural one: $\llbracket [a, b) \rrbracket = \{z \in \mathbb{D} \mid a \leq z \text{ and } z < b\}$. If $a \geq b$ then $\llbracket [a, b) \rrbracket = \emptyset$ and we have that $[a, b)$ is semantically equivalent to \perp .

The Propositional Algebra is defined with respect to a set $AP = \{p_1, p_2, \dots, p_k\}$ of atomic propositions. The set of *atomic predicates* \mathbb{P}_0 consists of the atomic propositions and their negations as well as \top and \perp . The domain \mathbb{D} consists of all the possible valuations for these propositions, thus it is \mathbb{B}^k where $\mathbb{B} = \{0, 1\}$. The semantics of an atomic predicate p is given by $\llbracket p_i \rrbracket = \{v \in \mathbb{B}^k \mid v[i] = 1\}$, and similarly $\llbracket \neg p_i \rrbracket = \{v \in \mathbb{B}^k \mid v[i] = 0\}$.²

2.1.1. *Predicate Size.* In order to reason about the complexity of operations over the Boolean algebra (and later, the efficient learnability of SFAs using such Boolean algebras), we need some measure of the size of predicates. We assume the algebra is associated with a function $size^{\mathbb{P}} : \mathbb{P} \rightarrow \mathbb{N}$ returning for each predicate its size. If the algebra is defined via a set of atomic propositions, one can assume the existence of functions $size^{\mathbb{P}} : \mathbb{P}_0 \rightarrow \mathbb{N}$, $size_{\wedge}^{\mathbb{P}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$, $size_{\vee}^{\mathbb{P}} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$, $size_{\neg}^{\mathbb{P}} : \mathbb{N} \rightarrow \mathbb{N}$ according to which the size of predicates can be inductively computed. Note that the size is a property of the predicate, not the set of concrete elements it represents.

Example 2.1. For the interval algebra, we define the size of one interval to be 1, and the size of a general predicate as the size of its parse tree, where leaves are single intervals (whose size is 1). Thus for example $size^{\mathbb{P}}(\llbracket ([0, 50) \vee [100, 200) \wedge [20, 60) \rrbracket$ is 5 whereas the size of the semantically equivalent predicate $\llbracket [20, 50) \rrbracket$ is 1.


Similarly, for the propositional algebra we define the size of a predicate to be the size of its parse tree. Note that Boolean functions from \mathbb{B}^k to \mathbb{B} can be represented in other ways as well, e.g., using Binary Decision Diagrams (BDDs) [Bry86]. This would result in a different Boolean algebra (where predicates are BDDs) with a different size measure for predicates.

2.2. **Symbolic Automata.** A *symbolic finite automaton* (SFA) is a tuple $\mathcal{M} = (\mathcal{A}, Q, q_i, F, \Delta)$ where \mathcal{A} is a Boolean algebra, Q is a finite set of states, $q_i \in Q$ is the initial state, $F \subseteq Q$ is the set of final states, and $\Delta \subseteq Q \times \mathbb{P}_{\mathcal{A}} \times Q$ is a finite set of transitions, where $\mathbb{P}_{\mathcal{A}}$ is the set of predicates of \mathcal{A} .

We use the term *letters* for elements of \mathbb{D} where \mathbb{D} is the domain of \mathcal{A} , and the term *words* for elements of \mathbb{D}^* . A *run* of \mathcal{M} on a word $a_1 a_2 \dots a_n$ is a sequence of transitions $\langle q_0, \psi_1, q_1 \rangle \langle q_1, \psi_2, q_2 \rangle \dots \langle q_{n-1}, \psi_n, q_n \rangle$ satisfying that $a_i \in \llbracket \psi_i \rrbracket$, that $\langle q_i, \psi_{i+1}, q_{i+1} \rangle \in \Delta$ and that $q_0 = q_i$. Such a run is said to be *accepting* if $q_n \in F$. A word $w = a_1 a_2 \dots a_n$ is said to be *accepted* by \mathcal{M} if there exists an accepting run of \mathcal{M} on w . The set of words accepted by an SFA \mathcal{M} is denoted $\mathcal{L}(\mathcal{M})$. We use $\hat{\mathcal{L}}(\mathcal{M})$ for the set of labeled words $\hat{\mathcal{L}}(\mathcal{M}) = \{(w, 1) \mid w \in \mathcal{L}(\mathcal{M})\} \cup \{(w, 0) \mid w \notin \mathcal{L}(\mathcal{M})\}$.

An SFA is said to be *deterministic* if for every state $q \in Q$ and every letter $a \in \mathbb{D}$ we have that $|\{\langle q, \psi, q' \rangle \in \Delta \mid a \in \llbracket \psi \rrbracket\}| \leq 1$, namely from every state and every concrete letter there exists at most one transition. It is said to be *complete* if $|\{\langle q, \psi, q' \rangle \in \Delta \mid a \in \llbracket \psi \rrbracket\}| \geq 1$ for every $q \in Q$ and $a \in \mathbb{D}$, namely from every state and every concrete letter there exists at least one transition. It is not hard to see that, as is the case for finite automata (over concrete

²In this case a basic formula is a *monomial*.

FIGURE 1. The SFA \mathcal{M} over $\mathcal{A}_{\mathbb{N}}$

alphabets), non-determinism does not add expressive power but does add succinctness. When \mathcal{A} is deterministic we use $\Delta(q, w)$ to denote the state \mathcal{A} reaches on reading the word w from state q . If $\Delta(q, w) = q$ then w is termed an *access word to state q* . If w is the smallest access word according to lexicographic order we say that w is the *lex-access word to state q* .

Example 2.2. Consider the SFA \mathcal{M} given in Figure 1. It is defined over the algebra $\mathcal{A}_{\mathbb{N}}$ which is the interval algebra restricted to the domain $\mathbb{D} = \mathbb{N} \cup \{\infty\}$. The language of \mathcal{M} is the set of all words over \mathbb{D} of the form $w_1 \cdot d \cdot w_2$ where w_1 is some word over the domain \mathbb{D} , the letter d satisfies $0 \leq d < 100$ and all letters of the word w_2 are numbers smaller than 200. The lex-access word to state q_0 is ϵ , and 0 is the lex-access word to state q_1 .

3. TYPES OF SYMBOLIC AUTOMATA

Since the complexity of a learning algorithm for a class of languages \mathbb{L} using some representation \mathbb{R} is measured with respect to the size of the smallest representation $R \in \mathbb{R}$ for the unknown language $L \in \mathbb{L}$, we first need to agree how to measure the size of an SFA. Subsection 3.1 explains why the number of states is not a sufficient measure, and proposes an alternative using three parameters. Optimizing different parameters leads to different special forms which are discussed in subsection 3.2.

3.1. Size of an SFA. We note that there is a trade-off between the number of transitions and the complexity of the transition predicates. The size of an automaton (not a symbolic one) is typically measured by its number of states. This is since for DFAs, the size of the alphabet is assumed to be a given constant, and the rest of the parameters, in particular the transition relation, are at most quadratic in the number of states. In the case of SFAs the situation is different, as the size of the predicates labeling the transitions can vary greatly. In fact, if we measure the size of a predicate by the number of nodes in its parse DAG, then the size of a formula can grow unboundedly (and the same is true for other reasonable size measures for predicates). The size and structure of the predicates influence the complexity of their satisfiability check, and thus the complexity of the corresponding algorithms. On the other hand there might be a trade-off between the size of the transition predicates and the number of transitions; e.g., a predicate of the form $\psi_1 \vee \psi_2 \dots \vee \psi_k$ can be replaced by k transitions, each one labeled by ψ_i for $1 \leq i \leq k$.

Therefore, we measure the size of an SFA by three parameters: the number of states (n), the maximal out-degree of a state (m) and the largest size of a predicate (l).

In addition, in order to analyze the complexity of automata algorithms discussed in subsection 5.1 and subsection 5.2, for a class \mathbb{P} of predicates over a Boolean algebra \mathcal{A} , we use the complexity measure $\text{sat}^{\mathbb{P}}(l)$, which is the complexity of satisfiability check for a predicate of size l in \mathbb{P} . We also use $\text{sat}^{\mathbb{P}^0}(l)$ for the respective complexities when restricted to atomic predicates.

3.2. Special Form SFAs. We turn to define special types of SFAs, which affect the complexity of related procedures.

Neat and Normalized SFAs The literature defines an SFA as *normalized* if for every two states q and q' there exists at most one transition from q to q' . This definition prefers fewer transitions at the cost of potentially complicated predicates. By contrast, preferring simple transitions at the cost of increasing the number of transitions, leads to *neat* SFAs. We define an SFA to be *neat* if all transition predicates are basic predicates.

Feasibility The second distinction concerns the fact that an SFA can have transitions with unsatisfiable predicates. A symbolic automaton is said to be *feasible* if for every $\langle q, \psi, q' \rangle \in \Delta$ we have that $\llbracket \psi \rrbracket \neq \emptyset$. Feasibility is an orthogonal property to being neat or normalized.

Monotonicity The third distinction we make concerning the nature of a given SFA regards its underlying algebra. A Boolean algebra \mathcal{A} over domain \mathbb{D} is said to be *monotonic* if the following conditions hold.

- (1) There exists a total order $<$ on the elements of \mathbb{D} ; and
- (2) There exist two elements $d_{-\infty}$ and d_{∞} such that $d_{-\infty} \leq d$ and $d \leq d_{\infty}$ for all $d \in \mathbb{D}$; and
- (3) An atomic predicate $\psi \in \mathbb{P}_0$ can be associated with two concrete values a and b such that $\llbracket \psi \rrbracket = \{d \in \mathbb{D} : a \leq d < b\}$. Henceforth, we denote an atomic predicate ψ over a monotonic algebra as $\psi = [a, b)$ where $\llbracket \psi \rrbracket = \{d \in \mathbb{D} : a \leq d < b\}$. If $b \leq a$ then we have that $\llbracket \psi \rrbracket = \emptyset$ and thus the predicate is equivalent to \perp .

The interval algebra is clearly monotonic, as is the similar algebra obtained using \mathbb{R} (the real numbers) instead of \mathbb{Z} (the integers). On the other hand, the propositional algebra is clearly non-monotonic.

Example 3.1. The SFA \mathcal{M} from Example 2.2 (Figure 1) is defined over a monotonic algebra, and is neat, normalized, deterministic and complete.

4. TRANSFORMATIONS TO SPECIAL FORMS

We now address the task of transforming SFAs into their special forms as presented in section 3. We discuss transformations to the special forms *neat*, *normalized* and *feasible* automata, measured as suggested using $\langle n, m, l \rangle$ — the number of states, the maximal out-degree of a state, and the largest size of a predicate.

4.1. Neat Automata. Since each predicate in a neat SFA is a conjunction of atomic predicates, neat automata are intuitive, and the number of transitions in the SFA reflects the complexity of the different operations, as opposed to the situation with normalized SFAs. For the class \mathbb{P}_0 of basic formulas, $\text{sat}^{\mathbb{P}_0}(l)$ is usually more efficient than $\text{sat}^{\mathbb{P}}(l)$, and in particular is polynomial for the algebras we consider here. This is since, for a basic predicate φ that is a conjunction of l atomic predicates, satisfiability testing can be reduced to checking that there are no two atomic predicates that contradict each other. Since satisfiability checking directly affects the complexity of various algorithms discussed in subsection 5.1, neat SFAs allow for efficient automata operations, as we show in subsection 5.2.

4.1.1. *Transforming to Neat.* Given a general SFA \mathcal{M} of size $\langle n, m, l \rangle$, we can construct a neat SFA \mathcal{M}' of size $\langle n, m \cdot 2^l, l \rangle$, by transforming each transition predicate to a DNF formula, and turning each disjunct into an individual transition. The number of states, n , remains the same. However, the number of transitions can grow exponentially due to the transformation to DNF. In the worst case, the size of the most complex predicate can remain the same after the transformation, resulting in the same l parameter for both automata.

Note that there is not necessarily a unique minimal neat SFA. For instance, a predicate ψ over the propositional algebra with $AP = \{p_1, p_2, p_3\}$, satisfying $\llbracket \psi \rrbracket = \{[100], [101], [111]\}$ can be represented using the two basic transitions $(p_1 \wedge \neg p_2)$ and $(p_1 \wedge p_2 \wedge p_3)$; or alternatively using the two basic transitions $(p_1 \wedge p_3)$ and $(p_1 \wedge \neg p_2 \wedge \neg p_3)$, though it cannot be represented using one basic transition.³

Although in the general case, the transformation from normalized to neat SFAs is exponential, for monotonic algebras we have the following lemma, which follows directly from the definition of monotonic algebras and basic predicates.

Lemma 4.1. *Over a monotonic algebra, the conjunction of two atomic predicates is also an atomic predicate; inductively, any basic formula that does not contain negations, over a monotonic algebra, is an atomic predicate. In addition, the negation of an atomic predicate is a disjunction of at most 2 atomic predicates.*

Lemma 4.2. *Let \mathcal{M} be a normalized SFA over a monotonic algebra \mathcal{A}_{mon} . Then, transforming \mathcal{M} into a neat SFA \mathcal{M}' is linear in the size of \mathcal{M} .*

Since a DNF formula with m disjunctions is a natural representation of m basic transitions, Lemma 4.2 follows from the following property of monotonic algebras.

Lemma 4.3. *Let ψ be a general formula over a monotonic algebra \mathcal{A}_{mon} . Then, there exists an equivalent DNF formula ψ_d of size linear in $|\psi|$.*

Proof. First, we transform ψ into a Negation Normal Form formula ψ_{NNF} , pushing negations inside the formula. When transforming to NNF, the number of atomic predicates (possibly under negation) remains the same, and so is the number of conjunctions and disjunctions. Since, by Lemma 4.1, a negation of an atomic predicate over a monotonic algebra, namely a negation of an interval, results in at most two intervals, we get that $|\psi_{NNF}| \leq 2 \cdot |\psi|$. Note that ψ_{NNF} does not contain any negations, as they were applied to the intervals. We now transform ψ_{NNF} into a DNF formula ψ_d recursively, operating on sub-formulas of ψ_{NNF} , distributing conjunctions over disjunctions.

We inductively prove that $\llbracket \psi_d \rrbracket = \llbracket \psi_{NNF} \rrbracket$ and $|\psi_d| \leq |\psi_{NNF}|$. For the base case, if ψ_{NNF} is a single interval $[a, b)$, then $[a, b)$ is in DNF and we are done.

For the induction step, consider the following two cases.

- (1) Assume $\psi_{NNF} = \psi_1 \vee \psi_2$. By the induction hypothesis, there exists DNF formulas ψ_{1d} and ψ_{2d} such that $\llbracket \psi_{id} \rrbracket = \llbracket \psi_i \rrbracket$ and $|\psi_{id}| \leq |\psi_i|$ for $i = 1, 2$. Then, $\psi_d = \psi_{1d} \vee \psi_{2d}$ is equivalent to ψ_{NNF} and at most of the same size.
- (2) Assume $\psi_{NNF} = \psi_1 \wedge \psi_2$. Again, by the induction hypothesis, instead of $\psi_1 \wedge \psi_2$ we can consider $\psi_{1d} \wedge \psi_{2d}$ where ψ_{1d} and ψ_{2d} are in DNF. That is $\psi_{1d} = \bigvee_{i=1}^k [a_i, b_i)$ and $\psi_{2d} = \bigvee_{j=1}^l [c_j, d_j)$. Then, we have the following:

³This is related to the fact that there is no unique DNF formula – a neat automaton “breaks” the DNF to transitions according to the disjunctions in the formula. If there is no unique DNF formula, then there is no unique neat SFA.

$$\psi_{1d} \wedge \psi_{2d} = \left(\bigvee_{i=1}^k [a_i, b_i] \right) \wedge \left(\bigvee_{j=1}^l [c_j, d_j] \right) = \bigvee_{i=1}^k \bigvee_{j=1}^l ([a_i, b_i] \wedge [c_j, d_j])$$

From properties of intervals, each conjunction $[a_i, b_i] \wedge [c_j, d_j]$ is of the form $[\max\{a_i, c_j\}, \min\{b_i, d_j\}]$. The intervals in $\{[a_i, b_i] : 1 \leq i \leq k\}$ do not intersect (otherwise it would have resulted in a longer single interval), and the same for $\{[c_j, d_j] : 1 \leq j \leq l\}$. Thus, every element a_i or c_j can define at most one interval of the form $[\max\{a_i, c_j\}, \min\{b_i, d_j\}]$. That is, the DNF formula $\psi_d = \bigvee_{i=1}^k \bigvee_{j=1}^l ([a_i, b_i] \wedge [c_j, d_j])$ contains at most $k + l$ intervals, as the others are empty intervals. Since the size of the original ψ_{NNF} is $k + l$, we have that $|\psi_d| \leq |\psi_{NNF}|$.

To conclude, since ψ_{NNF} is linear in the size of ψ , and the size of ψ_d is at most the size of ψ_{NNF} , we have that the translation of ψ into the DNF formula ψ_d is linear. \square

4.2. Normalized Automata. Neat automata stand in contrast to normalized ones. In a normalized SFA, there is at most one transition between every pair of states, which allows for a succinct formulation of the condition to transit from one state to another. On the other hand, this makes the predicates on the transitions structurally more complicated. Given a general SFA \mathcal{M} with parameters $\langle n, m, l \rangle$, we can easily construct a normalized SFA \mathcal{M}' as follows. For every pair of states q and q' , construct a single edge labeled with the predicate $\bigvee_{\langle q, \varphi, q' \rangle \in \delta} \varphi$. Then, \mathcal{M}' has size $\langle n, \min(n^2, m), \text{size}_{\bigvee_m}^{\mathbb{P}}(l) \rangle$, where we use $\text{size}_{\bigvee_m}^{\mathbb{P}}(l)$ to denote the size of m disjunctions of predicates of size at most l .

Note that there is no unique minimal normalized automaton either, since in general a Boolean formula can have two semantically equivalent, yet syntactically different expressions in the underlying representation system, e.g., two distinct BDDs can represent the same formula. However, in subsection 5.2 we show that over monotonic algebras there is a canonical minimal normalized SFA.

The complexity of $\text{sat}^{\mathbb{P}}(l)$ for general formulas (corresponding to normalized SFAs) is usually exponentially higher than for basic predicates (and thus for neat SFAs). In addition, as we show above, generating a normalized automaton is an easy operation. This motivates working with neat automata, and generating normalized automata as a last step, if desired (e.g., for presenting a graphical depiction of the automaton).

4.3. Feasible Automata. The motivation for feasible automata is clear; if the automaton contains unsatisfiable transitions, then its size is larger than necessary, and the redundancy of transitions makes it less interpretable. Thus, infeasible SFAs add complexity both algorithmically and for the user, as they are more difficult to understand. In order to generate a feasible SFA from a given SFA \mathcal{M} , we need to traverse the transitions of \mathcal{M} and test the satisfiability of each transition. The parameters $\langle n, m, l \rangle$ of the SFA remain the same since there is no change in the set of states, and there might be no change in transitions as well (if they are all satisfiable).

In the following, we usually assume that the automata are feasible, and when applying algorithms, we require the output to be feasible as well.

Operation	$\langle n, m, l \rangle$
product construction $\mathcal{M}_1, \mathcal{M}_2$	$\langle n_1 \times n_2, m_1 \times m_2, size_{\wedge}^{\mathbb{P}}(l_1, l_2) \rangle$
complementation of deterministic \mathcal{M}_1 ⁴	$\langle n_1 + 1, m_1 + 1, size_{\neg}^{\mathbb{P}}(size_{\vee m_1}^{\mathbb{P}}(l_1)) \rangle$
determinization of \mathcal{M}_1	$\langle 2^{n_1}, 2^{m_1}, size_{\wedge n_1 \times m_1}^{\mathbb{P}}(l_1) \rangle$ ⁵
minimization of \mathcal{M}_1	$\langle n_1, m_1, size_{\wedge m_1}^{\mathbb{P}}(l_1) \rangle$

TABLE 1. Analysis of standard automata procedures on SFAs.

Decision Procedures	Time Complexity
emptiness	linear in n, m
emptiness + feasibility	$n \times m \times sat^{\mathbb{P}}(l)$
membership of $\gamma_1 \cdots \gamma_t \in \mathbb{D}^*$	$\sum_{i=1}^t sat^{\mathbb{P}}(size_{\wedge}^{\mathbb{P}}(l, \psi_{\gamma_i}))$ ⁶
inclusion $\mathcal{M}_1 \subseteq \mathcal{M}_2$	$(n_1 \times n_2) \times (m_1 \times m_2) \times sat^{\mathbb{P}}(size_{\wedge}^{\mathbb{P}}(l_1, l_2))$

TABLE 2. Analysis of time complexity of decision procedures for SFAs.

5. COMPLEXITY OF STANDARD AUTOMATA PROCEDURES ON SFAS

In this section we analyze the complexity of automata procedures on SFAs, in terms of their effect on the parameters $\langle n, m, l \rangle$. We start in subsection 5.1 with examining general SFAs, and then in subsection 5.2 discuss the effects on special SFAs.

5.1. Complexity of Automata Procedures for General SFAs. We turn to discuss Boolean operations, determinization and minimization, and decision procedures (such as emptiness and equivalence) for the different types of SFAs. For intersection and union, the product construction of SFAs was studied in [VdHT10, HV11]. There, the authors assume a normalized SFAs as an input, and do not delve into the effect of the construction on the number of transitions and the complexity of the resulting predicates. Determinization of SFAs was studied in [VdHT10], and [DV14] study minimization of SFAs, assuming the given SFA is normalized.

Table 1 shows the sizes of the SFAs resulting from the mentioned operations, in terms of $\langle n, m, l \rangle$. The analysis applies to all types of SFAs, not just normalized ones. The time complexity for each operation is given in terms of the parameters $\langle n, m, l \rangle$ and the complexity of feasibility tests for the resulting SFA, as discussed in subsection 4.3. Table 2 summarizes the time complexity of decision procedures for SFAs: emptiness, inclusion, and membership. Again, the analysis applies to all types of SFAs. We note that in many applications of learning in verification, the challenging part is implementing the teacher (e.g., in [PGB⁺08, CKKS20, FGPS20, FGPS22]). In such cases the complexity of membership and equivalence queries as well as standard automata operations plays a major role.

In both tables we consider two SFAs \mathcal{M}_1 and \mathcal{M}_2 with parameters $\langle n_i, m_i, l_i \rangle$ for $i = 1, 2$, over algebra \mathcal{A} with predicates \mathbb{P} . We use $size_{\wedge m}^{\mathbb{P}}(l)$ for an upper bound on the size of m conjunctions of predicates of size at most l . All SFAs are assumed to be deterministic, except of course for the input for determinization.

⁴For complementation, no feasibility check is needed, since we assume a feasible input.

⁵To determinize transitions, conjunction may be applied $n_1 \times m_1$ times, according to the number of states that correspond to a new deterministic state.

⁶Where ψ_{γ_i} is a predicate describing γ_i .

We now briefly describe the algorithms we analyze in both tables.

Product Construction [VdHT10, HV11] The product construction for SFAs is similar to the product of DFAs — the set of states is the product of the states of \mathcal{M}_1 and \mathcal{M}_2 ; and a transition is a synchronization of transitions of \mathcal{M}_1 and \mathcal{M}_2 . That is, a transition from $\langle q_1, q_2 \rangle$ to $\langle p_1, p_2 \rangle$ can be made while reading a concrete letter γ , iff $\langle q_1, \psi_1, p_1 \rangle \in \Delta_1$ and $\langle q_2, \psi_2, p_2 \rangle \in \Delta_2$ and γ satisfies both ψ_1 and ψ_2 . Therefore, the predicates labeling transitions in the product construction are conjunctions of predicates from the two SFAs \mathcal{M}_1 and \mathcal{M}_2 .

Complementation In order to complement a deterministic SFA \mathcal{M}_1 , we first need to make \mathcal{M}_1 complete. In order to do so, we add one state which is a non-accepting sink, and from each state we add at most one transition which is the negation of all other transitions from that state. If \mathcal{M}_1 is complete, then complementation simply switches accepting and non-accepting states, resulting in the same parameters $\langle n_1, m_1, l_1 \rangle$.

Determinization [VdHT10] In order to make an SFA deterministic, the algorithm of [VdHT10] uses the subset construction for DFAs, resulting in an exponential blowup in the number of states. However, in the case of SFAs this is not enough, and the predicates require special care. Let $P = \{q_1, \dots, q_t\}$ be a state in the deterministic SFA, where q_1, \dots, q_t are states of the original SFA \mathcal{M}_1 , and let ψ_1, \dots, ψ_t be some predicates labelling outgoing transitions from q_1, \dots, q_t , correspondingly. Then, in order to determinize transitions, the algorithm of [VdHT10] computes the conjunction $\bigwedge_{i=1}^t \psi_i$, which labels a single transition from the state P .

Minimization [DV14] Given a deterministic SFA \mathcal{M}_1 , the output of minimization is an equivalent deterministic SFA with a minimal number of states. When constructing such an SFA, the number of states and transitions cannot grow. However, as in determinization, if two states of \mathcal{M}_1 are replaced with one state, then outgoing transitions might overlap, resulting in a non-deterministic SFA. D’Antoni and Veanes [DV14] suggest several algorithms to cope with this difficulty. One of their approaches is to compute *minterms*, which are the smallest conjunctions of outgoing transitions. Minterms then do not intersect, and thus the output is deterministic. Their other approaches avoid computing minterms, but are able to achieve the same goal.

Emptiness If we assume a feasible SFA \mathcal{M} as an input, then in order to check for emptiness we need to find an accepting state which is reachable from the initial state (as in DFAs). If we do not assume a feasible input, we need to test the satisfiability of each transition, thus the complexity depends on the complexity measure $\text{sat}^{\text{P}}(l)$.

Membership Similarly to emptiness, in order to check if a concrete word $\gamma_1 \dots \gamma_n$ is in $\mathcal{L}(\mathcal{M})$, we need not only check if it reaches an accepting state but also locally consider the satisfiability of each transition. In the case of membership, we need to check whether the letter γ_i satisfies the predicate on the corresponding transition.

Inclusion Deciding inclusion amounts to checking emptiness and feasibility of $\mathcal{M}_1 \cap \overline{\mathcal{M}_2}$. We assume here that both \mathcal{M}_1 and \mathcal{M}_2 are deterministic and complete.

5.2. Complexity of Automata Procedures for Special SFAs. We now discuss the advantages of neat SFAs and of monotonic algebras, in the context of the algorithms presented in the tables, and show that, in general, they are more efficient to handle compared to other SFA types.

5.2.1. *Neat SFAs.* As can be observed from Table 2, almost all decision procedures regarding SFAs depend on $\text{sat}^{\mathbb{P}}(l)$. For neat SFAs it is more precise to say that they depend on $\text{sat}^{\mathbb{P}^0}(l)$, namely on the satisfiability of atomic predicates rather than arbitrary predicates. Since $\text{sat}^{\mathbb{P}^0}(l)$ is usually less costly than $\text{sat}^{\mathbb{P}}(l)$, most decision procedures are more efficient on neat automata. Here, we claim that applying automata algorithms on neat SFAs preserves their neatness, thus suggesting that neat SFAs may be preferable in many applications.

Lemma 5.1. *Let \mathcal{M}_1 and \mathcal{M}_2 be neat SFAs. Then the algorithms for their product construction, complementation, determinization and minimization discussed in subsection 5.1 result in a neat SFA.*

Proof. Observing the procedures for product construction [VdHT10, HV11], determinization [VdHT10] and minimization [DV14] constructions, one can see that they use only conjunctions in order to construct the predicates on the output SFAs. Thus, if the predicates on the input SFAs are basic, then so are the output predicates. \square

5.2.2. *Monotonic Algebras.* We now consider the class $\mathbb{M}_{\mathcal{A}_{\text{mon}}}$ of SFAs over a monotonic algebra \mathcal{A}_{mon} with predicates \mathbb{P} . We first discuss $\text{size}_{\wedge}^{\mathbb{P}}(l_1, l_2)$ and $\text{sat}^{\mathbb{P}}(l)$, as they are essential measures in automata operations. Then we show that for \mathcal{M}_1 and \mathcal{M}_2 in the class $\mathbb{M}_{\mathcal{A}_{\text{mon}}}$, the product construction is linear in the number of transitions, adding to the efficiency of SFAs over monotonic algebras.

Lemma 5.2. *Let ψ_1 and ψ_2 be formulas over a monotonic algebra \mathcal{A}_{mon} . Then $\text{size}_{\wedge}^{\mathbb{P}}(|\psi_1|, |\psi_2|)$ is linear in $|\psi_1| + |\psi_2|$ and $\text{sat}^{\mathbb{P}}(|\psi_1|)$ is linear in $|\psi_1|$.*

Proof. Transforming to DNF is linear, as follows from Lemma 4.3. There, we show that the conjunction of two DNF formulas of sizes k and l has size $k + l$, which implies that the conjunction of general formulas has linear size. In addition, $\text{sat}^{\mathbb{P}}(l)$ is trivial for a single interval, and following Lemma 4.3, is linear for general formulas. For an interval $[a, b)$, satisfiability checking amounts to the question “is $a < b$?”. \square

Lemma 5.3. *Let \mathcal{M}_1 and \mathcal{M}_2 be deterministic SFAs over a monotonic algebra \mathcal{A}_{mon} . Then the out-degree of their product SFA \mathcal{M} is at most $m = 2 \cdot (m_1 + m_2)$.*

Proof. From Lemma 4.2 and Lemma 4.3, we can construct neat SFAs \mathcal{M}'_1 and \mathcal{M}'_2 of sizes $\langle n_i, 2m_i, l_i \rangle$ for $i \in \{1, 2\}$, that have the same languages as \mathcal{M}_1 and \mathcal{M}_2 , respectively. Similarly to the proof of Lemma 4.3, each transition $\langle \langle q_1, q_2 \rangle, [a, b) \wedge [c, d), \langle p_1, p_2 \rangle \rangle$ in the product SFA results in a predicate $[\max\{a, c\}, \min\{b, d\})$. Then, for $q_1 \in Q_1$, every minimal element in the set of q_1 's outgoing transitions can define at most one transition in \mathcal{M} , and the same for a state $q_2 \in Q_2$, and so the number of transitions from $\langle q_1, q_2 \rangle$ is at most $m_1 + m_2$, as required. \square

Lemma 5.4. *Let \mathcal{M} be a neat SFA over a monotonic algebra. Then, transforming \mathcal{M} into a complete SFA \mathcal{M}' is polynomial in the size of \mathcal{M} .*

Proof. In order to complete \mathcal{M} , we add a non-accepting sink r in case it does not already exist, and at most $m + 1$ transitions from each state q to r , when m is the out-degree of the SFA: Let $[a, b)$ and $[c, d)$ be two predicates labeling outgoing transitions of q , where c is the minimal left end-point of a predicate such that $b < c$. Then, in order to complete \mathcal{M} , we need to add a transition to the sink, labeled by the predicate $[b, c)$. In addition, for the predicate $[a, b)$ where there is no $c > b$ that defines another predicate, if $b \neq d_{\infty}$ we add

$[b, d_\infty)$, and similarly we add $[d_{-\infty}, a)$. Then, for each state we add at most $m + 1$ new transitions, resulting in at most $|Q| \times (m + 1)$ new transitions. \square

Definition 5.5. For predicates over a monotonic algebra, we define a *canonical representation* of a predicate ψ as the simplified DNF formula which is the disjunction of all maximal disjoint intervals satisfying ψ .

Note that every predicate ψ over a monotonic algebra defines a unique partition of the domain into maximal disjoint intervals. This unique partition corresponds to a simplified DNF formula, which is exactly the canonical representation of ψ .

Example 5.6. The canonical representation of $\psi = [0, 100) \wedge ([50, 150) \vee [20, 40))$ is $[20, 40) \vee [50, 100)$.

Lemma 5.7. *Let \mathcal{M} be an SFA over a monotonic algebra. Then:*

- (1) *There is a unique minimal-state neat SFA \mathcal{M}' such that $\mathcal{L}(\mathcal{M}) = \mathcal{L}(\mathcal{M}')$.*
- (2) *There is a canonical minimal-state normalized SFA \mathcal{M}'' such that $\mathcal{L}(\mathcal{M}) = \mathcal{L}(\mathcal{M}'')$.*

Proof. For a language $\mathcal{L} = \mathcal{L}(\mathcal{M})$ for some SFA \mathcal{M} , the minimal number of states in an SFA corresponds, similarly to DFAs, to the number of equivalence classes in the equivalence relation N defined by $(u, v) \in N \iff \forall z \in \mathbb{D}^* : (uz \in \mathcal{L} \iff vz \in \mathcal{L})$ [Myh57, Ner58]. Indeed if $(u, v) \in N$ then there is no reason that reading them (from the initial state) should end up in different states, and if $(u, v) \notin N$ then reading them (from the initial state) must lead to different states.

As for transitions, we have the following.

- (1) Let ψ be a general predicate labeling a transition in \mathcal{M} . Then ψ defines a unique partition of the domain into maximal disjoint intervals, which are exactly the transitions in a neat SFA. Then, the minimal state neat SFA is unique, where its transitions correspond exactly to these maximal disjoint intervals.
- (2) For normalized transitions, we can use Lemma 4.3 to transform a general predicate labeling a transition to a DNF predicate in linear time. A DNF predicate over a monotonic algebra is in-fact a disjunction of disjoint intervals, where the construction of Lemma 4.3 obtains the maximal disjoint intervals. Then, to obtain a canonical representation, we order these intervals by order of their minimal elements. \square

6. LEARNING SFAS

We turn to discuss the learnability of symbolic automata. In grammatical inference, loosely speaking, we are interested in learning a class of languages \mathbb{L} over an alphabet Σ , from examples which are words over Σ . Examples for classes of languages can be the set of regular languages, the set of context-free languages, etc. A learning algorithm, aka a *learner*, is expected to output some concise representation of the language from a class of representations \mathbb{R} for the class \mathbb{L} . For instance, in learning the class \mathbb{L}_{reg} of regular languages one might consider the class \mathbb{R}_{DFA} of DFAs, or the class \mathbb{R}_{LIN} of right linear grammars, since both are capable of expressing all regular languages.⁷ We often say that a class of representations \mathbb{R} is learnable (or not) when we mean that class of languages \mathbb{L} is learnable (or not) via the class of representations \mathbb{R} . The complexity of learning an unknown language

⁷The class of regular languages was shown learnable via various representations including DFAs [Ang87a], NFAs [BHKLO9], and AFAs (alternating finite automata) [AEF15].

$L \in \mathbb{L}$ via \mathbb{R} is typically measured with respect to the size of the smallest representation $R_L \in \mathbb{R}$ for L . For instance, when learning \mathbb{L}_{reg} via \mathbb{R}_{DFA} a learner is expected to output a DFA for an unknown language in time that is polynomial in the number of states of the minimal DFA for L .

In our setting we are interested in learning regular languages using as a representation a class of SFAs over a certain algebra. To measure complexity we must agree on how to measure the size of an SFA. Thus, as discussed in subsection 3.1 we represent the size of an SFA using the parameters $\langle n, m, l \rangle$ of the number of states of the SFA, the number of transitions, and the size of the largest predicate. Another important factor regarding size and canonical forms of SFAs, is the underlying algebra, specifically, whether it is monotonic or not.

Learning Paradigms The exact definition regarding learnability of a class depends on the *learning paradigm*. In this work we consider two widely studied paradigms: *identification in the limit using polynomial time and data* and *learning with membership and equivalence queries*. Their definitions are provided in sections 7 and 12, respectively. Note that in general, a positive or negative result in one paradigm, does not imply the same result in another paradigm. We discuss this further in section 13.

Basic SFAs To provide results regarding the learnability of SFAs, we study classes of SFAs that contain all basic SFAs, defined as follows.

Definition 6.1. An SFA \mathcal{M} over a Boolean Algebra \mathcal{A} with a set of predicates \mathbb{P} is termed *basic* if it is of the form $\mathcal{M}_\varphi = (\mathcal{A}, \{q_l, q_{ac}, q_{rj}\}, q_l, \{q_{ac}\}, \Delta)$ where $\varphi \in \mathbb{P}$ and $\Delta = \{\langle q_l, \varphi, q_{ac} \rangle, \langle q_l, \neg\varphi, q_{rj} \rangle, \langle q_{rj}, \top, q_{rj} \rangle, \langle q_{ac}, \top, q_{rj} \rangle\}$. Note that \mathcal{M}_φ accepts only words of length one consisting of a concrete letter satisfying φ , and it is minimal among all complete deterministic SFAs accepting this language (minimal in both number of states and number of transitions).

In the sequel, our results are regarding classes of SFAs that contain all basic SFAs \mathcal{M}_φ for all $\varphi \in \mathbb{P}$.

7. EFFICIENT IDENTIFIABILITY

While in the better-known setting of *active learning* (namely, query learning with MQs and EQs) the learner can select any word and query about its membership in the unknown language, in *passive learning* the learner is given a set of words, and for each word w in the set, a label b_w indicating whether w is in the unknown language or not. Formally, a *sample* for a language L is a finite set \mathcal{S} consisting of labeled examples, that is, pairs of the form $\langle w, b_w \rangle$ where w is a word and $b_w \in \{0, 1\}$ is its label, satisfying that $b_w = 1$ if and only if $w \in L$. The words that are labeled 1 are termed *positive* words, and those that are labeled 0 are termed *negative* words. Note that if L is recognized by an automaton \mathcal{M} , we have that $\mathcal{S} \subseteq \hat{\mathcal{L}}(\mathcal{M})$ (as defined in subsection 2.2). If \mathcal{S} is a sample for L we often say that \mathcal{S} *agrees with* L . Given two words w, w' , we say that w and w' are *not equivalent* with respect to \mathcal{S} , denoted $w \not\sim_{\mathcal{S}} w'$, iff there exists z such that $\langle wz, b \rangle, \langle w'z, b' \rangle \in \mathcal{S}$ and $b \neq b'$. Otherwise we say that w and w' are *equivalent* with respect to \mathcal{S} , and denote $w \sim_{\mathcal{S}} w'$.

Given a sample \mathcal{S} for a language L over a concrete domain \mathbb{D} , it is possible to construct a DFA that agrees with \mathcal{S} in polynomial time. Indeed one can create the *prefix-tree automaton*, a simple automaton that accepts all and only the positively labeled words in the sample. Clearly the constructed automaton may not be the minimal automaton that agrees

with \mathcal{S} . There are several algorithms that infer such a minimal automaton, in particular the popular RPNI [OG92], that merges the states of the prefix-tree automaton and results in an automaton that may accept an infinite language. Obviously though, this procedure is not guaranteed to return an automaton for the unknown language, as the sample may not provide sufficient information. For instance if $L = aL_1 \cup bL_2$ and the sample contains only words starting with a , there is no way for the learner to infer L_2 and hence also L correctly. One may thus ask, given a language L , what should a sample contain in order for a passive learning algorithm to infer L correctly, and can such sample be of polynomial size with respect to a minimal representation (e.g., a DFA) for the language.

One approach to answer these questions is captured in the paradigm of *identification in the limit using polynomial time and data*. This model was proposed by Gold [Gol78], who also showed that it admits learning of regular languages represented by DFAs. We follow de la Higuera’s more general definition [dlH97].⁸ This definition requires that for any language L in a class of languages \mathbb{L} represented by \mathbb{R} , there exists a sample \mathcal{S}_L of size polynomial in the size of the smallest representation $R \in \mathbb{R}$ of L (e.g., the smallest DFA for L), such that a valid learner can infer the unknown language L from the information contained in \mathcal{S}_L . The set \mathcal{S}_L is then termed a *characteristic sample*.⁹ Here, a valid learner is an algorithm that learns the target language exactly and efficiently. In particular, a valid learner produces in polynomial time a representation that agrees with the provided sample. The learner also has to correctly learn the unknown language L when given the characteristic sample \mathcal{S}_L as input. Moreover, if the input sample \mathcal{S} subsumes \mathcal{S}_L yet is still consistent with L , the additional information in the sample should not “confuse” the learner; the latter still has to output a correct representation for L . (Intuitively, this requirement precludes situations in which the sample consists of some smart encoding of the representation that the learner simply deciphers. In particular, the learner will not be confused if an adversary “contaminates” the characteristic sample by adding labeled examples for the target language.) We provide the formal definition after the following informal example.

Example 7.1. For the class of DFAs, let us consider the regular language $L = a^*$ over the alphabet $\{a, b\}$. Further, consider a sample set $\mathcal{S} = \{\langle \epsilon, 1 \rangle, \langle a, 1 \rangle, \langle b, 0 \rangle, \langle bb, 0 \rangle, \langle ba, 0 \rangle\}$ for L . There is a valid learner for the class of all DFAs that uses the sample \mathcal{S} as a characteristic sample for L . By definition, such a learner has to output a DFA for L when fed with \mathcal{S} , but also has to output equivalent DFAs whenever given any superset of \mathcal{S} as input, as long as this superset agrees with L . Naturally, the sample \mathcal{S} is also consistent with the regular language $L' = \{\epsilon, a\}$. However, this does not pose any problem, since the same learner can use a characteristic sample for L' that disagrees with L , for example, $\mathcal{S}' = \{\langle \epsilon, 1 \rangle, \langle a, 1 \rangle, \langle aa, 0 \rangle\}$. When defining a system of characteristic samples like that, the core requirement is that the size of a sample be bounded from above by a function that is polynomial in the size of the smallest DFA for the respective target language.

⁸This paradigm may seem related to conformance testing. The relation between conformance testing for Mealy machines and automata learning of DFAs has been explored in [BGJ⁺05].

⁹De la Higuera’s notion of characteristic sample is a core concept in grammatical inference, for various reasons. Firstly, it addresses shortcomings of several other attempts to formulate polynomial-time learning in the limit [Ang87b, Pit89]. Secondly, this notion has inspired the design of popular algorithms for learning formal languages such as, for example, the RPNI algorithm [OG92]. Thirdly, it was shown to bear strong relations to a classical notion of machine teaching [GM96]; models of the latter kind are currently experiencing increased attention in the machine learning community [ZSZR18].

Definition 7.2 (identification in the limit using polynomial time and data [dlH97]). A class of languages \mathbb{L} is said to be *identified in the limit using polynomial time and data* via representations in a class \mathbb{R} if there exists a learning algorithm \mathbf{A} such that the following two requirements are met.

- (1) Given a finite sample \mathcal{S} of labeled examples, \mathbf{A} returns a hypothesis $\mathcal{R} \in \mathbb{R}$ that agrees with \mathcal{S} in polynomial time.
- (2) For every language $L \in \mathbb{L}$, there exists a sample \mathcal{S}_L , termed a *characteristic sample*, of size polynomial in the minimal representation $\mathcal{R} \in \mathbb{R}$ for L such that the algorithm \mathbf{A} returns a correct hypothesis when run on any sample \mathcal{S} for L that subsumes \mathcal{S}_L .

Note that the first condition ensures polynomial time and the second polynomial data. However, the latter is not a worst-case measure; the algorithm may fail to return a correct hypothesis on arbitrarily large finite samples (if they do not subsume a characteristic set).

Note also that the definition does not require the existence of an efficient algorithm that constructs a characteristic sample for each language in the underlying class. When such an algorithm is also available we say that the class is *efficiently identifiable*. The following result shows that efficient identifiability does not trivially follow from identifiability; in fact it makes the much stronger statement that not even *computability* of characteristic sets follows from their existence.

Theorem 7.3. *There exists a class of languages that possesses polynomial-size characteristic sets, yet without the ability to construct such sets effectively.*

We first prove Theorem 7.3 and then provide a definition for efficient identification.

Proof. We present a class of recursive languages that is identifiable in the limit from polynomial time and data, while there is no (polynomial-time or other) algorithm that constructs a characteristic sample for every language in the class, using a specific underlying representation of the languages in the class.

For the purpose of defining such a class, let φ be a Gödel numbering of all partial computable functions over the natural numbers, and let Φ be a corresponding Blum complexity measure. Here φ_i refers to the i th partial computable function in the numbering φ . Intuitively, $\Phi_i(j)$ is undefined if $\varphi_i(j)$ is undefined (i.e., the computation of $\varphi_i(j)$ does not terminate); otherwise $\Phi_i(j)$ is the number of computational steps required until the termination of the computation of $\varphi_i(j)$. The set $K = \{k \mid \varphi_k(k) \text{ is defined}\}$ is called the halting set for φ ; this set is recursively enumerable but membership in K is not decidable.

We now define two languages for each natural number k :

$$\begin{aligned} L_{k,1} &= a^* \cup \{b^k\} \\ L_{k,2} &= \begin{cases} \{a^i \mid i \leq \Phi_k(k)\} \cup \{b^k\}, & \text{if } k \in K, \\ a^* \cup \{b^k\}, & \text{otherwise.} \end{cases} \end{aligned}$$

Note that $L_{k,1} = L_{k,2}$ if and only if $k \notin K$. Now let \mathbb{L} consist of all languages $L_{k,q}$ for $k \in \mathbb{N}$ and $q \in \{1, 2\}$.

There is an effective algorithm that decides membership in $L_{k,q}$, given $k \in \mathbb{N}$ and $q \in \{1, 2\}$. To see that, note that, given k, q , and a word w , membership is trivial to decide when $q = 1$ or when w is not of the form a^i . If $w = a^i$, then w belongs to $L_{k,2}$ if and only if the computation of $\varphi_k(k)$ does not terminate within fewer than i steps, which can be checked effectively.

Moreover, every language in \mathbb{L} is regular and has a characteristic sample of size at most 2. In particular, $\{\langle b^k, 1 \rangle\}$ serves as a characteristic sample for $L_{k,1}$ (and thus also for $L_{k,2}$ in case $k \notin K$), while $\{\langle b^k, 1 \rangle, \langle a^{\Phi_k(k)+1}, 0 \rangle\}$ is a characteristic sample for $L_{k,2}$ in case $k \in K$. Thus, using the above representation, the class \mathbb{L} has polynomial-size (even constant-size) characteristic samples. However, there is no algorithm to construct such characteristic samples effectively, since otherwise such an algorithm could be used to decide membership in K . (The latter can be verified by noting that $L_{k,2} \subseteq L_{k,1}$ for all k . Therefore, a system of characteristic samples would need to distinguish $L_{k,2}$ from $L_{k,1}$ (when $k \in K$) by either (i) a negative example of the form $\langle a^i, 0 \rangle$ for $L_{k,2}$, or (ii) a positive example of the form $\langle a^i, 1 \rangle$ for $L_{k,1}$, where $a^i \notin L_{k,2}$. Thus, the presence or absence of such example in the characteristic samples for $L_{k,1}$, $L_{k,2}$ can be used to decide whether or not $k \in K$.) \square

Since we are concerned with learning classes of automata, we now formulate the definition of *efficient identification* directly over classes of automata.

Definition 7.4 (efficient identification). A class of automata \mathbb{M} over an alphabet Σ is said to be *efficiently identified* if the following two requirements are met.

- (1) There exists a polynomial time learning algorithm $\mathbf{Infer} : 2^{\Sigma^* \times \{0,1\}} \rightarrow \mathbb{M}$ such that, for any sample \mathcal{S} , we have $\mathcal{S} \subseteq \hat{\mathcal{L}}(\mathbf{Infer}(\mathcal{S}))$.
- (2) There exists a polynomial time algorithm $\mathbf{Char} : \mathbb{M} \rightarrow 2^{\Sigma^* \times \{0,1\}}$ such that, for every $\mathcal{M} \in \mathbb{M}$ and every sample \mathcal{S} satisfying $\mathbf{Char}(\mathcal{M}) \subseteq \mathcal{S} \subseteq \hat{\mathcal{L}}(\mathcal{M})$, the automaton $\mathbf{Infer}(\mathcal{S})$ recognizes the same language as \mathcal{M} .

When we apply this definition for a class of SFAs over a Boolean algebra \mathcal{A} with domain \mathbb{D} and predicates \mathbb{P} , the characteristic sample is defined over the concrete set of letters \mathbb{D} rather than the set of predicates \mathbb{P} as this is the alphabet of the words accepted by an SFA. (Inferring an SFA from a set of words labeled by predicates can be done using the methods for inferring DFAs, by considering the alphabet to be the set of predicates.)

Throughout this section we study whether a class of SFAs \mathbb{M} is efficiently identifiable. That is, we are interested in the existence of algorithms $\mathbf{Infer}_{\mathbb{M}}$ and $\mathbf{Char}_{\mathbb{M}}$ satisfying the requirements of Definition 7.4. In section 8 we provide a necessary condition for a class of SFAs to be identified in the limit using polynomial time and data. In section 9 we provide a sufficient condition for a class of SFAs to be efficiently identifiable. On the positive side, we show in section 10 that the class of SFAs over the interval algebra is efficiently identifiable. On the negative side, we show in section 11 that SFAs over the general propositional algebra cannot be identified in the limit using polynomial time and data. All classes of SFAs that we study are assumed to contain all basic SFAs (as per Definition 6.1).

7.1. Efficient Identification of DFAs. Before investigating efficient identification of SFAs, it is worth noting that DFAs are efficiently identifiable. We state a result that provides more details about the nature of these algorithms, since we need it later, in section 10, to provide our positive result. Intuitively, it says that there exists a valid learner such that if \mathcal{D} is a minimal DFA recognizing a certain language L then the learner can infer L from a characteristic sample consisting of access words to each state of \mathcal{D} and their extensions with distinguishing words (words showing each pair of states cannot be merged) as well as one letter extensions of the access words that are required to retrieve the transition relation. For completeness we give a proof of this theorem in Appendix A.

Theorem 7.5 [OG92]. *The class of DFAs is efficiently identifiable via procedures **CharDFA** and **InferDFA**. Furthermore, these procedures satisfy that if \mathcal{D} is a minimal and complete DFA and $\mathbf{CharDFA}(\mathcal{D}) = \mathcal{S}_{\mathcal{D}}$ then the following hold:*

- (1) $\mathcal{S}_{\mathcal{D}}$ contains a prefix-closed set A of access words. Moreover, A can be chosen to contain only lex-access words, i.e., only the lexicographically smallest access word for each state.
- (2) For every $u_1, u_2 \in A$ it holds that $u_1 \not\sim_{\mathcal{S}_{\mathcal{D}}} u_2$.
- (3) For every $u, v \in A$ and $\sigma \in \Sigma$, if $\Delta(q_i, u\sigma) \neq \Delta(q_i, v)$ then $u\sigma \not\sim_{\mathcal{S}_{\mathcal{D}}} v$.

We briefly describe **CharDFA** and **InferDFA**.

The algorithm **CharDFA** works as follows. It first creates a prefix-closed set of access words to states. This can be done by considering the graph of the automaton and running an algorithm for finding a spanning tree from the initial state. Choosing one of the letters on each edge, the access word for a state is obtained by concatenating the labels on the unique path of the obtained tree that reaches that state. If we wish to work with lex-access words, we can use a depth-first search algorithm that spans branches according to the order of letters in Σ , starting from the smallest. The labels on the paths of the spanning tree constructed this way will form the set of lex-access words.

Let S be the set of access words (or lex-access words). Next the algorithm turns to find a distinguishing word $v_{i,j}$ for every pair of states $s_i, s_j \in S$ (where $s_i \neq s_j$). It holds that any pair of states of the minimal DFA has a distinguishing word of size quadratic in the size of the DFA. Let E be the set of all such distinguishing words. Then the algorithm returns the set $\mathcal{S}_{\mathcal{D}} = \{\langle w, \mathcal{D}(w) \rangle \mid w \in (S \cdot E) \cup (S \cdot \Sigma \cdot E)\}$ where $\mathcal{D}(w)$ is the label \mathcal{D} gives w (i.e., 1 if it is accepted, and 0 otherwise). It is easy to see that $\mathcal{S}_{\mathcal{D}}$ satisfies the properties of Theorem 7.5.

The algorithm **InferDFA**, given a sample of words \mathcal{S} , infers from it in polynomial time a DFA that agrees with \mathcal{S} . Moreover, if \mathcal{S} subsumes the characteristic set $\mathcal{S}_{\mathcal{D}}$ of a DFA \mathcal{D} then **InferDFA** returns a DFA that recognizes \mathcal{D} . Let W be the set of words in the given sample \mathcal{S} (without their labels). Let R be the set of prefixes of W and C the set of suffixes of W . Note that $\epsilon \in R$ and $\epsilon \in C$. Let r_0, r_1, \dots be some enumeration of R and c_0, c_1, \dots some enumeration of C where $r_0 = c_0 = \epsilon$. The algorithm builds a matrix M of size $|R| \times |C|$ whose entries take values in $\{0, 1, ?\}$, and sets the value of entry (i, j) as follows. If $r_i c_j$ is not in W , it is set to ?. Otherwise it is set to 1 iff the word $r_i c_j$ is labeled 1 in \mathcal{S} . We get that $r_i \sim_{\mathcal{S}} r_j$ iff for every k such that both $M(i, k)$ and $M(j, k)$ are different than ? we have that $M(i, k) = M(j, k)$. The algorithm sets $R_0 = \{\epsilon\}$. Once R_i is constructed, the algorithm tries to establish whether for $r \in R_i$ and $\sigma \in \Sigma$, $r\sigma$ is distinguished from all words in R_i . It does so by considering all other words $r' \in R_i$ and checking whether $r \sim_{\mathcal{S}} r'$. If $r\sigma$ is found to be distinct from all words in R_i , then R_{i+1} is set to $R_i \cup \{r\sigma\}$. The algorithm proceeds until no new words are distinguished.¹⁰

Let k be the minimal iteration such that $R_k = R_{k'}$ for all $k' > k$. If not all words in R_k are in W (that is $M(i, 0) = ?$ for some $r_i \in R_k$), the algorithm returns the prefix-tree automaton. Otherwise, the states of the constructed DFA are set to be the words in R_k . The initial state is ϵ and a state r_i is classified as accepting iff $M(i, 0) = 1$ (recall that the entry $M(i, 0)$ stands for the value of $r_i \cdot \epsilon$ in \mathcal{S}). To determine the transitions, for every $r \in R_k$ and $\sigma \in \Sigma$, recall that there exists at least one state $r' \in R$ that cannot be distinguished from $r\sigma$. The algorithm then adds a transition from r on σ to r' .

¹⁰There is some resemblance between the partial matrix and the apartness relation of [VGRW22], used in query learning of DFAs.

8. NECESSARY CONDITION

We make use of the following definitions. A sequence $\langle \Gamma_1, \dots, \Gamma_m \rangle$ consisting of finite sets of concrete letters $\Gamma_i \subseteq \mathbb{D}$ is termed a *concrete partition* of \mathbb{D} if the sets are pairwise disjoint (namely $\Gamma_i \cap \Gamma_j = \emptyset$ for every $i \neq j$). Note that we do not require that in addition $\bigcup_{1 \leq i \leq k} \Gamma_i = \mathbb{D}$. We use $\Pi_{\text{conc}}(\mathbb{D})$ to define the set of all concrete partitions over \mathbb{D} . A sequence of predicates $\langle \psi_1, \dots, \psi_k \rangle$ over a Boolean algebra \mathcal{A} on a domain \mathbb{D} is termed a *predicate partition* if $\llbracket \psi_i \rrbracket \cap \llbracket \psi_j \rrbracket = \emptyset$ for every $i \neq j$, and in addition $\bigcup_{1 \leq i \leq k} \llbracket \psi_i \rrbracket = \mathbb{D}$. That is, here we do require the assignments to the predicates cover the domain. We use $\Pi_{\text{pred}}(\mathbb{P})$ to define the set of all predicate partitions over \mathbb{P} . For both concrete partition and predicate partition, we do not require that the sets Γ_i or $\llbracket \psi_i \rrbracket$ are non-empty.

Definition 8.1.

- A function $f_c : \Pi_{\text{pred}}(\mathbb{P}) \rightarrow \Pi_{\text{conc}}(\mathbb{D})$ is termed *concretizing* if $f_c(\langle \psi_1, \dots, \psi_m \rangle) = \langle \Gamma_1, \dots, \Gamma_k \rangle$ implies $k = m$ and $\Gamma_i \subseteq \llbracket \psi_i \rrbracket$ for all $1 \leq i \leq m$.
- A function $f_g : \Pi_{\text{conc}}(\mathbb{D}) \rightarrow \Pi_{\text{pred}}(\mathbb{P})$ is termed *generalizing* if $f_g(\langle \Gamma_1, \dots, \Gamma_m \rangle) = \langle \psi_1, \dots, \psi_k \rangle$ implies $k = m$ and $\llbracket \psi_i \rrbracket \supseteq \Gamma_i$ for all $1 \leq i \leq m$.

Note that f_g and f_c are defined over partitions of any size. In Theorem 8.2 we use their *dyadic* restriction, that is, a concretizing and a generalizing functions that are defined only over partitions of size two.

We say that f_g (resp. f_c) is *efficient* if it can be computed in polynomial time. Note that if f_c is efficient then the sets Γ_i in the constructed concrete partition are of polynomial size.

We are now ready to provide a necessary condition for identifiability in the limit using polynomial time and data.

Theorem 8.2. *Let $\mathbb{M}_{\mathcal{A}}$ be a class of SFAs over a Boolean algebra \mathcal{A} , that contain all basic SFAs over \mathcal{A} . If $\mathbb{M}_{\mathcal{A}}$ is identified in the limit using polynomial time and data, then there exist efficient dyadic concretizing and generalizing functions $\text{Concretize}_{\mathcal{A}} : \Pi_{\text{pred}}(\mathbb{P}) \rightarrow \Pi_{\text{conc}}(\mathbb{D})$ and $\text{Generalize}_{\mathcal{A}} : \Pi_{\text{conc}}(\mathbb{D}) \rightarrow \Pi_{\text{pred}}(\mathbb{P})$ satisfying that*

$$\begin{aligned} & \text{if } \text{Concretize}_{\mathcal{A}}(\langle \psi_1, \psi_2 \rangle) = \langle \Gamma_1, \Gamma_2 \rangle \\ & \text{and } \text{Generalize}_{\mathcal{A}}(\langle \Gamma'_1, \Gamma'_2 \rangle) = \langle \varphi_1, \varphi_2 \rangle \\ & \text{where } \Gamma_i \subseteq \Gamma'_i \text{ for every } 1 \leq i \leq 2 \\ & \text{then } \llbracket \varphi_i \rrbracket = \llbracket \psi_i \rrbracket \text{ for every } 1 \leq i \leq 2. \end{aligned}$$

Proof. Assume that $\mathbb{M}_{\mathcal{A}}$ is identified in the limit using polynomial time and data. That is, there exist two algorithms $\text{CharSFA} : \mathbb{M}_{\mathcal{A}} \rightarrow 2^{\mathbb{D}^* \times \{0,1\}}$ and $\text{InferSFA} : 2^{\mathbb{D}^* \times \{0,1\}} \rightarrow \mathbb{M}_{\mathcal{A}}$ satisfying the requirements of Definition 7.2. We show that efficient dyadic concretizing and generalizing functions do exist.

We start with the definition of $\text{Concretize}_{\mathcal{A}}$. Let $\langle \varphi_1, \varphi_2 \rangle$ be the argument of $\text{Concretize}_{\mathcal{A}}$. Note that $\varphi_2 = \neg \varphi_1$ by the definition of a predicate partition. The implementation of $\text{Concretize}_{\mathcal{A}}$ invokes CharSFA on the SFA \mathcal{M}_{φ_1} accepting all words of length one consisting of a concrete letter satisfying φ_1 , as defined in Definition 6.1. Let \mathcal{S} be the returned sample. Let Γ_1 be the set of positively labeled words in the sample. Note that all such words are of size one, namely they are letters. Let Γ_2 be the set of letters that are first letters in a negative word in the sample. Then $\text{Concretize}_{\mathcal{A}}$ returns $\langle \Gamma_1, \Gamma_2 \rangle$.

We turn to the definition of $\text{Generalize}_{\mathcal{A}}$. Given $\langle \Gamma_1, \Gamma_2 \rangle$ the implementation of $\text{Generalize}_{\mathcal{A}}$ invokes InferSFA on sample $\mathcal{S} = \{(\gamma, 1) \mid \gamma \in \Gamma_1\} \cup \{(\gamma, 0) \mid \gamma \in \Gamma_2\} \cup \{(\gamma\gamma', 0) \mid \gamma, \gamma' \in \Gamma_1 \cup \Gamma_2\}$. That is, all one-letter words in Γ_1 are positively labeled, all one-letter words in Γ_2 are

negatively labeled, and all words of length 2 using some of the given concrete letters are negatively labeled. Let \mathcal{M} be the returned SFA when given \mathcal{S}' , such that $\mathcal{S}' \supseteq \mathcal{S}$, as an input. Let Ψ_1 be the set of all predicates labeling some edge from the initial state to an accepting state, and let Ψ_2 be the set of all predicates labeling some edge from the initial state to a rejecting state. Let $\varphi = (\bigvee_{\psi \in \Psi_1} \psi) \wedge (\bigwedge_{\psi \in \Psi_2} \neg\psi)$. Then $\text{Generalize}_{\mathcal{A}}$ returns $\langle \varphi, \neg\varphi \rangle$.

It is not hard to verify that the constructed methods $\text{Concretize}_{\mathcal{A}}$ and $\text{Generalize}_{\mathcal{A}}$ satisfy the requirements of the theorem. \square

The following example shows the existence of functions Concretize and Generalize for the interval algebra.

Example 8.3. Consider the class $\mathbb{M}_{\mathcal{A}_{\mathbb{N}}}$ of SFAs over the algebra $\mathcal{A}_{\mathbb{N}}$ of Example 2.2 and consider the functions $\text{Concretize}_{\mathcal{A}_{\mathbb{N}}}(\langle [d_1, d'_1], [d_2, d'_2], \dots, [d_m, d'_m] \rangle) = \langle \{d_1\}, \dots, \{d_m\} \rangle$ and $\text{Generalize}_{\mathcal{A}_{\mathbb{N}}}(\langle \Gamma_1, \dots, \Gamma_m \rangle) = \langle [\min \Gamma_1, \min \hat{\Gamma}_1], [\min \Gamma_2, \min \hat{\Gamma}_2], \dots, [\min \Gamma_m, \infty] \rangle$ where $\hat{\Gamma}_i = \bigcup_{j \neq i} \Gamma_j$ for every $1 \leq i < m$. Then, $\text{Concretize}_{\mathcal{A}_{\mathbb{N}}}$ and $\text{Generalize}_{\mathcal{A}_{\mathbb{N}}}$ satisfy the variadic generalization of the conditions of Theorem 8.2.


We would like to relate the necessary condition on the learnability of a class of SFAs over a Boolean algebra \mathcal{A} to the learnability of the Boolean algebra \mathcal{A} itself. For this we need to first define efficient identifiability of a Boolean algebra \mathcal{A} . Since to learn an unknown predicate we need to supply two sets, one of negative examples and one of positive examples, it makes sense to say that a Boolean algebra \mathcal{A} with predicates \mathbb{P} over domain \mathbb{D} is *efficiently identifiable* if there exist efficient dyadic concretizing and generalizing functions, $\text{Concretize}_{\mathcal{A}} : \Pi_{\text{pred}}(\mathbb{P}) \rightarrow \Pi_{\text{conc}}(\mathbb{D})$ and $\text{Generalize}_{\mathcal{A}} : \Pi_{\text{conc}}(\mathbb{D}) \rightarrow \Pi_{\text{pred}}(\mathbb{P})$ satisfying the criteria of Theorem 8.2. Using this terminology we can state the following corollary.

Corollary 8.4. *Efficient identifiability of the Boolean algebra \mathcal{A} is a necessary condition for identification in the limit using polynomial time and data of any class of SFAs over \mathcal{A} , that contains all basic SFAs over \mathcal{A} .*

9. SUFFICIENT CONDITION

We turn to discuss a sufficient condition for the efficient identifiability of a class of SFAs $\mathbb{M}_{\mathcal{A}}$ over a Boolean algebra \mathcal{A} . To prove that $\mathbb{M}_{\mathcal{A}}$ is efficiently identifiable, we need to supply two algorithms $\text{CharSFA}_{\mathbb{M}_{\mathcal{A}}}$ and $\text{InferSFA}_{\mathbb{M}_{\mathcal{A}}}$ as required in Definition 7.4. The idea is to reduce the problem to efficient identifiability of DFAs, namely to use the algorithms CharDFA and InferDFA provided in Theorem 7.5. The implementation of CharSFA , given an SFA \mathcal{M} , will transform it into a DFA $\mathcal{D}_{\mathcal{M}}$ by applying $\text{Concretize}_{\mathcal{A}}$ on the partitions induced by the states of the SFA. The resulting DFA $\mathcal{D}_{\mathcal{M}}$ will not be equivalent to the given SFA \mathcal{M} , but it may be used to create a sample of words $\mathcal{S}_{\mathcal{M}}$ that is a characteristic set for \mathcal{M} , see Figure 2.

To implement InferSFA we would like to use InferDFA to obtain, as a first step, a DFA from the given sample, then at the second step, apply $\text{Generalize}_{\mathcal{A}}$ on the concrete-partitions induced by the DFA states. A subtle issue that we need to cope with is that inference should succeed also on samples subsuming the characteristic sample. The fact that this holds for inference of the DFA does not suffice, since we are guaranteed that the inference of the DFA will not be confused if the sample contains more labeled words, as long as the new words are over the same alphabet. In our case the alphabet of the sample can be a

FIGURE 2. A schematic description of algorithms **CharSFA** and **InferSFA**

Input: An SFA \mathcal{M} , function $\text{Concretize}_{\mathcal{A}}$
Output: A DFA $\mathcal{D}_{\mathcal{M}}$

```

1 function CONCRETIZE $\mathbb{M}_{\mathcal{A}}$ ( $\mathcal{M} = \langle \mathcal{A}, Q, q_i, F, \Delta \rangle$ )
2    $\Gamma_{\mathcal{M}} := \bigcup_{q \in Q} \text{Concretize}_{\mathcal{A}}(\pi_q)$ 
3    $\Delta_{\mathcal{D}} := \emptyset$ 
4   for all  $q, q' \in Q, \psi \in \pi_q, d \in \Gamma_{\mathcal{M}}$  do
5     if  $\langle q, \psi, q' \rangle \in \Delta$  and  $d \in \llbracket \psi \rrbracket$  then
6        $\Delta_{\mathcal{D}} := \Delta_{\mathcal{D}} \cup \langle q, d, q' \rangle$ 
7   return  $\mathcal{D}_{\mathcal{M}} := \langle \Gamma_{\mathcal{M}}, Q, q_i, F, \Delta_{\mathcal{D}} \rangle$ 

```

ALGORITHM 1. $\text{Concretize}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{M})$

Input: A DFA \mathcal{D} , function $\text{Generalize}_{\mathcal{A}}$
Output: An SFA \mathcal{M}

```

1 function GENERALIZE $\mathbb{M}_{\mathcal{A}}$ ( $\mathcal{D} = \langle \Sigma, Q, q_i, F, \Delta_{\mathcal{D}} \rangle$ )
2    $\Delta_{\mathcal{M}} := \emptyset$ 
3   for all  $q \in Q$  do
4     for all  $q_i \in Q$  do  $\Gamma_i := \{ \gamma \mid \langle q, \gamma, q_i \rangle \in \Delta_{\mathcal{D}} \}$ 
5      $\langle \psi_1, \dots, \psi_n \rangle := \text{Generalize}_{\mathcal{A}}(\langle \Gamma_1, \dots, \Gamma_n \rangle)$ 
6     for all  $q_i \in Q$  do  $\Delta_{\mathcal{M}} := \Delta_{\mathcal{M}} \cup \langle q, \psi_i, q_i \rangle$ 
7   return  $\mathcal{M} := \langle \mathcal{A}, Q, q_i, F, \Delta_{\mathcal{M}} \rangle$ 

```

ALGORITHM 2. $\text{Generalize}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{M})$

strict subset of the concrete letters \mathbb{D} (and if \mathbb{D} is infinite, this surely will be the case). Example 10.7 in section 10 illustrates this problem for the class of SFAs over a monotonic algebra \mathcal{A}_m , for which the respective methods $\text{Concretize}_{\mathcal{A}_m}$ and $\text{Generalize}_{\mathcal{A}_m}$ exist. So, we need an additional step to remove words from the given sample if they are not over the alphabet of the characteristic sample. We call a method implementing this $\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}}}$.

Formally, we first define the extension of $\text{Concretize}_{\mathcal{A}}$ and $\text{Generalize}_{\mathcal{A}}$ to automata instead of partitions, which we term $\text{Concretize}_{\mathbb{M}_{\mathcal{A}}}$ and $\text{Generalize}_{\mathbb{M}_{\mathcal{A}}}$ (with \mathbb{M} in the subscript).

- The formal definition of $\text{Concretize}_{\mathbb{M}_{\mathcal{A}}}$ is given in Algorithm 1. Let $\mathcal{M} = (\mathcal{A}, Q, q_i, F, \Delta)$ be an SFA. Then $\text{Concretize}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{M})$ is the DFA $\mathcal{D}_{\mathcal{M}} = (\Sigma, Q, q_i, F, \Delta_{\mathcal{D}})$ where $\Delta_{\mathcal{D}}$ is defined as follows. For each state $q \in Q$ let $\pi_q = \langle \psi_1, \dots, \psi_m \rangle$ be the predicate partition consisting of all predicates labeling a transition exiting q in \mathcal{M} . Intuitively, in \mathcal{D} , the outgoing transitions of each state q correspond to $\text{Concretize}_{\mathcal{A}}(\pi_q)$. That is, let $\text{Concretize}_{\mathcal{A}}(\pi_q) = \langle \Gamma_1, \dots, \Gamma_m \rangle$. Then, if $\langle q, \psi_i, q' \rangle \in \Delta$, then $\langle q, \gamma, q' \rangle \in \Delta_{\mathcal{D}}$ for every $\gamma \in \Gamma_i$.
- The formal definition of $\text{Generalize}_{\mathbb{M}_{\mathcal{A}}}$ is given in Algorithm 2. Let $\mathcal{D} = (\Sigma, Q, q_i, F, \Delta_{\mathcal{D}})$ be a DFA. We define $\text{Generalize}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{D})$ with respect to an algebra \mathcal{A} as follows. Let $\mathcal{M} = (\mathcal{A}, Q, q_i, F, \Delta_{\mathcal{M}})$ where $\Delta_{\mathcal{M}}$ is defined as follows. For each state $q \in Q$ let $\langle \Gamma_1, \dots, \Gamma_m \rangle$ be the concrete partition consisting of letters labeling outgoing transitions from q . Note that $\langle \Gamma_1, \dots, \Gamma_m \rangle$ is a concrete partition, since \mathcal{D} is a DFA. Let $\text{Generalize}_{\mathcal{A}}(\langle \Gamma_1, \dots, \Gamma_m \rangle) = \langle \psi_1, \dots, \psi_m \rangle$. Then, $\langle q, \psi_i, q' \rangle \in \Delta_{\mathcal{M}}$ if Γ_i is the set of letters labeling transitions from q to q' in \mathcal{D} .

We are now ready to define the conditions the *decontaminating* function has to satisfy. We recall that the role of the decontaminating function is to identify words in the sample that are not over the alphabet $\Gamma_{\mathcal{D}}$ of the characteristic sample (note that $\Gamma_{\mathcal{D}}$ is not known to the decontaminating function).

Definition 9.1. A function $f_d : 2^{(\mathbb{D}^* \times \{0,1\})} \rightarrow 2^{(\mathbb{D}^* \times \{0,1\})}$ is called *decontaminating* for a class of SFAs \mathbb{M} and a respective $\text{Concretize}_{\mathbb{M}}$ function if the following holds. Let $\mathcal{M} \in \mathbb{M}$ be

an SFA, and let $\mathcal{D} = \text{Concretize}_{\mathbb{M}}(\mathcal{M})$. Let $\mathcal{S}_{\mathcal{D}} = \text{CharDFA}(\mathcal{D})$. Then, for every $\mathcal{S}' \supseteq \mathcal{S}_{\mathcal{D}}$ such that \mathcal{S}' agrees with \mathcal{M} , it holds that $\mathcal{S}_{\mathcal{D}} \subseteq f_d(\mathcal{S}') \subseteq (\mathcal{S}' \cap \Gamma_{\mathcal{D}})$, where $\Gamma_{\mathcal{D}}$ is the alphabet of $\mathcal{S}_{\mathcal{D}}$.

As before we say that f_d is *efficient* if it can be computed in polynomial time.

Example 9.2. Intuitively, **InferDFA** is only promised to be correct if it is applied on a sample set \mathcal{S}' over the alphabet of the original DFA. For DFAs, this is always the case. However for SFAs, the concrete alphabet \mathbb{D} usually contains more letters than appear in the characteristic set \mathcal{S} . If $\Gamma_{\mathcal{D}}$ is the set of letters in \mathcal{S} then \mathcal{S}' might contain letters in $\mathbb{D} \setminus \Gamma_{\mathcal{D}}$, i.e., letters that are not from the alphabet of the characteristic set. Consider, for example, the characteristic set \mathcal{S} over the interval algebra:

$$\mathcal{S} = \{\langle \epsilon, 0 \rangle, \langle 0, 1 \rangle, \langle 100, 0 \rangle, \langle 200, 0 \rangle, \langle 0 \cdot 0, 1 \rangle, \langle 0 \cdot 100, 1 \rangle, \langle 0 \cdot 200, 0 \rangle\}$$

and consider the set $\mathcal{S}' \supseteq \mathcal{S}$ that contains, in addition to the words in \mathcal{S} , also the word $150 \cdot 100$. Since the letter 150 is not part of any word in the original set \mathcal{S} , we cannot apply **InferDFA** as it is, but we first need to remove it from \mathcal{S}' . Note that words that are not in \mathcal{S} but are over the alphabet $\{0, 100, 200\}$ do not pose a problem, as **InferDFA** can handle supersets over the same alphabet as the set \mathcal{S} . See Examples 10.6 and 10.7 in section 10 for more details.

We now provide the sufficient condition for efficient identifiability.

Theorem 9.3. *Let $\mathbb{M}_{\mathcal{A}}$ be a class of SFAs over a Boolean algebra \mathcal{A} . If there exist an efficient decontaminating function $\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}}}$ and efficient functions $\text{Concretize}_{\mathcal{A}}$ and $\text{Generalize}_{\mathcal{A}}$ satisfying that*

$$\begin{aligned} & \text{if } \text{Concretize}_{\mathcal{A}}(\langle \psi_1, \dots, \psi_m \rangle) = \langle \Gamma_1, \dots, \Gamma_m \rangle \\ & \text{and } \text{Generalize}_{\mathcal{A}}(\langle \Gamma'_1, \dots, \Gamma'_m \rangle) = \langle \varphi_1, \dots, \varphi_m \rangle \\ & \text{where } \Gamma_i \subseteq \Gamma'_i \text{ for every } 1 \leq i \leq m \\ & \text{then } \llbracket \varphi_i \rrbracket = \llbracket \psi_i \rrbracket \text{ for every } 1 \leq i \leq m \end{aligned}$$

then the class $\mathbb{M}_{\mathcal{A}}$ is efficiently identifiable.

Given functions $\text{Concretize}_{\mathcal{A}}$, $\text{Generalize}_{\mathcal{A}}$ and $\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}}}$ for a class $\mathbb{M}_{\mathcal{A}}$ of SFAs over a Boolean algebra \mathcal{A} , meeting the criteria of Theorem 9.3, we show that $\mathbb{M}_{\mathcal{A}}$ can be efficiently identified by providing two algorithms **CharSFA** and **InferSFA**, described below. These algorithms make use of the respective algorithms **CharDFA** and **InferDFA** guaranteed in Theorem 7.5, as well as the methods provided by the theorem.

We briefly describe these two algorithms, and then turn to prove Theorem 9.3. The algorithm **CharSFA** receives an SFA $\mathcal{M} \in \mathbb{M}$, and returns a characteristic sample for it. It does so by applying $\text{Concretize}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{M})$ (Algorithm 1) to construct a DFA $\mathcal{D}_{\mathcal{M}}$ and generating the sample $\mathcal{S}_{\mathcal{M}}$ using the algorithm **CharDFA** applied on the DFA $\mathcal{D}_{\mathcal{M}}$.

Algorithm **InferSFA**, given a sample set \mathcal{S} , if \mathcal{S} subsumes a characteristic set of an SFA \mathcal{M} , returns an equivalent SFA. Otherwise **InferSFA** returns an SFA that agrees with the sample \mathcal{S} . First, it applies $\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}}}$ to find a subset $\mathcal{S}' \subseteq \mathcal{S}$ over the alphabet of the subsumed characteristic sample, if such a subsumed sample exists. Then it uses \mathcal{S}' to construct a DFA by applying the inference algorithm **InferDFA** on \mathcal{S}' . From this DFA it constructs an SFA, $\mathcal{M}_{\mathcal{S}}$, by applying $\text{Generalize}_{\mathbb{M}_{\mathcal{A}}}$ (Algorithm 2). If the resulting automaton disagrees with the given sample it resorts to returning the prefix-tree automaton. In order to construct the symbolic prefix-tree automaton we first construct the prefix-tree DFA \mathcal{A} for the set \mathcal{S} , and then apply $\text{Generalize}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{A})$ to get an SFA that agrees with \mathcal{S} .

In brief, we define:

- **CharSFA**(\mathcal{M}) = **CharDFA**(**Concretize** $_{\mathbb{M}_{\mathcal{A}}}$ (\mathcal{M}))
- **InferSFA**(\mathcal{S}) = $\begin{cases} \mathcal{M}_{\mathcal{S}} := \text{Generalize}_{\mathbb{M}_{\mathcal{A}}}(\text{InferDFA}(\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{S}))) & \text{if } \mathcal{S} \subseteq \hat{\mathcal{L}}(\mathcal{M}_{\mathcal{S}}) \\ \text{The symbolic prefix-tree automaton of } \mathcal{S} & \text{otherwise} \end{cases}$

In section 10 we provide methods **Concretize** $_{\mathcal{A}}$, **Generalize** $_{\mathcal{A}}$ and **Decontaminate** $_{\mathbb{M}_{\mathcal{A}}}$ for SFAs over monotonic algebras, deriving their identification in the limit result. We now prove Theorem 9.3.

Proof of Theorem 9.3. Given functions **Concretize** $_{\mathcal{A}}$, **Generalize** $_{\mathcal{A}}$, and **Decontaminate** $_{\mathbb{M}_{\mathcal{A}}}$, we show that the algorithms **CharSFA** and **InferSFA** satisfy the requirements of Definition 7.4.

For the first condition, given that **CharDFA**, **Decontaminate** $_{\mathbb{M}_{\mathcal{A}}}$ and **Generalize** $_{\mathcal{A}}$ run in polynomial time, and that the prefix-tree automaton can be constructed in polynomial time, it is clear that so does **InferSFA**. In addition, the test performed in the definition of **InferSFA** ensures the output agrees with the sample.

For the second condition, note that the sample generated by **CharSFA** is polynomial in the size of $\mathcal{D}_{\mathcal{M}}$, from the correctness of **CharDFA**. In addition, since **Concretize** $_{\mathcal{A}}$ is efficient, $\mathcal{D}_{\mathcal{M}}$ is polynomial in the size of \mathcal{M} , and thus $\mathcal{S}_{\mathcal{M}}$ generated by **CharSFA** is polynomial in \mathcal{M} as well. It is left to show that given $\mathcal{S}_{\mathcal{M}}$ is the concrete sample produced by **CharSFA** when running on an SFA \mathcal{M} , then when **InferSFA** runs on any sample $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{M}}$ it returns an SFA for $\mathcal{L}(\mathcal{M})$. Since **Decontaminate** $_{\mathbb{M}_{\mathcal{A}}}$ is a decontaminating function, and $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{M}}$, it holds that the set $\mathcal{S}' = \text{Decontaminate}_{\mathbb{M}_{\mathcal{A}}}(\mathcal{S})$ is such that $\mathcal{S}' \supseteq \mathcal{S}_{\mathcal{M}}$ and is only over the alphabet $\Gamma_{\mathcal{M}}$, which is the alphabet of the DFA $\mathcal{D}_{\mathcal{M}}$ generated in Algorithm 1.

From the correctness of **InferDFA**, given $\mathcal{S}' \supseteq \mathcal{S}_{\mathcal{M}}$, applying **InferDFA** on the output \mathcal{S}' of **Decontaminate** $_{\mathbb{M}_{\mathcal{A}}}$ results in a DFA \mathcal{D} that is equivalent to $\mathcal{D}_{\mathcal{M}}$ constructed in Algorithm 1. Since $\mathcal{D}_{\mathcal{M}}$ is complete with respect to its alphabet $\Gamma_{\mathcal{M}}$, for state q of \mathcal{D} , the concrete partition $\langle \Gamma_1, \dots, \Gamma_n \rangle$ generated in Algorithm 2, line 4, covers $\Gamma_{\mathcal{M}}$ and subsumes the output of **Concretize** $_{\mathbb{M}_{\mathcal{A}}}$ on π_q (Algorithm 1, line 2). Thus, since **Generalize** $_{\mathcal{A}}$ and **Concretize** $_{\mathcal{A}}$ satisfy the criteria of Theorem 9.3, it holds that the constructed predicates agree with the original predicates. In addition, since \mathcal{S} , and therefore \mathcal{S}' , agrees with \mathcal{M} , the test performed in the definition of **InferSFA** succeeds and the returned SFA is equivalent to \mathcal{M} . \square

10. POSITIVE RESULT

We present the following positive result regarding monotonic algebras.

Theorem 10.1. *Let $\mathbb{M}_{\mathcal{A}_m}$ be the set of SFAs over a monotonic Boolean algebra \mathcal{A}_m . Then $\mathbb{M}_{\mathcal{A}_m}$ is efficiently identifiable.*

In order to prove Theorem 10.1, we show that the sufficient condition holds for the case of monotonic algebras. Example 10.6 demonstrates how to apply **CharSFA** and **InferSFA** in order to learn an SFA over the algebra $\mathcal{A}_{\mathbb{N}}$.

Proposition 10.2. *There exist functions **Concretize** $_{\mathcal{A}_m}$ and **Generalize** $_{\mathcal{A}_m}$ for a monotonic Boolean algebra \mathcal{A}_m , satisfying the criteria of Theorem 9.3.*

Proof. Let \mathbb{D} be the domain of \mathcal{A}_m . We provide the functions **Concretize** $_{\mathcal{A}_m}$ and **Generalize** $_{\mathcal{A}_m}$ and prove that the criteria of Theorem 9.3 hold for them. For ease of presentation, for the function **Concretize** $_{\mathcal{A}_m}$ we consider basic predicates. Note that for monotonic algebras,




FIGURE 3. Visualisation of Example 10.3

basic predicates are in fact intervals, as a conjunction of intervals is an interval. We can assume all predicates are basic since, as we show in Lemma 4.3, for monotonic algebras the transformation from a general formula to a DNF formula of basic predicates is linear. Then, each basic predicate in the formula corresponds to a different predicate in the predicate partition. The definitions of $\text{Concretize}_{\mathcal{A}_m}$ and $\text{Generalize}_{\mathcal{A}_m}$ are generalizations of the functions $\text{Concretize}_{\mathcal{A}_{\mathbb{N}}}$ and $\text{Generalize}_{\mathcal{A}_{\mathbb{N}}}$ given in Example 8.3. We define $\text{Concretize}_{\mathcal{A}_m}(\langle \psi_1, \dots, \psi_m \rangle) = \langle \Gamma_1, \dots, \Gamma_m \rangle$ where we set $\Gamma_i = \{ \min\{d \in \mathbb{D} \mid d \in \llbracket \psi_i \rrbracket\} \}$ for $1 \leq i \leq m$. Since \mathcal{A}_m is monotonic, Γ_i is well-defined and contains a single element, thus $\text{Concretize}_{\mathcal{A}_m}$ is an efficient concretizing function.

We define $\text{Generalize}_{\mathcal{A}_m}(\langle \Gamma_1, \dots, \Gamma_m \rangle) = \langle \psi_1, \dots, \psi_m \rangle$, where ψ_i is defined as follows. Let $\Gamma = \bigcup_{1 \leq i \leq m} \Gamma_i$. First, for all $1 \leq i \leq m$ we set $\psi_i = \perp$. Then, we iteratively look for the minimal element $\gamma \in \Gamma$. Let i be such that $\gamma \in \Gamma_i$, and let γ' be the minimal element in Γ satisfying $\gamma' \notin \Gamma_i$. We then set $\psi_i = \psi_i \vee [\gamma, \gamma')$, and remove all elements $\gamma \leq \gamma'' < \gamma'$ from Γ . We repeat the process until for the found $\gamma \in \Gamma_j$, there is no $\gamma' > \gamma$ such that $\gamma' \notin \Gamma_j$. In that case, we define $\psi_j = \psi_j \vee [\gamma, d_\infty)$. Then, $\Gamma_i \subseteq \llbracket \psi_i \rrbracket$ and the predicates are disjoint, thus $\text{Generalize}_{\mathcal{A}_m}$ is an efficient generalizing function. See Example 10.3 and Figure 3.

Now, let $\langle \Gamma_1, \dots, \Gamma_m \rangle$ be the concrete partition obtained from $\text{Concretize}_{\mathcal{A}_m}$ when applied on the predicate partition $\langle \psi_1, \dots, \psi_m \rangle$. Assume further that the predicate partition $\langle \Gamma'_1, \dots, \Gamma'_m \rangle$ satisfies $\Gamma_i \subseteq \Gamma'_i \subseteq \llbracket \psi_i \rrbracket$ for $1 \leq i \leq m$. In particular, $\min(\Gamma'_i) = \min(\Gamma_i)$, since Γ_i contains the minimal elements in $\llbracket \psi_i \rrbracket$, and $\Gamma_i \subseteq \Gamma'_i \subseteq \llbracket \psi_i \rrbracket$. Thus applying $\text{Generalize}_{\mathcal{A}_m}$ will result in the same interval, satisfying the criterion of Theorem 9.3. \square

Example 10.3. Let $\Gamma_1 = \{0, 100, 400, 500\}$ and $\Gamma_2 = \{150, 200\}$ over the algebra $\mathcal{A}_{\mathbb{N}}$ with domain $\mathbb{N} \cup \{\infty\}$, see Figure 3a. Then, $\text{Generalize}_{\mathcal{A}_{\mathbb{N}}}$ sets $\Gamma = \{0, 100, 150, 200, 400, 500\}$, and finds the minimal element in Γ , which is 0. Since $0 \in \Gamma_1$, it then looks for the minimal element $\gamma \in \Gamma$ such that $\gamma \notin \Gamma_1$, and finds $150 \in \Gamma_2$. Therefore $\psi_1 = [0, 150)$ and we remove from Γ all elements in $[0, 150)$, that is, we remove 0 and 100, and we have $\Gamma = \{150, 200, 400, 500\}$. Next, it finds the minimal element in the updated Γ , which is 150 and is in Γ_2 . The minimal element in $\Gamma \setminus \Gamma_2$ is 400. Then, ψ_2 is set to be $\psi_2 = [150, 400)$ and $\Gamma = \{400, 500\}$. Now, $\psi_1 = [0, 150) \vee [400, \infty)$ since $400 \in \Gamma_1$ and there is no greater element that is not in Γ_1 (Figure 3b).

To show that any class of SFAs $\mathbb{M}_{\mathcal{A}_m}$ over a monotonic algebra \mathcal{A}_m is efficiently identifiable, we define in Algorithm 3 an algorithm that implements a decontaminating function

Input: set \mathcal{S} over alphabet Σ
Output: set \mathcal{S}' over alphabet Σ'

```

1 function DECONTAMINATE $_{\mathbb{M}_{\mathcal{A}_m}}$ ( $\mathcal{S}$ )
2    $A_w := \{\epsilon\}$ ,  $\Sigma' := \{d_{-\infty}\}$ ,  $\sigma_{max} := d_{-\infty}$ 
3   repeat
4     for all  $u \in A_w$ , by lexicographic order do
5       for all  $\sigma \in \Sigma$ , by lexicographic order do
6         if  $\sigma > \sigma_{max}$  and  $u\sigma \not\sim_{\mathcal{S}} u\sigma_{max}$  then
7           if  $\forall \sigma'. \sigma_{max} < \sigma' < \sigma : u\sigma' \sim_{\mathcal{S}} u\sigma_{max}$  then
8              $\Sigma' := \Sigma' \cup \{\sigma\}$ 
9             if  $\forall u' \in A_w. u\sigma \not\sim_{\mathcal{S}} u'$  then  $A_w := A_w \cup \{u\sigma\}$ 
10             $\sigma_{max} := \sigma$ 
11           $\sigma_{max} := d_{-\infty}$ 
12    until  $\Sigma'$  is remained unchanged
13  return  $\mathcal{S}' := \mathcal{S} \cap \Sigma'^*$ 

```

ALGORITHM 3. Decontaminate $_{\mathbb{M}_{\mathcal{A}_m}}$ – finding the necessary letters for a characteristic sample

Decontaminate $_{\mathbb{M}_{\mathcal{A}_m}}$, that fulfills the requirements of Theorem 9.3. Loosely speaking, the idea of the algorithm is to simultaneously collect elements into two sets A_w and Σ' such that A_w will consist of the minimal representative according to the lexicographic order of each equivalence class in $\sim_{\mathcal{S}}$ and Σ' will consist of minimal letters aiding to distinguish these words. When this process terminates the algorithm returns the subset of words in the sample that consist of only letters in Σ' .

Lemma 10.4. *Assume the input to Decontaminate $_{\mathbb{M}_{\mathcal{A}_m}}$ is \mathcal{S} , where $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{M}}$ for some $\mathcal{M} \in \mathbb{M}_{\mathcal{A}_m}$ such that $\mathcal{S}_{\mathcal{M}} = \mathbf{CharDFA}(\mathbf{Concretize}_{\mathbb{M}_{\mathcal{A}_m}}(\mathcal{M}))$, and $\mathcal{D}_{\mathcal{M}} = \mathbf{Concretize}_{\mathbb{M}_{\mathcal{A}_m}}(\mathcal{M})$ is over the alphabet $\Gamma_{\mathcal{M}}$. Then for Σ' constructed by Decontaminate $_{\mathbb{M}_{\mathcal{A}_m}}$ (Algorithm 3) it holds that $\Sigma' = \Gamma_{\mathcal{M}}$.*

Proof. Let $\mathcal{M} = (\mathcal{A}, Q, q_l, F, \Delta_{\mathcal{M}})$, $\mathcal{D}_{\mathcal{M}} = \mathbf{Concretize}_{\mathbb{M}_{\mathcal{A}_m}}(\mathcal{M})$ and $\mathcal{S}_{\mathcal{M}} = \mathbf{CharDFA}(\mathcal{D}_{\mathcal{M}})$. Assume $\mathcal{D}_{\mathcal{M}} = (\Gamma_{\mathcal{M}}, Q, q_l, F, \Delta_{\mathcal{D}})$. We inductively show that for Decontaminate $_{\mathbb{M}_{\mathcal{A}_m}}$ given in Algorithm 3, if its input \mathcal{S} satisfies $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{M}}$ then the set A_w is exactly the set of all lex-access words (lexicographically smallest access words) of states in $\mathcal{D}_{\mathcal{M}}$ and that $\Sigma' = \Gamma_{\mathcal{M}}$ (where $\Gamma_{\mathcal{M}}$ is the alphabet of $\mathcal{D}_{\mathcal{M}}$).

First, we show that every $u \in A_w$ is a lex-access word and that $\Sigma' \subseteq \Gamma_{\mathcal{M}}$.

For the base case, we consider $A_w = \{\epsilon\}$ and $\Sigma' = \{d_{-\infty}\}$. From item 1 of Theorem 7.5, we can assume access words are minimal according to the lexicographic order. Thus, $\epsilon \in A_w$ is indeed a lex-access word (of the state q_l). For $d_{-\infty} \in \Sigma'$, it holds that $\Gamma_{\mathcal{M}}$ contains the minimal element of \mathbb{D} since it contains all concretizations of intervals, the SFA is complete and Concretize $_{\mathbb{M}_{\mathcal{A}_m}}$ returns the minimal element of each interval. Therefore $d_{-\infty} \in \Gamma_{\mathcal{M}}$.

For the induction step, assume that A_w contains only lex-access words and that the current Σ' is a subset of $\Gamma_{\mathcal{M}}$. Then, when considering $u \in A_w$ in line 4, it holds that u is a lex-access word of some state q . Then, σ is added to Σ' only if $u\sigma \not\sim_{\mathcal{S}} ud_{-\infty}$. Since \mathcal{S} agrees with \mathcal{M} , it holds that $\Delta_{\mathcal{M}}(q_l, u\sigma) \neq \Delta_{\mathcal{M}}(q_l, ud_{-\infty})$ and σ is a minimal element with that property. Then, σ must be a minimal element of an interval labeling an outgoing transition from q , therefore is in $\Gamma_{\mathcal{M}}$. Inductively this holds for all elements added to Σ' in the current

iteration. This proves that $\Sigma' \subseteq \Gamma_{\mathcal{M}}$. Assume now that A_w contains only lex-access words and let $u\sigma$ be a word added to A_w in line 9. Thus, for all $u' \in A_w$ it holds that $u\sigma \not\sim_{\mathcal{S}} u'$.

Claim. In this setting, $u\sigma \not\sim_{\mathcal{S}} u'$ implies $u\sigma \not\sim_{\mathcal{S}_{\mathcal{M}}} u'$.

Proof. Since we assume all words already in A_w are lex-access words, then in particular u is a lex-access word. In addition, $\sigma \in \Sigma'$ and thus $\sigma \in \Gamma_{\mathcal{M}}$. Since $u\sigma \not\sim_{\mathcal{S}} u'$ there is $z \in \Sigma^*$ such that $\langle u\sigma z, b_1 \rangle, \langle u'z, b_2 \rangle \in \mathcal{S}$ and $b_1 \neq b_2$. Since \mathcal{S} agrees with \mathcal{M} it holds that $\Delta_{\mathcal{M}}(q_l, u\sigma) \neq \Delta_{\mathcal{M}}(q_l, u')$. Now, $u\sigma$ and u' are all in $\Gamma_{\mathcal{M}}$ since from line 9 we have $A_w \subseteq \Sigma'^*$, and thus $\Delta_{\mathcal{D}_{\mathcal{M}}}(q_l, u\sigma) \neq \Delta_{\mathcal{D}_{\mathcal{M}}}(q_l, u')$, and from item 3 of Theorem 7.5 it holds that $u\sigma \not\sim_{\mathcal{S}_{\mathcal{M}}} u'$. This proves the claim.

Then, for all $u' \in A_w$ we have $\Delta_{\mathcal{D}}(q_l, u\sigma) \neq \Delta_{\mathcal{D}}(q_l, u')$. Then, since we traverse words and letters in lexicographic order, $u\sigma$ is a lex-access word for $\Delta_{\mathcal{D}}(q_l, u\sigma)$. This concludes the first direction.

For the second direction, we show that every lex-access word is in A_w and that $\Gamma_{\mathcal{M}} \subseteq \Sigma'$.

Let $u\sigma$ be a lex-access word (for ϵ it holds that $\epsilon \in A_w$). For all lex-access words u' found in previous iterations it holds that $u\sigma \not\sim_{\mathcal{S}_{\mathcal{M}}} u'$ from item 2 of Theorem 7.5, and thus $u\sigma \not\sim_{\mathcal{S}} u'$ since $\mathcal{S}_{\mathcal{M}} \subseteq \mathcal{S}$. Then, $u\sigma$ satisfies the condition of line 9 and is added to A_w .




To prove $\Gamma_{\mathcal{M}} \subseteq \Sigma'$, let $\sigma \in \Gamma_{\mathcal{M}}$. From the construction of $\text{Concretize}_{\mathcal{A}_m}$ it holds that σ is the left endpoint of some interval that is an outgoing transition from q_l . Then, indeed σ is found in the first iteration of line 4. Inductively, let σ label an outgoing transition of q for some $q \in Q$ and let u_q be the lex-access word of q . Since A_w contains all lex-access words, it holds that $u_q \in A_w$, and then the outgoing transitions of q will be considered in some following iteration. Thus, all minimal letters indicating new intervals are added to Σ' and we have that $\Gamma_{\mathcal{M}} \subseteq \Sigma'$. We conclude that $\Sigma' = \Gamma_{\mathcal{M}}$. \square

Proposition 10.5. *The sufficient condition of Theorem 9.3 holds for the class $\mathbb{M}_{\mathcal{A}_m}$ of SFAs over a monotonic Boolean algebra \mathcal{A}_m .*

Proof. In Proposition 10.2 we have shown that there exist functions $\text{Concretize}_{\mathcal{A}_m}$ and $\text{Generalize}_{\mathcal{A}_m}$ for a monotonic Boolean algebra \mathcal{A}_m , satisfying the criteria of Theorem 9.3. It is left to show that $\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}_m}}$ is an efficient decontaminating function. Assume that $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{M}}$ where $\mathcal{S}_{\mathcal{M}} = \text{CharDFA}(\text{Concretize}_{\mathbb{M}_{\mathcal{A}_m}}(\mathcal{M}))$, and $\text{Concretize}_{\mathbb{M}_{\mathcal{A}_m}}(\mathcal{M})$ is over alphabet $\Gamma_{\mathcal{M}}$. In Lemma 10.4 we showed that under these assumptions it holds that the alphabet Σ' of the returned sample \mathcal{S}' is $\Gamma_{\mathcal{M}}$. Then, for the set \mathcal{S}' returned in line 13 (Algorithm 3) it holds that $\mathcal{S}' = \mathcal{S} \cap \Gamma_{\mathcal{M}}^*$. Since $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{M}}$ and $\Gamma_{\mathcal{M}}^* \supseteq \mathcal{S}_{\mathcal{M}}$, it holds that $\mathcal{S}' \supseteq \mathcal{S}_{\mathcal{M}}$ and \mathcal{S}' is defined over the alphabet $\Gamma_{\mathcal{M}}$. Therefore, $\text{Decontaminate}_{\mathbb{M}_{\mathcal{A}_m}}$ is a decontaminating function. In addition, it runs in time polynomial in the size of \mathcal{S} , thus the conditions of Theorem 9.3 are met. \square

We now provide an example for efficient identification of SFAs over the interval algebra (Example 10.6), and an example for the need of a decontaminating function (Example 10.7).

Example 10.6. Continuing Example 8.3, let \mathcal{M} be the SFA from Figure 1 and consider the class $\mathbb{M}_{\mathcal{A}_{\mathbb{N}}}$ of SFAs over the interval algebra. Algorithm **CharSFA** computes the set $\Gamma_{\mathcal{M}}$ using the function $\text{Concretize}_{\mathbb{M}_{\mathcal{A}_{\mathbb{N}}}}$ given in Example 8.3. That is, for the predicates labeling outgoing transitions from q_0 we have $\text{Concretize}_{\mathcal{A}_{\mathbb{N}}}(\langle [0, 100), [100, \infty) \rangle) = \langle \{0\}, \{100\} \rangle$; and for outgoing transitions from q_1 , it holds that $\text{Concretize}_{\mathcal{A}_{\mathbb{N}}}(\langle [0, 200), [200, \infty) \rangle) = \langle \{0\}, \{200\} \rangle$. Then, $\Gamma_{\mathcal{M}} = \{0, 100, 200\}$, and **CharSFA** constructs the DFA over $\Gamma_{\mathcal{M}}$ where concrete transitions agree with symbolic transitions of the original SFA. See Figure 4 for the resulting

FIGURE 4. The DFA $\mathcal{D}_{\mathcal{M}}$ constructed in **CharSFA**FIGURE 5. The SFA \mathcal{M}_1 FIGURE 6. The prefix-tree automaton for the set \mathcal{S}' of Example 10.7

DFA. Then, using **CharDFA**, it returns the sample set:


$$\mathcal{S}_{\mathcal{M}} = \{\langle \epsilon, 0 \rangle, \langle 0, 1 \rangle, \langle 100, 0 \rangle, \langle 200, 0 \rangle, \langle 0 \cdot 0, 1 \rangle, \langle 0 \cdot 100, 1 \rangle, \langle 0 \cdot 200, 0 \rangle\}.$$

Now, assume algorithm **InferSFA** is given the set $\mathcal{S} = \mathcal{S}_{\mathcal{M}} \cup \{\langle 150, 0 \rangle\}$ over the alphabet $\Sigma = \{0, 100, 150, 200\}$. The algorithm applies $\text{Decontaminate}_{\mathbb{M}, \mathfrak{q}_{\mathbb{N}}}$ that generates the set \mathcal{S}' over $\Gamma_{\mathcal{M}}$. To do so, it first finds the set $\Gamma_{\mathcal{M}}$ of all elements that are a minimal left point of some interval, and then chooses from \mathcal{S} the words over $\Gamma_{\mathcal{M}}$. It does so as follows. First, note that $100 \sim_{\mathcal{S}} 150$, $100 \sim_{\mathcal{S}} 200$ and $150 \sim_{\mathcal{S}} 200$, while $0 \not\sim_{\mathcal{S}} 100, 150, 200$. Since 0 is the minimal element it has to be in $\Gamma_{\mathcal{M}}$; and since 100 is the minimal element that is not equivalent to 0 it has to define a new interval and thus is in $\Gamma_{\mathcal{M}}$ as well. Next, we consider suffixes of words over $\{0, 100\}$. These are $0 \cdot 0$ and $0 \cdot 100$ which are equivalent, and $0 \cdot 200$ which is not equivalent to the former. Since 100 is equivalent to 0 it does not define a new interval now, but 200 does as it is the minimal (and only) element that is not equivalent to 0 when considering suffixes of 0. Then, we deduce that $\Gamma_{\mathcal{M}} = \{0, 100, 200\}$ and thus $\mathcal{S}' = \{\langle \epsilon, 0 \rangle, \langle 0, 1 \rangle, \langle 100, 0 \rangle, \langle 200, 0 \rangle, \langle 0 \cdot 0, 1 \rangle, \langle 0 \cdot 100, 1 \rangle, \langle 0 \cdot 200, 0 \rangle\}$.

Now Algorithm **InferSFA** is applied to the set \mathcal{S}' and the resulting DFA would be the DFA $\mathcal{D}_{\mathcal{M}}$ of Figure 4. Then it applies $\text{Generalize}_{\mathbb{M}, \mathfrak{q}_{\mathbb{N}}}$ described in Example 8.3 and the result will be the original SFA of Figure 1. That is, for outgoing transitions of q_0 it applies $\text{Generalize}_{\mathfrak{q}_{\mathbb{N}}}(\langle \{0\}, \{100, 200\} \rangle) = \langle [0, 100], [100, \infty) \rangle$ and for outgoing transitions of q_1 it applies $\text{Generalize}_{\mathfrak{q}_{\mathbb{N}}}(\langle \{0, 100\}, \{200\} \rangle) = \langle [0, 200], [200, \infty) \rangle$ and uses these predicates to annotate the corresponding transitions in the SFA.

Example 10.7. Consider the SFA \mathcal{M}_1 of Figure 5. Applying $\text{CharDFA}(\text{Concretize}_{\mathbb{M}, \mathfrak{q}_{\mathbb{N}}}(\mathcal{M}_1))$ results in the following \mathcal{S} .

$$\mathcal{S} = \{\langle \epsilon, 0 \rangle, \langle 0, 1 \rangle, \langle 100, 1 \rangle, \langle 0 \cdot 0, 1 \rangle, \langle 0 \cdot 100, 0 \rangle, \langle 100 \cdot 100, 1 \rangle, \langle 100 \cdot 0, 0 \rangle, \langle 0 \cdot 100 \cdot 0, 0 \rangle\}$$

FIGURE 7. Two DFAs that are consistent with the set \mathcal{S}' of Example 10.7.

Now, let $\mathcal{S}' = \mathcal{S} \cup \{\langle 150, 1 \rangle, \langle 150 \cdot 250, 1 \rangle\}$. Note that \mathcal{S}' is consistent with $\mathcal{L}(\mathcal{M}_1)$. When trying to learn an SFA from \mathcal{S}' and applying **InferDFA**(\mathcal{S}'), the algorithm cannot distinguish between the words 150 and 100, as well as between 0 and 150. The same holds also for $0 \cdot 0$ vs. $150 \cdot 250$, and $100 \cdot 100$ vs. $150 \cdot 250$.

Figure 7 presents two DFAs that are both consistent with the set \mathcal{S}' . Since \mathcal{S}' does not contain any characteristic sample for DFAs over the alphabet $\{0, 100, 150, 250\}$, **InferDFA** concludes that \mathcal{S}' does not subsume any characteristic sample, and returns the prefix-tree automaton, given in Figure 6.

This example illustrates that **InferSFA** cannot classify 150 and $150 \cdot 250$ by only applying **InferDFA**, without reasoning about the predicates of the algebra. To this end we provide the function **Decontaminate** $_{\mathbb{M}_{\mathcal{A}}}$, which is able, in the case of a monotonic algebra, to find which letters are the ones that should be used to define new predicates.

11. NEGATIVE RESULT

The result of Theorem 10.1 does not extend to the non-monotonic case, as stated in Theorem 11.1 regarding SFAs over the general propositional algebra. Let $\mathbb{D}_{\mathbb{B}} = \{\mathbb{B}^k\}_{k \in \mathbb{N}}$. Recall that $\mathbb{B} = \{0, 1\}$ and \mathbb{B}^k is the set of all valuations of k atomic propositions. Let $\mathbb{P}_{\mathbb{B}} = \{\mathbb{P}_{\mathbb{B}^k}\}_{k \in \mathbb{N}}$ where $\mathbb{P}_{\mathbb{B}^k}$ is the set of predicates over at most k atomic propositions. Let $\mathcal{A}_{\mathbb{B}}$ be the Boolean algebra defined over the discrete domain $\mathbb{D}_{\mathbb{B}}$ and the set of predicates $\mathbb{P}_{\mathbb{B}}$, and the usual operators \vee , \wedge and \neg . Let $\mathbb{M}_{\mathcal{A}_{\mathbb{B}}}$ be the class of SFAs over the Boolean algebra $\mathcal{A}_{\mathbb{B}}$. We show that unless $P = NP$, this class of SFAs is not efficiently identifiable.¹¹

Theorem 11.1. *The class $\mathbb{M}_{\mathcal{A}_{\mathbb{B}}}$ is not efficiently identifiable unless $P = NP$.*

Proof. We show that there is no pair of efficient dyadic concretizing and generalizing functions $f_c : \Pi_{\text{pred}}(\mathbb{P}_{\mathbb{B}}) \rightarrow \Pi_{\text{conc}}(\mathbb{D}_{\mathbb{B}})$ and $f_g : \Pi_{\text{conc}}(\mathbb{D}_{\mathbb{B}}) \rightarrow \Pi_{\text{pred}}(\mathbb{P}_{\mathbb{B}})$ unless $P = NP$. From Theorem 8.2 it follows that $\mathbb{M}_{\mathbb{B}}$ is not efficiently identifiable unless $P = NP$.

¹¹This result may be contrasted with [AD18] who provide a positive learnability result regarding SFAs over the OBDDs algebra. The result of [AD18] is with respect to query learning, while Theorem 11.1 concerns efficient identifiability in the limit. As we discuss in section 13, one cannot derive efficient learnability from a positive result in the query learning setting. Moreover, in Theorem 11.1 (as well as in Corollary 12.3 and the discussion in section 12) we refer to efficient learnability with respect the propositional algebra as defined in subsection 2.1 where the size is measured with respect to the number of atomic propositions, while [AD18] refer to the size of the SFAs in which the predicates are OBDDs, whose size is measured by the number of nodes in the OBDD. However, the number of nodes in an OBDD can be exponential in the number of atomic propositions. Therefore, our result has no conflict with the result of [AD18].

Assume towards contradiction that such a pair of functions exist. We provide a polynomial time algorithm A_{SAT} for SAT. On a predicate φ , the algorithm A_{SAT} invokes $f_c(\langle\varphi, \neg\varphi\rangle)$. Suppose the returned concrete partition is $\langle\Gamma_1, \Gamma_2\rangle$. Then A_{SAT} returns “true” if and only if $\Gamma_1 \neq \emptyset$. Correctness follows from the fact that if there exists a system of characteristic samples for $\mathbb{P}_{\mathbb{B}}$ then the set of positive examples associated with a satisfiable predicate φ must be non-empty, as otherwise f_g cannot distinguish φ from \perp . \square

12. QUERY LEARNING

The paradigm of *query learning* stipulates that the *learner* can interact with an *oracle* (*teacher*) by asking it several types of allowed queries. In this section we consider these queries to be *membership queries* (MQ) and *equivalence queries* (EQ). We say that a class \mathbb{M} of automata is efficiently learnable using MQs / EQs / both MQs and EQs if there is an algorithm that for every language L with a representation in \mathbb{M} asks a polynomial number of MQs / EQs / both MQs and EQs and outputs an automaton in \mathbb{M} that is polynomial in the minimal representation of L in \mathbb{M} .

Angluin showed, on the negative side, that regular languages cannot be efficiently learned (in the exact model) from only MQs [Ang81] or only EQs [Ang90]. On the positive side, she showed that regular languages, represented as DFAs, can be efficiently learned using both MQs and EQs [Ang87a]. The celebrated algorithm, termed \mathbf{L}^* , was extended to learning many other classes of languages and representations, e.g., [Sak90, BV96, AV10, BHKL09, AEF15, MP95, AF16, NFZ21]. See the survey [Fis18] for more references.

In particular, an extension of \mathbf{L}^* , termed \mathbf{MAT}^* , to learn SFAs was provided in [AD18] which proved that SFAs over an algebra \mathcal{A} can be efficiently learned using \mathbf{MAT}^* if and only if the underlying algebra is efficiently learnable, and the size of disjunctions of k predicates does not grow exponentially in k .¹² From this it was concluded that SFAs over the following underlying algebras are efficiently learnable: Boolean algebras over finite domains, equality algebra, tree automata algebra, and SFAs algebra. Efficient learning of SFAs over a monotonic algebra using MQs and EQs was established in [CDYS17], which improved the results of [MM14, MM17] by using a binary search instead of a helpful teacher.

The result of [AD18] provides means to establish new positive results on learning classes of SFAs using MQs and EQs, but it does not provide means for obtaining negative results for query learning of SFAs using MQs and EQs. We strengthen this result by providing a learnability result that is independent of the use of a specific learning algorithm. In particular, we show that efficient learnability of a Boolean algebra \mathcal{A} using MQs and EQs is a necessary condition for the learnability of a class of SFAs over \mathcal{A} , as we state in Theorem 12.1.

Theorem 12.1. *A class of SFAs \mathbb{M} over a Boolean algebra \mathcal{A} , that contains all basic SFAs over \mathcal{A} , is polynomially learnable using MQs and EQs, only if \mathcal{A} is polynomially learnable using MQs and EQs.*

Proof. Assume that \mathbb{M} is polynomially learnable using MQs and EQs, using an algorithm $\mathbf{Q}_{\mathbb{M}}$. We show that there exists a polynomial learning algorithm $\mathbf{Q}_{\mathcal{A}}$ for the algebra \mathcal{A} using MQs and EQs. The algorithm $\mathbf{Q}_{\mathcal{A}}$ uses $\mathbf{Q}_{\mathbb{M}}$ as a subroutine, and behaves as a teacher for $\mathbf{Q}_{\mathbb{M}}$. Whenever $\mathbf{Q}_{\mathbb{M}}$ asks an \mathbb{M} -MQ on word $\gamma_1 \dots \gamma_k$, if $k > 1$ then $\mathbf{Q}_{\mathcal{A}}$ answers “no”. If $k = 1$ then the \mathbb{M} -MQ is essentially an \mathcal{A} -MQ, thus $\mathbf{Q}_{\mathcal{A}}$ issues this query and passes the answer

¹²As is the case, for instance, in the OBDD (Ordered Binary Decisions Diagrams) algebra [Bry86].

to $\mathbf{Q}_{\mathbb{M}}$. Whenever $\mathbf{Q}_{\mathbb{M}}$ asks a \mathbb{M} -EQ on SFA \mathcal{M} , if \mathcal{M} is not of the form \mathcal{M}_ψ for some ψ (as defined in Definition 6.1) then $\mathbf{Q}_{\mathcal{A}}$ answers “no” to the \mathbb{M} -EQ and returns some word $w \in \mathcal{L}(\mathcal{M})$ s.t. $|w| > 1$ and w was not provided before, as a counterexample. To this aim it can record the largest counterexample given so far (according to the lexicographic order) and return the next one in this order. Otherwise (if the SFA is of the form \mathcal{M}_ψ for some ψ) $\mathbf{Q}_{\mathcal{A}}$ asks an \mathcal{A} -EQ on ψ . If the answer is “yes” then $\mathbf{Q}_{\mathcal{A}}$ terminates and returns ψ as the result of the learning algorithm; if the answer to the \mathcal{A} -EQ on ψ is “no”, then the provided counterexample $\langle \gamma, b_\gamma \rangle$ is passed back to $\mathbf{Q}_{\mathbb{M}}$ together with the answer “no” to the \mathbb{M} -EQ. It is easy to verify that $\mathbf{Q}_{\mathcal{A}}$ terminates correctly in polynomial time. \square

From Theorem 12.1 we derive what we believe to be the first negative result on learning SFAs from MQs and EQs, as we show that SFAs over $\mathcal{A}_{\mathbb{B}_k}$, the propositional algebra over k variables, are not polynomially learnable using MQs and EQs. Polynomiality is measured with respect to the parameters $\langle n, m, l \rangle$ representing the size of the SFA and the number k of atomic propositions. Note that the algebra $\mathcal{A}_{\mathbb{B}_k}$ is a restriction of the algebra $\mathcal{A}_{\mathbb{B}}$ considered in section 11 and therefore implies a negative result also with regard to the algebra $\mathcal{A}_{\mathbb{B}}$ considered there.

We achieve this by showing that no learning algorithm \mathbf{A} for the propositional algebra using MQs and EQs can do better than asking 2^k MQs/EQs, where k is the number of atomic propositions.¹³ We assume the learning algorithm is *sound*, that is, if S_i^+ and S_i^- are the sets of positive and negative examples observed by the algorithm up to stage i , then at stage $i+1$ the algorithm will not ask a MQ for a word in $S_i^+ \cup S_i^-$ or an EQ for an automaton that rejects a word in S_i^+ or accepts a word in S_i^- .

Proposition 12.2. *Let \mathbf{A} be a sound learning algorithm for the propositional algebra over \mathbb{B}^k . There exists a target predicate ψ of size k , for which \mathbf{A} will be forced to ask at least $2^k - 1$ queries (either MQ or EQ).*

Proof. Since \mathbf{A} is sound, at stage $i+1$ we have $S_{i+1}^+ \supseteq S_i^+$ and $S_{i+1}^- \supseteq S_i^-$ and at least one inclusion is strict. Since the size of the concrete alphabet is 2^k , for every round $i < 2^k$, an adversarial teacher can answer both MQs and EQs negatively. In the case of an EQ there must be an element in $\mathbb{B}^k \setminus (S_i^- \cup S_i^+)$ with which the provided automaton disagrees. The adversary will return one such element as a counterexample. This forces \mathbf{A} to ask at least $2^k - 1$ queries. Note that for any element v in \mathbb{B}^k there exists a predicate φ_v of size k such that $\llbracket \varphi_v \rrbracket = \{v\}$. \square

Corollary 12.3. *SFAs over the propositional algebra $\mathcal{A}_{\mathbb{B}_k}$ with k propositions cannot be learned in $\text{poly}(k)$ time using MQs and EQs.*

The propositional algebra $\mathcal{A}_{\mathbb{B}_k}$ is a special case of the n -dimensional boxes algebra. Learning n -dimensional boxes was studied using MQs and EQs [GGM94, BGGM98, BK98], as well as in the PAC setting [BK00]. The algorithms presented in [GGM94, BGGM98, BK98, BK00] are mostly exponential in n . Alternatively, [GGM94, BGGM98] suggest algorithms that are exponential in the number of boxes in the union. In [BK98] a linear query learning algorithm for unions of disjoint boxes is presented. Since n -dimensional boxes subsume the propositional algebra, Corollary 12.3 implies the following.

¹³In [Nak00] Boolean formulas represented using OBDDs are claimed to be polynomially learnable with MQs and EQs. However, [Nak00] measures the size of an OBDD by its number of nodes, which can be exponential in the number of propositions.

Corollary 12.4. *The class of SFAs over the n -dimensional boxes algebra cannot be learned in $\text{poly}(n)$ time using MQs and EQs.*

13. DISCUSSION

We examined the question of learnability of a class of SFAs over certain algebras where the main focus of our study is on passive learning. We provided a necessary condition for identification of SFAs in the limit using polynomial time and data, as well as a necessary condition for efficient learning of SFAs using membership and equivalence queries. We note that a positive result on learning deterministic SFAs using MQs and EQs implies a positive result for identification of deterministic SFAs in the limit using polynomial time and data. The latter follows because a systematic set of characteristic samples $\{S_L\}_{L \in \mathbb{L}}$ for a class of languages \mathbb{L} may be obtained by collecting the words observed by the query learner when learning L , and given the SFA is deterministic, the words in the sample can be restricted to ones of polynomial size, thus the size of the sample is polynomial in the size of the SFA.¹⁴ However, it does not imply a positive result regarding the stronger notion of efficient identifiability, as the latter requires the set to be also constructed efficiently, and the complexity analysis for query learning does not include the complexity of the teacher in computing queries, e.g., deciding equivalence and in constructing counterexamples. We thus provided a sufficient condition for efficient identification of a class of SFAs, and showed that the class of SFAs over any monotonic algebra satisfies these conditions.

We hope that these sufficient or necessary conditions will help to obtain more positive and negative results for learning of SFAs, and spark an interest in investigating characteristic samples in other automata models used in verification.

ACKNOWLEDGMENTS

This research was partially supported by the Israel Science Foundation grant 2507/21. S. Zilles acknowledges financial support by the Natural Sciences and Engineering Research Council (NSERC) of Canada, both through the Canada Research Chairs program and the Discovery Grants program. She was also supported through the Canada CIFAR AI Chairs program as an Affiliate Chair with the Alberta Machine Intelligence Institute (Amii).

REFERENCES

- [AD18] George Argyros and Loris D’Antoni. The learnability of symbolic automata. In Hana Chockler and Georg Weissenbacher, editors, *Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I*, volume 10981 of *Lecture Notes in Computer Science*, pages 427–445. Springer, 2018. doi:10.1007/978-3-319-96145-3_23.
- [AEF15] Dana Angluin, Sarah Eisenstat, and Dana Fisman. Learning regular languages via alternating automata. In Qiang Yang and Michael J. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 3308–3314. AAAI Press, 2015. URL: <http://ijcai.org/Abstract/15/466>.

¹⁴For non-deterministic SFAs, as is the case for NFAs, it might be the case that some exponentially long words are required to learn the language [dIH97].

- [AF16] Dana Angluin and Dana Fisman. Learning regular omega languages. *Theor. Comput. Sci.*, 650:57–72, 2016. doi:10.1016/j.tcs.2016.07.031.
- [AFH99] Rajeev Alur, Limor Fix, and Thomas A. Henzinger. Event-clock automata: A determinizable class of timed automata. *Theor. Comput. Sci.*, 211(1-2):253–273, 1999. doi:10.1016/S0304-3975(97)00173-4.
- [Ang81] Dana Angluin. A note on the number of queries needed to identify regular languages. *Inf. Control.*, 51(1):76–87, 1981. doi:10.1016/S0019-9958(81)90090-5.
- [Ang87a] Dana Angluin. Learning regular sets from queries and counterexamples. *Inf. Comput.*, 75(2):87–106, 1987. doi:10.1016/0890-5401(87)90052-6.
- [Ang87b] Dana Angluin. Queries and concept learning. *Mach. Learn.*, 2(4):319–342, 1987. doi:10.1007/BF00116828.
- [Ang90] Dana Angluin. Negative results for equivalence queries. *Mach. Learn.*, 5:121–150, 1990. doi:10.1007/BF00116034.
- [ASJ⁺16] George Argyros, Ioannis Stais, Suman Jana, Angelos D. Keromytis, and Aggelos Kiayias. SFADiff: Automated evasion attacks and fingerprinting using black-box differential automata learning. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 1690–1701. ACM, 2016. doi:10.1145/2976749.2978383.
- [ASKK16] George Argyros, Ioannis Stais, Aggelos Kiayias, and Angelos D. Keromytis. Back in black: Towards formal, black box analysis of sanitizers and filters. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016*, pages 91–109. IEEE Computer Society, 2016. doi:10.1109/SP.2016.14.
- [AV10] Fides Aarts and Frits W. Vaandrager. Learning I/O automata. In Paul Gastin and François Laroussinie, editors, *CONCUR 2010 - Concurrency Theory, 21th International Conference, CONCUR 2010, Paris, France, August 31-September 3, 2010. Proceedings*, volume 6269 of *Lecture Notes in Computer Science*, pages 71–85. Springer, 2010. doi:10.1007/978-3-642-15375-4_6.
- [BGGM98] Nader H. Bshouty, Paul W. Goldberg, Sally A. Goldman, and H. David Mathias. Exact learning of discretized geometric concepts. *SIAM J. Comput.*, 28(2):674–699, 1998. doi:10.1137/S0097539794274246.
- [BGJ⁺05] Therese Berg, Olga Grinchtein, Bengt Jonsson, Martin Leucker, Harald Raffelt, and Bernhard Steffen. On the correspondence between conformance testing and regular inference. In Maura Cerioli, editor, *Fundamental Approaches to Software Engineering, 8th International Conference, FASE 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*, volume 3442 of *Lecture Notes in Computer Science*, pages 175–189. Springer, 2005. doi:10.1007/978-3-540-31984-9_14.
- [BHKL09] Benedikt Bollig, Peter Habermehl, Carsten Kern, and Martin Leucker. Angluin-style learning of NFA. In Craig Boutilier, editor, *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, pages 1004–1009, 2009. URL: <http://ijcai.org/Proceedings/09/Papers/170.pdf>.
- [BK98] Amos Beimel and Eyal Kushilevitz. Learning boxes in high dimension. *Algorithmica*, 22(1/2):76–90, 1998. doi:10.1007/PL00013835.
- [BK00] Amos Beimel and Eyal Kushilevitz. Learning unions of high-dimensional boxes over the reals. *Inf. Process. Lett.*, 73(5-6):213–220, 2000. doi:10.1016/S0020-0190(00)00024-7.
- [BK08] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- [Bry86] Randal E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers*, 35(8):677–691, 1986. doi:10.1109/TC.1986.1676819.
- [BV96] Francesco Bergadano and Stefano Varricchio. Learning behaviors of automata from multiplicity and equivalence queries. *SIAM J. Comput.*, 25(6):1268–1280, 1996. doi:10.1137/S009753979326091X.
- [CDYS17] K. Chubachi, Diptarama, R. Yoshinaka, and A. Shinohara. Query learning of regular languages over large ordered alphabets. In *1st Workshop on Learning and Automata (LearnAut)*, 2017.
- [CGP01] Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. *Model checking, 1st Edition*. MIT Press, 2001. URL: <http://books.google.de/books?id=Nmc4wEaLXFEC>.

- [CHYS19] Kaizaburo Chubachi, Diptarama Hendrian, Ryo Yoshinaka, and Ayumi Shinohara. Query learning algorithm for residual symbolic finite automata. In Jérôme Leroux and Jean-François Raskin, editors, *Proceedings Tenth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2019, Bordeaux, France, 2-3rd September 2019*, volume 305 of *EPTCS*, pages 140–153, 2019. doi:10.4204/EPTCS.305.10.
- [CKKS20] Hana Chockler, Pascal Kesseli, Daniel Kroening, and Ofer Strichman. Learning the language of software errors. *J. Artif. Intell. Res.*, 67:881–903, 2020. doi:10.1613/jair.1.11798.
- [DD17] Samuel Drews and Loris D’Antoni. Learning symbolic automata. In Axel Legay and Tiziana Margaria, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Part I*, volume 10205 of *Lecture Notes in Computer Science*, pages 173–189, 2017. doi:10.1007/978-3-662-54577-5_10.
- [dlH97] Colin de la Higuera. Characteristic sets for polynomial grammatical inference. *Mach. Learn.*, 27(2):125–138, 1997. doi:10.1023/A:1007353007695.
- [DV14] Loris D’Antoni and Margus Veanes. Minimization of symbolic automata. In Suresh Jagannathan and Peter Sewell, editors, *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’14, San Diego, CA, USA, January 20-21, 2014*, pages 541–554. ACM, 2014. doi:10.1145/2535838.2535849.
- [DV16] Loris D’Antoni and Margus Veanes. Minimization of symbolic tree automata. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS ’16, New York, NY, USA, July 5-8, 2016*, pages 873–882. ACM, 2016. doi:10.1145/2933575.2933578.
- [DVL14] Loris D’Antoni, Margus Veanes, Benjamin Livshits, and David Molnar. Fast: a transducer-based language for tree manipulation. In Michael F. P. O’Boyle and Keshav Pingali, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’14, Edinburgh, United Kingdom - June 09 - 11, 2014*, pages 384–394. ACM, 2014. doi:10.1145/2594291.2594309.
- [FFZ22] Dana Fisman, Hadar Frenkel, and Sandra Zilles. Inferring symbolic automata. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14-19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPICs*, pages 21:1–21:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.CSL.2022.21.
- [FGPS20] Hadar Frenkel, Orna Grumberg, Corina S. Pasareanu, and Sarai Sheinvald. Assume, guarantee or repair. In Armin Biere and David Parker, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 26th International Conference, TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings, Part I*, volume 12078 of *Lecture Notes in Computer Science*, pages 211–227. Springer, 2020. doi:10.1007/978-3-030-45190-5_12.
- [FGPS22] Hadar Frenkel, Orna Grumberg, Corina S. Pasareanu, and Sarai Sheinvald. Assume, guarantee or repair: a regular framework for non regular properties. *Int. J. Softw. Tools Technol. Transf.*, 24(5):667–689, 2022. doi:10.1007/s10009-022-00669-9.
- [Fis18] Dana Fisman. Inferring regular languages and ω -languages. *J. Log. Algebraic Methods Program.*, 98:27–49, 2018. doi:10.1016/j.jlamp.2018.03.002.
- [GGM94] Paul W. Goldberg, Sally A. Goldman, and H. David Mathias. Learning unions of boxes with membership and equivalence queries. In Manfred K. Warmuth, editor, *Proceedings of the Seventh Annual ACM Conference on Computational Learning Theory, COLT 1994, New Brunswick, NJ, USA, July 12-15, 1994*, pages 198–207. ACM, 1994. doi:10.1145/180139.181102.
- [GJL10] Olga Grinchtein, Bengt Jonsson, and Martin Leucker. Learning of event-recording automata. *Theor. Comput. Sci.*, 411(47):4029–4054, 2010. doi:10.1016/j.tcs.2010.07.008.
- [GM96] Sally A. Goldman and H. David Mathias. Teaching a smarter learner. *J. Comput. Syst. Sci.*, 52(2):255–267, 1996. doi:10.1006/jcss.1996.0020.
- [Gol78] E. Mark Gold. Complexity of automaton identification from given data. *Inf. Control.*, 37(3):302–320, 1978. doi:10.1016/S0019-9958(78)90562-4.

- [HD17] Qinheping Hu and Loris D'Antoni. Automatic program inversion using symbolic transducers. In Albert Cohen and Martin T. Vechev, editors, *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, pages 376–389. ACM, 2017. doi:10.1145/3062341.3062345.
- [HLM⁺11] Pieter Hooimeijer, Benjamin Livshits, David Molnar, Prateek Saxena, and Margus Veanes. Fast and precise sanitizer analysis with BEK. In *20th USENIX Security Symposium, San Francisco, CA, USA, August 8-12, 2011, Proceedings*. USENIX Association, 2011. URL: http://static.usenix.org/events/sec11/tech/full_papers/Hooimeijer.pdf.
- [HSM11] Falk Howar, Bernhard Steffen, and Maik Merten. Automata learning with automated alphabet abstraction refinement. In Ranjit Jhala and David A. Schmidt, editors, *Verification, Model Checking, and Abstract Interpretation - 12th International Conference, VMCAI 2011, Austin, TX, USA, January 23-25, 2011. Proceedings*, volume 6538 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 2011. doi:10.1007/978-3-642-18275-4_19.
- [HV11] Pieter Hooimeijer and Margus Veanes. An evaluation of automata algorithms for string analysis. In Ranjit Jhala and David A. Schmidt, editors, *Verification, Model Checking, and Abstract Interpretation - 12th International Conference, VMCAI 2011, Austin, TX, USA, January 23-25, 2011. Proceedings*, volume 6538 of *Lecture Notes in Computer Science*, pages 248–262. Springer, 2011. doi:10.1007/978-3-642-18275-4_18.
- [KT14] Matthias Keil and Peter Thiemann. Symbolic solving of extended regular expression inequalities. In Venkatesh Raman and S. P. Suresh, editors, *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, volume 29 of *LIPICs*, pages 175–186. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2014. doi:10.4230/LIPICs.FSTTCS.2014.175.
- [MM14] Oded Maler and Irini-Eleftheria Mens. Learning regular languages over large alphabets. In Erika Ábrahám and Klaus Havelund, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014. Proceedings*, volume 8413 of *Lecture Notes in Computer Science*, pages 485–499. Springer, 2014. doi:10.1007/978-3-642-54862-8_41.
- [MM17] Oded Maler and Irini-Eleftheria Mens. A generic algorithm for learning symbolic automata from membership queries. In Luca Aceto, Giorgio Bacci, Giovanni Bacci, Anna Ingólfssdóttir, Axel Legay, and Radu Mardare, editors, *Models, Algorithms, Logics and Tools - Essays Dedicated to Kim Guldstrand Larsen on the Occasion of His 60th Birthday*, volume 10460 of *Lecture Notes in Computer Science*, pages 146–169. Springer, 2017. doi:10.1007/978-3-319-63121-9_8.
- [MP95] Oded Maler and Amir Pnueli. On the learnability of infinitary regular sets. *Inf. Comput.*, 118(2):316–326, 1995. doi:10.1006/inco.1995.1070.
- [MRA⁺17] Konstantinos Mamouras, Mukund Raghothaman, Rajeev Alur, Zachary G. Ives, and Sanjeev Khanna. StreamQRE: modular specification and efficient evaluation of quantitative queries over streaming data. In Albert Cohen and Martin T. Vechev, editors, *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, pages 693–708. ACM, 2017. doi:10.1145/3062341.3062369.
- [Myh57] J. Myhill. Finite automata and the representation of events. Technical report, Wright Patterson AFB, Ohio, 1957.
- [Nak00] Atsuyoshi Nakamura. Query learning of bounded-width obdds. *Theor. Comput. Sci.*, 241(1-2):83–114, 2000. doi:10.1016/S0304-3975(99)00267-4.
- [Ner58] A. Nerode. Linear automaton transformations. In *Proceedings of the American Mathematical Society*, 9(4), page 541–544, 1958.
- [NFZ21] Dolav Nitay, Dana Fisman, and Michal Ziv-Ukelson. Learning of structurally unambiguous probabilistic grammars. 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*, pages 9170–9178. AAAI Press, 2021. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/17107>.

- [OG92] J. Oncina and P. García. Inferring regular languages in polynomial update time. In *Selected Papers from the 4th Spanish Symposium on Pattern Recognition and Image Analysis*, pages 49–61, 1992. doi:10.1142/9789812797902\0004.
- [PGB⁺08] Corina S. Pasareanu, Dimitra Giannakopoulou, Mihaela Gheorghiu Bobaru, Jamieson M. Cobleigh, and Howard Barringer. Learning to divide and conquer: applying the L* algorithm to automate assume-guarantee reasoning. *Formal Methods Syst. Des.*, 32(3):175–205, 2008. doi:10.1007/s10703-008-0049-6.
- [PGLM15] Mila Dalla Preda, Roberto Giacobazzi, Arun Lakhotia, and Isabella Mastroeni. Abstract symbolic automata: Mixed syntactic/semantic similarity analysis of executables. In Sriram K. Rajamani and David Walker, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 329–341. ACM, 2015. doi:10.1145/2676726.2676986.
- [Pit89] Leonard Pitt. Inductive inference, DFAs, and computational complexity. In Klaus P. Jantke, editor, *Analogical and Inductive Inference, International Workshop AII '89, Reinhardsbrunn Castle, GDR, October 1-6, 1989, Proceedings*, volume 397 of *Lecture Notes in Computer Science*, pages 18–44. Springer, 1989. doi:10.1007/3-540-51734-0\50.
- [Sak90] Yasubumi Sakakibara. Learning context-free grammars from structural data in polynomial time. *Theor. Comput. Sci.*, 76(2-3):223–242, 1990. doi:10.1016/0304-3975(90)90017-C.
- [She19] Sarai Sheinvald. Learning deterministic variable automata over infinite alphabets. In Maurice H. ter Beek, Annabelle McIver, and José N. Oliveira, editors, *Formal Methods - The Next 30 Years - Third World Congress, FM 2019, Porto, Portugal, October 7-11, 2019, Proceedings*, volume 11800 of *Lecture Notes in Computer Science*, pages 633–650. Springer, 2019. doi:10.1007/978-3-030-30942-8\37.
- [SV17] Olli Saarikivi and Margus Veanes. Translating c# to branching symbolic transducers. In Thomas Eiter, David Sands, Geoff Sutcliffe, and Andrei Voronkov, editors, *IWIL@LPAR 2017 Workshop and LPAR-21 Short Presentations, Maun, Botswana, May 7-12, 2017*, volume 1 of *Kalpa Publications in Computing*, pages 86–99. EasyChair, 2017. doi:10.29007/7jff1.
- [Vaa17] Frits W. Vaandrager. Model learning. *Commun. ACM*, 60(2):86–95, 2017. doi:10.1145/2967606.
- [VdHT10] Margus Veanes, Peli de Halleux, and Nikolai Tillmann. Rex: Symbolic regular expression explorer. In *Third International Conference on Software Testing, Verification and Validation, ICST 2010, Paris, France, April 7-9, 2010*, pages 498–507. IEEE Computer Society, 2010. doi:10.1109/ICST.2010.15.
- [VGRW22] F.W. Vaandrager, B. Garhewal, J. Rot, and T. Wißmann. A new approach for active automata learning based on apartness. In Dana Fisman and Grigore Rosu, editors, *Tools and Algorithms for the Construction and Analysis of Systems TACAS 2022*, volume 13243 of *Lecture Notes in Computer Science*, pages 223–243. Springer, 2022. doi:10.1007/978-3-030-99524-9\12.
- [ZSZR18] Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N. Rafferty. An overview of machine teaching. *CoRR*, abs/1801.05927, 2018. URL: <http://arxiv.org/abs/1801.05927>, arXiv:1801.05927.

APPENDIX A. EFFICIENT IDENTIFICATION OF DFAS

It was shown by [Gol78, OG92] that DFAs are identifiable in the limit using polynomial time and data. Since the proof of Proposition 10.5 relies on some properties of the involved procedures and for completeness of the presentation, we provide a complete description of the procedures showing that DFAs are identifiable in the limit using polynomial time and data, and that they satisfy the required properties.

Theorem 7.5 (restated) [OG92] *I. The class of DFAs is efficiently identifiable via procedures **CharDFA** and **InferDFA**. II. Furthermore, these procedures satisfy that if \mathcal{D} is a minimal and complete DFA and $\mathbf{CharDFA}(\mathcal{D}) = \mathcal{S}_{\mathcal{D}}$ then the following holds:*

- (1) $\mathcal{S}_{\mathcal{D}}$ contains a prefix-closed set A of access words. Moreover, A can be chosen to contain only lex-access words, i.e., only the lexicographically smallest access word for each state.
- (2) For every $u_1, u_2 \in A$ it holds that $u_1 \not\sim_{\mathcal{S}_{\mathcal{D}}} u_2$.
- (3) For every $u, v \in A$ and $\sigma \in \Sigma$, if $\Delta(q_i, u\sigma) \neq \Delta(q_i, v)$ then $u\sigma \not\sim_{\mathcal{S}_{\mathcal{D}}} v$.

To prove Theorem 7.5 we first show, in subsection A.1, that given a DFA $\mathcal{D} = \langle \Sigma, Q, q_i, F, \Delta \rangle$ we can construct a polynomial-sized sample of words $\mathcal{S}_{\mathcal{D}}$ that agrees with \mathcal{D} and satisfies the required properties. In subsection A.2 we show an algorithm that (i) can infer in polynomial time from a given sample \mathcal{S} a DFA that agrees with \mathcal{S} , and (ii) if it is given the set $\mathcal{S}_{\mathcal{D}}$, or any set $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{D}}$ that agrees with \mathcal{D} , then it infers a DFA that is equivalent to \mathcal{D} . All this together proves Theorem 7.5 (and explains why we can refer to $\mathcal{S}_{\mathcal{D}}$ as the *characteristic sample*).

A.1. Constructing a characteristic set. The algorithm **CharDFA** works as follows. It first creates a prefix-closed set of access words to states. This can be done by considering the graph of the automaton and running an algorithm for finding a spanning tree from the initial state. Choosing one of the letters on each edge, the access word for a state is obtained by concatenating the labels on the unique path of the obtained tree that reaches that state. If we wish to work with lex-access words, we can use a depth-first search algorithm that spans branches according to the order of letters in Σ , starting from the smallest. The labels on the paths of the spanning tree constructed this way will form the set of lex-access words.

Let S be the set of access words (or lex-access words). Next the algorithm turns to find a distinguishing word $v_{i,j}$ for every pair of state $s_i, s_j \in S$ (where $s_i \neq s_j$). Lemma A.1 below states that any pair of states of the minimal DFA has a distinguishing word of size quadratic in the size of the DFA. Let E be the set of all such distinguishing words. We may assume $\epsilon \in E$.¹⁵ Then the algorithm returns the set $\mathcal{S}_{\mathcal{D}} = \{\langle w, \mathcal{D}(w) \rangle \mid w \in (S \cdot E) \cup (S \cdot \Sigma \cdot E)\}$ where $\mathcal{D}(w)$ is the label \mathcal{D} gives w (i.e., 1 if it is accepted, and 0 otherwise).

It is easy to see that $\mathcal{S}_{\mathcal{D}}$ satisfies the properties of Theorem 7.5.

Lemma A.1. *Let $\mathcal{D} = \langle \Sigma, Q, q_i, F, \Delta \rangle$ be a minimal DFA, and let $q_1, q_2 \in Q$ s.t. $q_1 \neq q_2$. There exists a polynomial time procedure that returns a word v of size at most $|Q|^2$ such that $\Delta(q_1, v)$ is accepting iff $\Delta(q_2, v)$ is rejecting.*

¹⁵Unless \mathcal{D} accepts all words or rejects all words, it has at least one accepting state and one rejecting state, and ϵ is the shortest word distinguishing these states. If all states of \mathcal{D} are accepting (or all rejecting) the algorithm returns $\mathcal{S}_{\mathcal{D}} = \{\langle \epsilon, 1 \rangle\}$ (resp. $\mathcal{S}_{\mathcal{D}} = \{\langle \epsilon, 0 \rangle\}$).

Proof. We can apply the product construction to $\mathcal{D}_i = \langle \Sigma, Q, q_i, F, \Delta \rangle$ for $i \in \{1, 2\}$ and search for a path from the initial state (q_1, q_2) to a state in $F \times (Q \setminus F)$ or $(Q \setminus F) \times F$ to find a word that leads to an accepting state when read from q_1 and a rejecting state when read from q_2 or vice versa. Since a shortest simple path in a graph is bounded by the number of nodes, the shortest such word is of length at most $|Q|^2$. The shortest path can be found using breadth-first search algorithms that run in time linear in the number of vertices and edges, thus polynomial in the size of the DFA. \square

Since computing a spanning tree (in particular via DFS) and finding shortest paths can be done in polynomial time this shows that for DFAs we can construct the characteristic set in polynomial time. That is, while Definition 7.2 only requires that the characteristic set be of polynomial size, for DFAs we can show that it can also be computed in polynomial time.

A.2. Inferring a DFA. Next we describe algorithm **InferDFA** that given a sample of words \mathcal{S} , infers from it in polynomial time a DFA that agrees with \mathcal{S} . And moreover, if \mathcal{S} subsumes the characteristic set $\mathcal{S}_{\mathcal{D}}$ of a DFA \mathcal{D} then **InferDFA** returns a DFA that recognizes \mathcal{D} .

Let W be the set of words in the given sample \mathcal{S} (without their labels). Let R be the set of prefixes of W and C the set of suffixes of W . Note that $\epsilon \in R$ and $\epsilon \in C$. Let r_0, r_1, \dots be some enumeration of R and c_0, c_1, \dots some enumeration of C where $r_0 = c_0 = \epsilon$. In the sequel we often use i_w for the index of w in R . The algorithm builds a matrix M of size $|R| \times |C|$ whose entries take values in $\{0, 1, ?\}$. The algorithm set the value of entry (i, j) as follows. If $r_i c_j$ is not in W , it is set to $?$. Otherwise it is set to 1 iff the word $r_i c_j$ is labeled 1 in \mathcal{S} . We get that $r_i \sim_{\mathcal{S}} r_j$ iff for every k such that both $M(i, k)$ and $M(j, k)$ are different from $?$ we have that $M(i, k) = M(j, k)$.

The algorithm sets $R_0 = \{\epsilon\}$. Once R_i is constructed, the algorithm tries to establish whether $r\sigma$ for $r \in R_i$ and $\sigma \in \Sigma$ is distinguished from all words in R_i . It does so by considering all other words $r' \in R_i$ and checking whether $r \sim_{\mathcal{S}} r'$. If $r\sigma$ is found to be distinct from all words in R_i , then R_{i+1} is set to $R_i \cup \{r\sigma\}$. The algorithm proceeds until no new words are distinguished. Let k be minimal such that $R_k = R_{k'}$ for all $k' > k$, and let $R = R_k$. If not all words in R are in W (that is $M(i, 0) = ?$ for some $r_i \in R$), the algorithm returns the prefix-tree automaton.¹⁶ Otherwise, the states of the constructed DFA are set to be the words in R . The initial state is ϵ and a state r_i is classified as accepting iff $M(i, 0) = 1$ (recall that the entry $M(i, 0)$ stands for the value of $r_i \cdot \epsilon$ in \mathcal{S}). To determine the transitions, for every $r \in R$ and $\sigma \in \Sigma$, recall that there exists at least one state $r' \in R$ that cannot be distinguished from $r\sigma$. The algorithm then adds a transition from r on σ to r' .

Proposition A.2.

- (1) Algorithm **InferDFA** runs in polynomial time and returns a DFA that agrees with the given sample \mathcal{S} .
- (2) Let $\mathcal{S}_{\mathcal{D}}$ be the sample constructed for a DFA \mathcal{D} by algorithm **CharDFA**, and let $\mathcal{S} \supseteq \mathcal{S}_{\mathcal{D}}$. Then algorithm **InferDFA** returns a DFA that recognized the same language as \mathcal{D} .

¹⁶The prefix-tree automaton is the automaton obtained by placing all words in a tree data structure (sharing common prefixes) and labeling a state accepting iff the unique word reaching that state is in the sample and is labeled 1.

Proof.

- (1) The number of prefixes (or suffixes) of a set of words is bounded by the size of the longest word times the size of the set. Thus M is of polynomial size, and so is its construction. The number of iterations required for converging the R_i sets is bounded by $|W|$. The prefix-tree automaton can be computed in polynomial time. Determining acceptance is polynomial in $|R|$, and determining the transitions is polynomial in $|R| \times |\Sigma|$. Thus the overall running time of the algorithm is polynomial.

Clearly, if the algorithm returns the prefix-tree automaton then it agrees with the given sample \mathcal{S} . We claim that it agrees with the given sample also in the second case. We show, by induction on the length of the word, that for every $w \in W$, if w reaches state r of the constructed DFA, then $w \sim_{\mathcal{S}} r$. Since w is in the sample, and r is in the sample (otherwise the algorithm would return the prefix-tree automaton), it follows that $M(i_w, 0) = M(i_r, 0)$ hence the DFA agrees with the sample on w .

For the base case we have that $|w| = 0$ then the DFA accepts if r_0 is accepting, which holds iff $M(0, 0) = 1$. Indeed this entry is filled with the label of ϵ in \mathcal{S} . Consider now $w = v\sigma$ for some $v \in \Sigma^*$ and $\sigma \in \Sigma$. Assume the DFA reaches state s_ℓ on reading v and s_m after reading w . By induction hypothesis, we know that $r_\ell \sim_{\mathcal{S}} v$. From the construction of the algorithm it follows that $r_\ell\sigma \sim_{\mathcal{S}} s_m$ as otherwise, reading σ from r_ℓ would lead to a different state. If $r_m \not\sim_{\mathcal{S}} w$ then there exists a suffix $c_i \in C$ s.t. $M(m, i) \neq M(i_w, i)$. But then σc_i is also in C , denote it by c_j . Then $M(\ell, j) \neq M(i_v, j)$ contradicting that $r_\ell \sim_{\mathcal{S}} v$.

- (2) Next we show that if \mathcal{S} subsumes $\mathcal{S}_{\mathcal{D}}$ then the returned DFA agrees with \mathcal{D} . Let w_1, \dots, w_n be the set of accessible words chosen by **CharDFA**. Since \mathcal{S} consists of a distinguished word for every pair of access words w_i, w_j of \mathcal{D} , algorithm **InferDFA** will determine $w_i \not\sim_{\mathcal{S}} w_j$ and R will consist of at least n states. It may not consist of more states, since the sample has to agree with the language of \mathcal{D} and every word w agrees with some state of \mathcal{D} on all possible suffixes, thus **InferDFA** cannot determine that w corresponds to a new distinct state. Since $S \cdot \Sigma \cdot E$ was placed in \mathcal{S} , for every distinguished state w and every $\sigma \in \Sigma$ the algorithm **InferDFA** can determine the transition from w upon reading σ . Since $S \cdot \epsilon$ is placed in \mathcal{S} , algorithm **InferDFA** can correctly label acceptance of states. Thus the obtained DFA is isomorphic to the original DFA. \square

We can thus conclude that DFAs are identifiable in the limit using polynomial time and data. Furthermore, they satisfy the properties of Theorem 7.5.