

CONSTRUCTING CONCISE CHARACTERISTIC SAMPLES FOR ACCEPTORS OF OMEGA REGULAR LANGUAGES

DANA ANGLUIN  AND DANA FISMAN 

^a Yale University, New Haven, CT, USA
e-mail address: dana.angluin@yale.edu

^b Ben-Gurion University, Be'er-Sheva, Israel
e-mail address: dana@cs.bgu.ac.il

ABSTRACT. A characteristic sample for a language L and a learning algorithm \mathbf{L} is a finite sample of words T_L labeled by their membership in L such that for any sample $T \supseteq T_L$ consistent with L , on input T the learning algorithm \mathbf{L} returns a hypothesis equivalent to L . Which omega automata have characteristic sets of polynomial size, and can these sets be constructed in polynomial time? We address these questions here.

In brief, non-deterministic omega automata of any of the common types, in particular Büchi, do not have characteristic samples of polynomial size. For deterministic omega automata that are isomorphic to their right congruence automata, the *fully informative languages*, polynomial time algorithms for constructing characteristic samples and learning from them are given.

The algorithms for constructing characteristic sets in polynomial time for the different omega automata (of types Büchi, coBüchi, parity, Rabin, Street, or Muller), require deterministic polynomial time algorithms for (1) equivalence of the respective omega automata, and (2) testing membership of the language of the automaton in the informative classes, which we provide.

CONTENTS

1. Introduction	3
2. Preliminaries	5
3. Notions of learnability and teachability	8
3.1. Teachers, learners, and characteristic samples	8
3.2. Computable teachers and learners	9
3.3. Polynomial time learners and teachers	12
3.4. Identification in the limit with polynomial time and data.	15
3.5. Relation to learning with equivalence and membership queries	16
4. Focusing on regular ω -languages	18
4.1. Examples and samples for ω -languages	18
4.2. Negative results for nondeterministic classes	19
4.3. Consequences of membership and equivalence algorithms	19
5. The informative classes are efficiently teachable	20

5.1.	Duality	20
5.2.	The default acceptor	21
5.3.	Strongly connected components	22
5.4.	Proving efficient teachability of the informative classes — overview	24
6.	The sample T_{Aut} for the automaton	25
6.1.	Existence of short distinguishing words	25
6.2.	Defining the sample T_{Aut} for the automaton	25
6.3.	Learning the automaton from T_{Aut}	26
7.	The samples T_{Acc} and learning algorithms for IMA and IBA	27
7.1.	Muller acceptors	27
7.2.	Büchi acceptors	27
8.	The sample T_{Acc} and the learning algorithm for IPA	28
8.1.	A decreasing forest of SCCs of an automaton	29
8.2.	Constructing the Canonical Forest and Coloring of a DPA	29
8.3.	Constructing T_{Acc}^{IPA}	32
8.4.	The learning algorithm \mathbf{L}_{Acc}^{IPA}	32
9.	The sample T_{Acc} and learning algorithm for IRA	34
9.1.	Singleton normal form for a Rabin acceptor	34
9.2.	An ordering on sets of states	34
9.3.	The learning algorithm \mathbf{L}_{Acc}^{IRA}	35
9.4.	Constructing T_{Acc}^{IRA}	35
9.5.	Correctness of \mathbf{L}_{Acc}^{IRA}	37
10.	Constructing characteristic samples in polynomial time	38
10.1.	Computing T_{Aut}	38
10.2.	Computing T_{Acc}	39
11.	Inclusion algorithms	40
12.	Inclusion and equivalence for DPAs, DBAs, DCAs	41
12.1.	Searching for w with given minimum colors in two acceptors	41
12.2.	An inclusion algorithm for DPAs	42
13.	An inclusion algorithm for DRAs	42
14.	An inclusion algorithm for DMAs	44
14.1.	Reduction of DMA inclusion to DBA/DMA inclusion	44
14.2.	A DBA/DMA inclusion algorithm	46
15.	Computing the automaton \mathcal{M}_{\sim_L}	48
16.	Testing membership in \mathbb{IXA}	49
16.1.	Testing membership in \mathbb{IBA}	49
16.2.	Testing membership in \mathbb{IPA}	50
16.3.	Testing membership in \mathbb{IRA}	51
16.4.	Testing membership in \mathbb{IMA}	54
16.5.	Variants of the testing algorithms	55
16.6.	Efficient teachability of the informative classes	55
17.	Discussion	56
	Acknowledgments	56
	References	56

1. INTRODUCTION

With the growing success of machine learning in efficiently solving a wide spectrum of problems, we are witnessing an increased use of machine learning techniques in formal methods for system design. One thread in recent literature uses general purpose machine learning techniques for obtaining more efficient verification/synthesis algorithms. Another thread, following the *automata theoretic approach to verification* [Var95, KVV00] works on developing grammatical inference algorithms for verification and synthesis purposes. *Grammatical inference* (aka *automata learning*) refers to the problem of automatically inferring from examples a finite representation (e.g. an automaton, a grammar, or a formula) for an unknown language [dlH10]. The term *model learning* [Vaa17] was coined for the task of learning an automaton model for an unknown system. A large body of works has developed learning techniques for different automata types (e.g. visibly-pushdown automata [KMV06], I/O automata [AV10], register automata [HSJC12], symbolic automata [DD17], program automata [MS18], probabilistic grammars [NFZ21], lattice automata [FS22]) and has shown its usability in a diverse range of tasks.¹

In grammatical inference, the learning algorithm does not learn a *language*, but rather a finite *representation* of it. The complexity of learning algorithms may vary greatly by switching representations. For instance, if one wishes to learn regular languages, she may consider representations using deterministic finite automata (DFAs), non-deterministic finite automata (NFAs), regular expressions, linear grammars, etc. Since the translation results between two such formalisms are not necessarily polynomial, a polynomial learnability result for one representation does not necessarily imply a polynomial learnability result for another representation. Let \mathbb{C} be a class of representations \mathcal{C} with a size measure $size(\mathcal{C})$ (e.g. for DFAs the size measure can be the number of states in the minimal DFA). We extend $size(\cdot)$ to the languages recognized by representations in \mathbb{C} by defining $size(L)$ to be the minimum of $size(\mathcal{C})$ over all \mathcal{C} representing L . In this paper we restrict attention to automata representations, namely, *acceptors*.

There are various learning paradigms considered in the grammatical inference literature, roughly classified into *passive* and *active*. We mention here the two central ones. In *passive learning* the model of *learning from finite data* refers to the following problem: given a finite sample $T \subseteq \Sigma^* \times \{0, 1\}$ of labeled words, a learning algorithm \mathbf{L} should return an acceptor \mathcal{C} that agrees with the sample T . That is, for every $(w, l) \in T$ the following holds: $w \in \llbracket \mathcal{C} \rrbracket$ iff $l = 1$ (where $\llbracket \mathcal{C} \rrbracket$ is the language accepted by \mathcal{C}). The class \mathbb{C} is *identifiable in the limit using polynomial time and data* if and only if there exists a polynomial time algorithm \mathbf{L} that takes as input a labeled sample T and outputs an acceptor $\mathcal{C} \in \mathbb{C}$ that is consistent with T , and \mathbf{L} also satisfies the following condition. If L is any language recognized by an automaton from class \mathbb{C} , then there exists a labeled sample T_L consistent with L of length bounded by a polynomial in $size(L)$, and for any labeled sample T consistent with L such that $T_L \subseteq T$, on input T the algorithm \mathbf{L} produces an acceptor \mathcal{C} that recognizes L . In this case, T_L is termed a *characteristic sample* for the algorithm \mathbf{L} . The definition of identifiability in the limit using polynomial time and data relates to learning paradigms considering a *teacher-learner* pair [GM96]. While identification in the limit using polynomial

¹E.g. , tasks such as black-box checking [PVY99], specification mining [ABL02], assume-guarantee reasoning [CGP03], regular model checking [HV05], learning verification fixed-points [VSVA05], learning interfaces [NA06], analyzing botnet protocols [CBSS10] or smart card readers [CPPdR14], finding security bugs [CPPdR14], error localization [CCK⁺15], and code refactoring [MNRS04, SHV16].

time and data does not require that the characteristic set can be computed in polynomial time, if one is interested in devising a teacher that can train a learner, having a teacher that can compute a characteristic set in polynomial time is also desired. The definition of *efficiently teachable* adds this requirement. In Section 3 we define several notions related to efficient teachability and learnability, the stronger one is *efficiently teachable*. The question which representations of regular ω -languages are efficiently teachable is the focus of this paper.

In *active learning* the model of *query learning* [Ang87] assumes the learner communicates with an *oracle* that can answer certain types of queries about the language. The most common types of queries are *membership queries* (is $w \in L$ where L is the unknown language) and *equivalence queries* (is $\llbracket \mathcal{A} \rrbracket = L$ where \mathcal{A} is the current hypothesis for an acceptor recognizing L). Equivalence queries are typically assumed to return a counterexample, i.e. a word in $\llbracket \mathcal{A} \rrbracket \setminus L$ or in $L \setminus \llbracket \mathcal{A} \rrbracket$.

With regard to ω -automata (automata on infinite words) most of the works consider *query learning* using *membership queries* and *equivalence queries*. The representations learned so far include: $(L)_\S$ [FCC⁺08], a non-polynomial reduction to finite words; families of DFAs (FDFA) [AF14, AF16, ABF16, LCZL21]; strongly unambiguous Büchi automata (SUBA) [AAF20]; mod-2-multiplicity automata (M2MA) [AAFG22]; and deterministic weak parity automata (DWPA) [MP95]. Among these only the latter two are known to be learnable in polynomial time using membership queries and proper equivalence queries.² We show in Section 4.3 that the classes M2MA, SUBA and DWPA are efficiently teachable.

One of the main obstacles in obtaining a polynomial learning algorithm for regular ω -languages is that they do not in general have a Myhill-Nerode characterization; that is, there is no theorem correlating the states of a minimal automaton of some of the common automata types (Büchi, parity, Muller, etc.) to the equivalence classes of the right congruence of the language. The *right congruence relation* for an ω -language L relates two finite words x and y iff there is no infinite suffix z differentiating them, that is $x \sim_L y$ (for $x, y \in \Sigma^*$) iff $\forall z \in \Sigma^\omega. xz \in L \iff yz \in L$. The quest for finding a polynomial query learning algorithm for a subclass of the regular ω -languages, led to studying subclasses of languages for which such a relation holds. These languages are termed *fully informative* [AF18]. We use $\mathbb{I}BA, \mathbb{I}CA, \mathbb{I}PA, \mathbb{I}RA, \mathbb{I}SA, \mathbb{I}MA$ to denote the classes of languages that are fully informative of type Büchi, coBüchi, parity, Rabin, Streett and Muller, respectively. A language L is said to be fully informative of type \mathbb{X} for $\mathbb{X} \in \{\mathbb{B}, \mathbb{C}, \mathbb{P}, \mathbb{R}, \mathbb{S}, \mathbb{M}\}$ if there exists a deterministic automaton of type \mathbb{X} that recognizes L and is isomorphic to the automaton derived from \sim_L . While many properties of these classes are now known, in particular that they span the entire hierarchy of regular ω -languages [Wag75], a polynomial learning algorithm for them is not known.

We show (in Sections 5-9) that the classes $\mathbb{I}BA, \mathbb{I}CA, \mathbb{I}PA, \mathbb{I}RA, \mathbb{I}SA, \mathbb{I}MA$ can be identified in the limit using polynomial time and data. We further show (in Section 10) that there is a polynomial time algorithm to compute a characteristic sample given an acceptor $\mathcal{C} \in \mathbb{I}XA$. To show that these classes are also efficiently teachable we need polynomial time algorithms for inclusion and equivalence of automata of these types, that also return shortlex counterexamples in case of inequivalence.³ Such an algorithm is known to exist for the classes $\mathbb{N}BA, \mathbb{N}CA, \mathbb{N}PA$, since these classes have inclusion algorithms in NL [Sch10]. For

²Query learning with an additional type of query, *loop-index queries*, was studied for deterministic Büchi automata [MO20].

³The formal definition of shortlex is deferred to Section 2.

the other classes a polynomial-time algorithm can be obtained following a reduction to model checking a certain fragment of CTL* formulas [CDK93]. However this reduction does not yield shortlex counterexamples. We provide such algorithms in Sections 11-14.

The last part of this paper (Sections 15-16) is devoted to the question of deciding whether a given automaton \mathcal{A} of type \mathbb{X} is isomorphic to its right congruence, or if this is not the case whether there exists an automaton \mathcal{A}' of the same type that recognizes the same language and is isomorphic to its right congruence, namely whether the given automaton recognizes a language in the class \mathbb{IXA} . Using this result we can show that a teacher can construct a characteristic sample not only given an acceptor which is isomorphic to the right congruence of the language, but also given an acceptor which is not, but is equivalent to such an acceptor. We conclude in Section 17 with a short discussion.

2. PRELIMINARIES

Automata. An *automaton* is a tuple $\mathcal{M} = \langle \Sigma, Q, q_\iota, \delta \rangle$ consisting of a finite alphabet Σ of symbols, a finite set Q of states, an initial state $q_\iota \in Q$, and a transition function $\delta : Q \times \Sigma \rightarrow 2^Q$. We extend δ to domain $Q \times \Sigma^*$ in the usual way: $\delta(q, \varepsilon) = q$ and $\delta(q, \sigma x) = \cup_{q' \in \delta(q, \sigma)} \delta(q', x)$ for all $q \in Q$ and $\sigma \in \Sigma$.

We define the *size* of an automaton to be $|\Sigma| \cdot |Q|$. A state $q \in Q$ is *reachable* iff there exists $x \in \Sigma^*$ such that $q \in \delta(q_\iota, x)$. For $q \in Q$, \mathcal{M}^q is the automaton \mathcal{M} with its initial state replaced by q . We say that \mathcal{A} is *deterministic* if $|\delta(q, \sigma)| \leq 1$ and *complete* if $|\delta(q, \sigma)| \geq 1$, for every $q \in Q$ and $\sigma \in \Sigma$. For deterministic automata we abbreviate $\delta(q, \sigma) = \{q'\}$ as $\delta(q, \sigma) = q'$. Two automata \mathcal{M} and \mathcal{M}' with the same alphabet Σ are isomorphic if there exists a bijection f from the states Q of \mathcal{M} to the states Q' of \mathcal{M}' such that $f(q_\iota) = q'_\iota$ and for every $q \in Q$ and $\sigma \in \Sigma$, $\{f(r) \mid r \in \delta(q, \sigma)\} = \delta'(f(q), \sigma)$.

We assume a fixed total ordering on Σ , which induces the *shortlex* total ordering on Σ^* , defined as follows. For $x, y \in \Sigma^*$, x precedes y in the shortlex ordering if $|x| < |y|$ or $|x| = |y|$ and x precedes y in the lexicographic ordering induced by the ordering on Σ .

A *run* of an automaton on a finite word $v = a_1 a_2 \dots a_n$ is a sequence of states q_0, q_1, \dots, q_n such that $q_0 = q_\iota$, and for each $i \geq 1$, $q_i \in \delta(q_{i-1}, a_i)$. A *run* on an infinite word is defined similarly and consists of an infinite sequence of states. For an infinite run $\rho = q_0, q_1, \dots$, we define the set of states visited infinitely often, denoted $\text{inf}_{\mathcal{M}}(\rho)$, as the set of $q \in Q$ such that $q = q_i$ for infinitely many indices $i \in \mathbb{N}$. This is abbreviated to $\text{inf}(\rho)$ if \mathcal{M} is understood.

The product of two automata. Let \mathcal{M}_1 and \mathcal{M}_2 be two deterministic complete automata with the same alphabet Σ , where for $i = 1, 2$, $\mathcal{M}_i = \langle \Sigma, Q_i, (q_\iota)_i, \delta_i \rangle$. Their product automaton, denoted $\mathcal{M}_1 \times \mathcal{M}_2$, is the deterministic complete automaton $\mathcal{M} = \langle \Sigma, Q, q_\iota, \delta \rangle$ such that $Q = Q_1 \times Q_2$ is the set of ordered pairs of states of \mathcal{M}_1 and \mathcal{M}_2 ; the initial state $q_\iota = ((q_\iota)_1, (q_\iota)_2)$ is the pair of initial states of the two automata; and for all $(q_1, q_2) \in Q$ and $\sigma \in \Sigma$, $\delta((q_1, q_2), \sigma) = (\delta_1(q_1, \sigma), \delta_2(q_2, \sigma))$. For $i = 1, 2$, let π_i be projection onto the i -th coordinate, so that for a subset S of Q , $\pi_1(S) = \{q_1 \in Q_1 \mid \exists q_2 \in Q_2. (q_1, q_2) \in S\}$, and analogously for π_2 .

Acceptors. By augmenting an automaton $\mathcal{M} = \langle \Sigma, Q, q_\iota, \delta \rangle$ with an acceptance condition α , obtaining a tuple $\mathcal{A} = \langle \Sigma, Q, q_\iota, \delta, \alpha \rangle$, we get an *acceptor*, a machine that accepts some words and rejects others. We may also denote \mathcal{A} by (\mathcal{M}, α) . An acceptor accepts a word if at least one of the runs on that word is accepting. If the automaton is not complete, a given word w may not have any run in the automaton, in which case w is rejected.

For finite words the acceptance condition is a set $F \subseteq Q$ and a run on a word v is accepting if it ends in an accepting state, i.e., if $\delta(q_\iota, v)$ contains an element of F . For infinite words, there are various acceptance conditions in the literature, and we consider six of them: Büchi, coBüchi, parity, Rabin, Streett and Muller, all based on the set of states visited infinitely often in a given run. For each model we define the related quantity of the size of the acceptor, taking into account the acceptance condition.

A *Büchi* or *coBüchi* acceptance condition is a set of states $F \subseteq Q$. A run ρ of a Büchi acceptor is accepting if it visits F infinitely often, that is, $\text{inf}(\rho) \cap F \neq \emptyset$. A run ρ of a coBüchi acceptor is accepting if it visits F only finitely many times, that is, $\text{inf}(\rho) \cap F = \emptyset$. The *size* of a Büchi or coBüchi acceptor is the size of its automaton.

A *parity* acceptance condition is a map $\kappa : Q \rightarrow \mathbb{N}$ assigning to each state a natural number termed a color (or priority). A run of a parity acceptor is accepting if the **minimum** color visited infinitely often is **odd**. The *size* of a parity acceptor is the size of its automaton.

A *Rabin* or *Streett* acceptance condition consists of a finite set of pairs of sets of states $\alpha = \{(G_1, B_1), \dots, (G_k, B_k)\}$ for some $k \in \mathbb{N}$ and $G_i \subseteq Q$ and $B_i \subseteq Q$ for $i \in [1..k]$. A run of a Rabin acceptor is accepting if there exists an $i \in [1..k]$ such that G_i is visited infinitely often and B_i is visited finitely often. A run of a Streett acceptor is accepting if for all $i \in [1..k]$, G_i is visited finitely often or B_i is visited infinitely often. The *size* of a Rabin or Streett acceptor is the sum of the size of its automaton and $k - 1$.

A *Muller* acceptance condition is a set of sets of states $\alpha = \{F_1, F_2, \dots, F_k\}$ for some $k \in \mathbb{N}$ and $F_i \subseteq Q$ for $i \in [1..k]$. A run of a Muller acceptor is accepting if the set S of states visited infinitely often in the run is a member of α . The *size* of a Muller acceptor is the sum of the size of its automaton and $k - 1$.

The set of words accepted by an acceptor \mathcal{A} is denoted by $\llbracket \mathcal{A} \rrbracket$. $L_1 \oplus L_2$ is the symmetric difference of sets L_1 and L_2 : $(L_1 \setminus L_2) \cup (L_2 \setminus L_1)$. Two acceptors \mathcal{A} and \mathcal{B} are *equivalent* if they accept the same language, that is, $\llbracket \mathcal{A} \rrbracket = \llbracket \mathcal{B} \rrbracket$. For a state q , the acceptor \mathcal{A}^q is the acceptor \mathcal{A} with its automaton initial state replaced by q . We say that the ω -word w is *accepted from state q* iff $w \in \llbracket \mathcal{A}^q \rrbracket$.

We use three-letter acronyms for automata and classes concerning the common ω -automata discussed above. The first letter is in $\{N, D, I\}$ and stands for *non-deterministic*, *deterministic* and *isomorphic* (or *fully informative*) which will be explained in the sequel. The second letter describes the acceptance condition, and the third letter A stands for *acceptor*. That is, we use NBA, NCA, NPA, NRA, NSA, NMA (resp., DBA, DCA, DPA, DRA, DSA, DMA) for non-deterministic (resp., deterministic) Büchi, coBüchi, parity, Rabin, Streett and Muller acceptors. We use blackboard font for the respective classes of representations. That is, we use NBA, NCA, NPA, NMA, NRA and NSA (resp., DBA, DCA, DPA, DRA, DSA and DMA) for the corresponding class of representations. It is known that NCA and DCA recognize the same languages and that the classes DCA and DBA are distinct proper subclasses of the regular ω -languages. The other classes are the full class of regular ω -languages.

Some relationships between the models. We observe the following known relationships [GTW02, Chapter 1].

Claim 2.1. Let \mathcal{A} be an acceptor of one of the types NBA, NCA, NPA, NRA, NSA, or NMA with n states over the alphabet Σ . There is an equivalent complete acceptor \mathcal{A}' of the same type whose size is at most $|\Sigma|$ larger. \mathcal{A}' may be taken to be deterministic if \mathcal{A} is deterministic.

- Claim 2.2.** (1) Let $\mathcal{B} = \langle \Sigma, Q, q_i, \delta, F \rangle$, where \mathcal{B} is an NBA. Define the NPA $\mathcal{P} = \langle \Sigma, Q, q_i, \delta, \kappa \rangle$ where $\kappa(q) = 1$ if $q \in F$ and $\kappa(q) = 2$ otherwise. Then \mathcal{B} and \mathcal{P} are equivalent and have the same size. \mathcal{P} is deterministic if \mathcal{B} is.
- (2) Let $\mathcal{B} = \langle \Sigma, Q, q_i, \delta, F \rangle$, where \mathcal{B} is an NBA. Define the NRA $\mathcal{R} = \langle \Sigma, Q, q_i, \delta, \{(F, \emptyset)\} \rangle$. Then \mathcal{B} and \mathcal{R} are equivalent and have the same size. \mathcal{R} is deterministic if \mathcal{B} is.
- (3) Let $\mathcal{C} = \langle \Sigma, Q, q_i, \delta, F \rangle$, where \mathcal{C} is an NCA. Define the NSA $\mathcal{S} = \langle \Sigma, Q, q_i, \delta, \{(F, \emptyset)\} \rangle$. Then \mathcal{S} and \mathcal{C} are equivalent and have the same size. \mathcal{S} is deterministic if \mathcal{C} is.
- (4) Let $\mathcal{B} = \mathcal{C} = \langle \Sigma, Q, q_i, \delta, F \rangle$, where \mathcal{B} is a complete DBA and \mathcal{C} is a complete DCA. Then \mathcal{B} and \mathcal{C} are the same size and the languages they recognize are complements of each other, that is, $\llbracket \mathcal{B} \rrbracket = \Sigma^\omega \setminus \llbracket \mathcal{C} \rrbracket$.
- (5) Let $\mathcal{R} = \mathcal{S} = \langle \Sigma, Q, q_i, \delta, \{(G_1, B_1), \dots, (G_k, B_k)\} \rangle$, where \mathcal{R} is a complete DRA and \mathcal{S} is a complete DSA. Then \mathcal{R} and \mathcal{S} have the same size and the languages they recognize are complements of each other, that is, $\llbracket \mathcal{R} \rrbracket = \Sigma^\omega \setminus \llbracket \mathcal{S} \rrbracket$.

Right congruence. An equivalence relation \sim on Σ^* is a *right congruence* if $x \sim y$ implies $x\sigma \sim y\sigma$ for all $x, y \in \Sigma^*$ and $\sigma \in \Sigma$. The *index* of \sim , denoted $|\sim|$ is the number of equivalence classes of \sim . For a word $x \in \Sigma^*$ the notation $[x]_\sim$ denotes the equivalence class of \sim that contains x .

With a right congruence \sim of finite index one can naturally associate a complete deterministic automaton $\mathcal{M}_\sim = \langle \Sigma, Q, q_i, \delta \rangle$ as follows: the set of states Q consists of the equivalence classes of \sim . The initial state q_i is the equivalence class $[\varepsilon]_\sim$. The transition function δ is defined by $\delta([u]_\sim, \sigma) = [u\sigma]_\sim$ for all $\sigma \in \Sigma$. Also, given a complete deterministic automaton $\mathcal{M} = \langle \Sigma, Q, q_i, \delta \rangle$, we can naturally associate with it a right congruence as follows: $x \sim_{\mathcal{M}} y$ iff \mathcal{M} reaches the same state of \mathcal{M} when reading x or y , that is, $\delta(q_i, x) = \delta(q_i, y)$.

Given a language $L \subseteq \Sigma^*$ its *canonical right congruence* \sim_L is defined as follows: $x \sim_L y$ iff $\forall z \in \Sigma^*. xz \in L \iff yz \in L$. The Myhill-Nerode theorem states that a language $L \subseteq \Sigma^*$ is regular iff \sim_L is of finite index. Moreover, if L is accepted by a complete DFA \mathcal{A} , then $\sim_{\mathcal{M}}$ refines \sim_L , where \mathcal{M} is the automaton of \mathcal{A} . Finally, any complete DFA of minimum size that accepts L has an automaton that is isomorphic to \mathcal{M}_{\sim_L} .

For an ω -language $L \subseteq \Sigma^\omega$, its *canonical right congruence* \sim_L is defined similarly, by quantifying over ω -words. That is, $x \sim_L y$ iff $\forall z \in \Sigma^\omega. xz \in L \iff yz \in L$. If L is a regular ω -language then \sim_L is of finite index, and for any complete DBA (resp., DCA, DPA, DRA, DSA, DMA) \mathcal{A} that accepts L , $\sim_{\mathcal{M}}$ refines \sim_L , where \mathcal{M} is the automaton of the acceptor.

However, for regular ω -languages, the relation \sim_L does not suffice to obtain a ‘‘Myhill-Nerode’’ characterization. In particular, for a regular ω -language L there may be no way to define an acceptance condition for \mathcal{M}_{\sim_L} that yields a DBA (resp., DCA, DPA, DRA, DSA, DMA) that accepts L . As an example consider the language $L = (a + b)^*(bba)^\omega$. Then \sim_L consists of just one equivalence class, because for any $x \in \Sigma^*$ and $w \in \Sigma^\omega$ we have that

Table 1. Summary of the definitions for teachability and learnability

Samples	Learner	Teacher	Definition
arb	arb	arb	characteristic samples
poly	arb	arb	concise characteristic samples
arb	poly	arb	efficiently learnable
arb	arb	poly	efficiently teachable (samples will be poly)
poly	poly	arb	identifiable in the limit using polynomial time and data
arb	poly	poly	efficiently teachable/learnable (samples will be poly)

$xw \in L$ iff w has $(bba)^\omega$ as a suffix. But a DBA (resp., DCA, DPA, DRA, DSA, DMA) that accepts L clearly needs more than a single state.

The fully informative classes. In light of the lack of a Myhill-Nerode result for regular ω -languages, we define a restricted type of deterministic Büchi (resp., coBüchi, parity, Rabin, Streett, Muller) acceptors. Following the three-letter acronym notation introduced in the *Acceptors* subsection, for $X \in \{B, C, P, R, S, M\}$ we say that a DXA \mathcal{A} recognizing a language L is *fully informative* if it is complete and its automaton is isomorphic to \mathcal{M}_{\sim_L} (the automaton that corresponds to the canonical right congruence \sim_L for the language, as defined above). A DXA is an IXA if it is fully informative. A DXA \mathcal{B} is in \mathbb{IXA} if there exists an IXA \mathcal{A} such that $\llbracket \mathcal{B} \rrbracket = \llbracket \mathcal{A} \rrbracket$. We note that every state of an IXA is reachable because every state of \mathcal{M}_{\sim_L} is reachable. Since every state of a minimal automaton for a language L in \mathbb{IXA} for $\mathbb{X} \in \{\mathbb{B}, \mathbb{C}, \mathbb{P}, \mathbb{R}, \mathbb{S}, \mathbb{M}\}$ corresponds to an equivalence class of \sim_L we refer to the \mathbb{IXA} classes as the *fully informative* classes.

Despite the fact that each of these classes is a proper subset of its corresponding deterministic class (e.g., \mathbb{IBA} is a proper subset of \mathbb{DBA}), these classes are more expressive than one might first conjecture. It was shown in [AF18] that in every class of the infinite Wagner hierarchy [Wag75] there are languages in \mathbb{IMA} and \mathbb{IPA} . Moreover, in a small experiment reported in [AF18], among randomly generated Muller automata, the vast majority turned out to be in \mathbb{IMA} .

3. NOTIONS OF LEARNABILITY AND TEACHABILITY

In this section we define and compare general notions of learnability and teachability, and some computable and polynomial time variants of them. A summary of the definitions is provided in Table 1.

3.1. Teachers, learners, and characteristic samples. We are concerned with examples and concepts that can be represented by finite binary strings as follows. $\mathcal{X} = \{0, 1\}^*$ is the domain of representations of examples. (In the next section we describe how we use finite strings as representations of ω -words.) A *concept* is any subset of \mathcal{X} . A *class of concepts* \mathbb{C} consists of the set $\{0, 1\}^*$ of representations of concepts, together with a mapping $\llbracket \cdot \rrbracket$ from $\{0, 1\}^*$ to concepts such that $\llbracket \mathcal{C} \rrbracket$ is the subset of \mathcal{X} that \mathcal{C} represents. Thus every finite binary string represents an example and also a concept. The *length* of the representation of

an example x or concept \mathcal{C} is its length as a string, that is, $|x|$ or $|\mathcal{C}|$. The *size* of a concept \mathcal{C} , denoted $\text{size}(\mathcal{C})$, is the minimum $|\mathcal{C}'|$ of any representation \mathcal{C}' such that $\llbracket \mathcal{C}' \rrbracket = \llbracket \mathcal{C} \rrbracket$.⁴

A *sample* S is a finite set of elements (x, b) where $x \in \mathcal{X}$ and $b \in \{0, 1\}$. The *length* of S is the sum of the lengths of the examples x that appear in it. A sample S is *consistent* with a concept \mathcal{C} iff for every $(x, b) \in S$ we have $b = 1$ iff $x \in \llbracket \mathcal{C} \rrbracket$.

A *learner* for \mathbb{C} is a function \mathbf{L} that maps a sample T to the representation of a concept $\mathbf{L}(T)$ in \mathbb{C} with the property that if T is consistent with at least one element of \mathbb{C} , then $\mathbf{L}(T)$ and T are consistent. A *teacher* for \mathbb{C} is a function \mathbf{T} that maps the representation of a concept \mathcal{C} in \mathbb{C} to a sample $\mathbf{T}(\mathcal{C})$ such that $\mathbf{T}(\mathcal{C})$ and \mathcal{C} are consistent. Note that learners and teachers need not be computable.

A sample T is a *characteristic sample* for \mathcal{C} and a learner \mathbf{L} if T is consistent with \mathcal{C} and for every sample $T' \supseteq T$ consistent with \mathcal{C} we have $\llbracket \mathbf{L}(T') \rrbracket = \llbracket \mathcal{C} \rrbracket$. The intuition is that additional information consistent with \mathcal{C} beyond T will not cause the learner to change its mind about the correct concept.

A class \mathbb{C} *has characteristic samples for a learner* \mathbf{L} if there exists a teacher \mathbf{T} such that for every \mathcal{C} in \mathbb{C} , $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} . A class \mathbb{C} *has characteristic samples* if it has characteristic samples for some learner. A well known property of characteristic samples is the following.

Lemma 3.1 (Key Property of Characteristic Samples). *Assume \mathbb{C} is a class of concepts, \mathbf{T} is a teacher, and \mathbf{L} is a learner such that \mathbf{T} gives a characteristic sample for \mathcal{C} and \mathbf{L} for every \mathcal{C} in \mathbb{C} . If \mathcal{C}_i and \mathcal{C}_j are any concepts from \mathbb{C} such that $\llbracket \mathcal{C}_i \rrbracket \neq \llbracket \mathcal{C}_j \rrbracket$ then there exists some $(x, b) \in \mathbf{T}(\mathcal{C}_i) \cup \mathbf{T}(\mathcal{C}_j)$ such that $x \in \llbracket \mathcal{C}_i \rrbracket \oplus \llbracket \mathcal{C}_j \rrbracket$.*

Proof. Assume to the contrary. Let $T = \mathbf{T}(\mathcal{C}_i) \cup \mathbf{T}(\mathcal{C}_j)$ and consider \mathbf{L} with input T . Because the characteristic sample for \mathcal{C}_i is contained in T and T is consistent with \mathcal{C}_i , \mathbf{L} must output a concept denoting $\llbracket \mathcal{C}_i \rrbracket$. The same is true of \mathcal{C}_j , but because $\llbracket \mathcal{C}_i \rrbracket \neq \llbracket \mathcal{C}_j \rrbracket$ it is impossible for \mathbf{L} to output a concept whose denotation is equal to both of them. \square

3.2. Computable teachers and learners. To discuss computability of teaching and learning, we consider the following three possible properties of a class of concepts.

- (C1) There is an algorithm to decide membership of x in $\llbracket \mathcal{C} \rrbracket$, given example representation x and concept representation \mathcal{C} .
- (C2) There is an algorithm to decide whether there exists a concept representation \mathcal{C} consistent with T , given a sample T .
- (C3) There is an algorithm to decide whether $\llbracket \mathcal{C}_i \rrbracket = \llbracket \mathcal{C}_j \rrbracket$, given two concept representations \mathcal{C}_i and \mathcal{C}_j .

The assumptions (C1) and (C2) are sufficient to guarantee the existence of characteristic samples for \mathbb{C} with a computable learner and a possibly non-computable teacher. The learner uses the algorithm of identification by enumeration [Gol67].

Theorem 3.2 (Identification by Enumeration). *Under the assumptions (C1) and (C2), there is a computable learner \mathbf{L} such that class \mathbb{C} has characteristic samples for \mathbf{L} .*

⁴The notions of size defined for acceptors in Section 2 polynomially relate to the notion of size defined here.

Proof. The learner \mathbf{L} enumerates the finite binary strings representing concepts in shortlex order as $\mathcal{C}_1, \mathcal{C}_2, \dots$. Given a sample T , the learner first uses (C2) to check whether there is any concept in \mathbb{C} consistent with T . If not, it outputs an arbitrary concept. Otherwise, the output of \mathbf{L} is \mathcal{C}_n for the least n such that \mathcal{C}_n is consistent with T . This is possible because by (C1), whether \mathcal{C} is consistent with T is decidable.

The teacher \mathbf{T} with input \mathcal{C} finds the least n such that $\llbracket \mathcal{C}_n \rrbracket = \llbracket \mathcal{C} \rrbracket$, and for each $m < n$ determines an example x_m that distinguishes $\llbracket \mathcal{C}_m \rrbracket$ from $\llbracket \mathcal{C}_n \rrbracket$. The sample output by \mathbf{T} consists of all x_m with $m < n$, labeled to be consistent with \mathcal{C} . \square

We cannot necessarily take the teacher \mathbf{T} in this construction to be computable; in particular, it must test the equivalence of two concepts. In fact, Fisman et al. [FFZ23] show that there exists a class \mathbb{C} that satisfies assumptions (C1) and (C2) and has characteristic samples such that there is no computable function \mathbf{T} to construct a characteristic sample for every language in the class.⁵

However, under the further assumption (C3) that equivalence of concepts is computable, the teacher \mathbf{T} in the proof may be made computable, since if two concepts \mathcal{C}_i and \mathcal{C}_j are determined to be inequivalent, it is safe to do an unbounded search for an example that distinguishes them. Thus we have the following.

Corollary 3.3. *Under the assumptions (C1), (C2) and (C3), there is are a computable teacher \mathbf{T} and a computable learner \mathbf{L} such that \mathbf{T} computes characteristic samples for \mathbb{C} and \mathbf{L} .*

For a concept \mathcal{C} , the cardinality of the characteristic sample constructed for \mathcal{C} in the proof of Theorem 3.2 may be as large as the number of distinct concepts preceding it in the enumeration $\mathcal{C}_1, \mathcal{C}_2, \dots$, which can be exponential in $|\mathcal{C}|$. Barzdin and Freivalds [BF72] proved that the majority vote algorithm could be used in a prediction setting to bound the number of mistakes of prediction linearly in the size of the target concept, which implies a corresponding bound for characteristic samples. (See also the surveys [FBP91] and [ZZ08].) For completeness, we describe their construction in the context of characteristic samples.

First we prove a lemma on using the majority vote algorithm with a finite set of concepts to construct a sample. If \mathbb{D} is a finite nonempty set of concepts, we define the majority vote concept of \mathbb{D} , denoted $\mathcal{C}_{maj}(\mathbb{D})$, to consist of all $x \in \mathcal{X}$ such that the cardinality of the set of concepts in \mathbb{D} containing x is at least as large as the cardinality of the set of concepts in \mathbb{D} not containing x . The concept $\mathcal{C}_{maj}(\mathbb{D})$ may or may not belong to \mathbb{D} or \mathbb{C} . Given a finite nonempty set of concepts \mathbb{D} and another concept \mathcal{C} , we define a sample, denoted $sample(\mathbb{D}, \mathcal{C})$ by executing the halving algorithm with \mathbb{D} as the concept set and \mathcal{C} as the target concept. In detail, initialize $\mathbb{D}_0 = \mathbb{D}$, $i = 0$ and $S_0 = \emptyset$. While \mathbb{D}_i is nonempty, compare $\mathcal{C}_{maj}(\mathbb{D}_i)$ and \mathcal{C} . If these are the same concept, then S_i is output as the value of $sample(\mathbb{D}, \mathcal{C})$. Otherwise, let x be the least element of \mathcal{X} on which they differ, and add the pair (x, b) to S_i to get S_{i+1} , where $b = 1$ if x is in $\llbracket \mathcal{C} \rrbracket$ and $b = 0$ otherwise. \mathbb{D}_{i+1} is set to those concepts in \mathbb{D}_i that are consistent with (x, b) , i is set to $i + 1$, and the while loop continues. If \mathbb{D}_i is empty, then S_i is the value of $sample(\mathbb{D}, \mathcal{C})$.

Lemma 3.4. *If \mathbb{D} is a finite nonempty set of concepts and \mathcal{C} any concept, $sample(\mathbb{D}, \mathcal{C})$ has at most $1 + \log_2(|\mathbb{D}|)$ elements. For any concept \mathcal{C}' in \mathbb{D} , we have that \mathcal{C} and \mathcal{C}' are both consistent with $sample(\mathbb{D}, \mathcal{C}) \cup sample(\mathbb{D}, \mathcal{C}')$ iff $\llbracket \mathcal{C} \rrbracket = \llbracket \mathcal{C}' \rrbracket$.*

⁵The characteristic samples for the class are of cardinality 2, but have no computable bound on their length.

Proof. Because each \mathbb{D}_{i+1} contains at most half as many concepts as \mathbb{D}_i , termination must occur with $i \leq \log_2(|\mathbb{D}|)$, and therefore $|\text{sample}(\mathbb{D}, \mathcal{C})| \leq 1 + \log_2(|\mathbb{D}|)$.

If $\llbracket \mathcal{C} \rrbracket = \llbracket \mathcal{C}' \rrbracket$, the same sample will be returned for each concept, and they will both be consistent with that sample. Conversely, if both are consistent with the union of their final samples, the sets \mathbb{D}_i and S_i will be the same for each i . Because \mathcal{C}' is in \mathbb{D} , termination cannot occur because \mathbb{D}_i becomes empty. Thus, termination must occur because both are equal to $\mathcal{C}_{\text{maj}}(\mathbb{D}_i)$ for some i , which means that they are the same concept. \square

Note in particular that if \mathcal{C} is not equivalent to any concept in \mathbb{D} , then for all \mathcal{C}' in \mathbb{D} the fact that $\llbracket \mathcal{C} \rrbracket \neq \llbracket \mathcal{C}' \rrbracket$ will be witnessed by some example (x, b) in $\text{sample}(\mathbb{D}, \mathcal{C}) \cup \text{sample}(\mathbb{D}, \mathcal{C}')$.

Theorem 3.5. *Let \mathbb{C} be a concept class. There exist a teacher \mathbf{T} and a learner \mathbf{L} such that for every concept representation \mathcal{C} in \mathbb{C} , $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} , and the cardinality of $\mathbf{T}(\mathcal{C})$ is bounded by $O(|\mathcal{C}|)$.*

Proof. The approach is to use the method of Lemma 3.4 on successive finite blocks of concepts of double exponential cardinality. Let \mathbb{B}_0 contain concepts \mathcal{C}_1 and \mathcal{C}_2 . For $i > 0$, let \mathbb{B}_i contain those concepts \mathcal{C}_k such that $2^{2^{i-1}} < k \leq 2^{2^i}$.

Then the teacher \mathbf{T} on input \mathcal{C} finds the least i such that for some k with \mathcal{C}_k in \mathbb{B}_i , $\llbracket \mathcal{C} \rrbracket = \llbracket \mathcal{C}_k \rrbracket$. The sample output for \mathcal{C} is the union of $\text{sample}(\mathbb{B}_j, \mathcal{C})$ for $j = 0, 1, \dots, i$.

The cardinality of the sample output is at most the sum of the cardinalities of the samples $\text{sample}(\mathbb{B}_j, \mathcal{C})$. We have $|\mathbb{B}_0| = 2$ and for $i > 0$

$$|\mathbb{B}_i| = 2^{2^i} - 2^{2^{i-1}} < 2^{2^i}.$$

For all j , the sample $\text{sample}(\mathbb{B}_j, \mathcal{C})$ thus has at most $1 + \log_2(|\mathbb{B}_j|) = 1 + 2^j$ elements. The sum of these bounds through i is

$$(1 + 2^0) + (1 + 2^1) + \dots + (1 + 2^i) = i + 2^{i+1} \quad (3.1)$$

The least i such that \mathbb{B}_i contains \mathcal{C}_k is at most $\log_2 \log_2(k)$ for $k > 1$. So for $k > 1$ by instantiating i with $\log_2 \log_2(k)$ in Equation (3.1) we get that the cardinality of $\mathbf{T}(\mathcal{C})$ is at most $2 \log_2(k) + \log \log(k)$. In the shortlex ordering of finite binary strings, the string \mathcal{C}_k has length at least $\log_2 k$, so the cardinality of $\mathbf{T}(\mathcal{C})$ is bounded by $O(|\mathcal{C}|)$.

The learner \mathbf{L} on input T finds the least \mathcal{C} , if any, such that \mathcal{C} is consistent with T and the sample $\mathbf{T}(\mathcal{C})$ is a subset of T . If such a \mathcal{C} is found, it is output. Otherwise, if there is some concept \mathcal{C} in \mathbb{C} consistent with T , \mathbf{L} outputs the least such, and otherwise outputs an arbitrary element of \mathbb{C} .

To see that $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} , suppose T contains $\mathbf{T}(\mathcal{C})$ and is consistent with \mathcal{C} . Then \mathbf{L} outputs \mathcal{C} unless there is some earlier concept \mathcal{C}' that occurs in an example string of T and is consistent with T and is such that $\mathbf{T}(\mathcal{C}')$ is contained in T . Consider the computation of the sample for \mathcal{C} , and the block \mathbb{B}_j that contains the least concept equivalent to \mathcal{C}' . Because \mathcal{C} and \mathcal{C}' are consistent with $\mathbf{T}(\mathcal{C})$, they are both consistent with $\text{sample}(\mathbb{B}_j, \mathcal{C})$ and because they are consistent with $\mathbf{T}(\mathcal{C}')$, they are consistent with $\text{sample}(\mathbb{B}_j, \mathcal{C}')$, so by Lemma 3.4 applied to the block \mathbb{B}_j , they denote the same concept. \square

We turn to the question of the computability of the teacher and the learner.

Corollary 3.6. *Let \mathbb{C} be a concept class satisfying (C1), (C2) and (C3). There exist a computable teacher \mathbf{T} and a computable learner \mathbf{L} such that for every concept representation \mathcal{C} in \mathbb{C} , $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} , and the cardinality of $\mathbf{T}(\mathcal{C})$ is bounded by $O(|\mathcal{C}|)$.*

Proof. To allow the construction of $sample(\mathbb{B}_i, \mathcal{C})$ to be computable, using (C3), the teacher may test the equivalence of pairs of concepts in \mathbb{B}_i , and pairs consisting of a concept from \mathbb{B}_i and \mathcal{C} , and for each inequivalent pair, may search for the least example x on which they disagree. Then the same computation as in Lemma 3.4 is done, using just these finitely many disagreement strings as the domain.

On input \mathcal{C} , to allow the learner to limit its search to a finite set of candidates, the teacher also adds to the sample some example (x, b) , where $|x| \geq |\mathcal{C}|$ and $b = 1$ iff x is in $\llbracket \mathcal{C} \rrbracket$.⁶ On input T , the learner searches among \mathcal{C} such that $|\mathcal{C}|$ is bounded by the length of the longest string in T to find the least \mathcal{C} consistent with T such that $\mathbf{T}(\mathcal{C})$ is a subset of T . If none is found, then (C2) allows the learner to output some concept consistent with T , if there is one. \square

3.3. Polynomial time learners and teachers. Moving from this general setting, we consider polynomial bounds on the length of characteristic samples and the running times of the teacher and learner. We consider the following possible properties of \mathbb{C} .

- (P1) *Polynomial time example membership.* There is a polynomial time algorithm to decide whether $x \in \llbracket \mathcal{C} \rrbracket$ given x and \mathcal{C} .
- (P2) *Polynomial time default hypothesis construction.* There is a polynomial time algorithm that returns \mathcal{C}_T in \mathbb{C} consistent with a given sample T , or determines that no concept in \mathbb{C} is consistent with T .
- (P3) *Polynomial time equivalence with least counterexamples.* There is a polynomial time algorithm that determines whether $\llbracket \mathcal{C}_i \rrbracket = \llbracket \mathcal{C}_j \rrbracket$ given concept representations \mathcal{C}_i and \mathcal{C}_j . In the case the concepts are not equal, the algorithm also returns the shortest string $x \in \llbracket \mathcal{C}_i \rrbracket \oplus \llbracket \mathcal{C}_j \rrbracket$.

A class \mathbb{C} is *concisely distinguishable* if there exists a polynomial $p(n)$ such that for every pair \mathcal{C}_i and \mathcal{C}_j such that $\llbracket \mathcal{C}_i \rrbracket \neq \llbracket \mathcal{C}_j \rrbracket$, there exists a string $x \in \llbracket \mathcal{C}_i \rrbracket \oplus \llbracket \mathcal{C}_j \rrbracket$ such that $|x| \leq p(\text{size}(\mathcal{C}_1) + \text{size}(\mathcal{C}_2))$.

Concise characteristic samples. A class \mathbb{C} *has concise characteristic samples* if there exist a polynomial $p(n)$, a teacher \mathbf{T} , and a learner \mathbf{L} such that for every $\mathcal{C} \in \mathbb{C}$, $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} , and $|\mathbf{T}(\mathcal{C})| \leq p(\text{size}(\mathcal{C}))$. Recall that the length of a sample is the sum of length of the strings in it. In Corollary 3.6 we considered just the cardinality of $\mathbf{T}(\mathcal{C})$ and not the length of the words in the sample, but in $|\mathbf{T}(\mathcal{C})|$ the length of the words in the sample matters. Note also that the polynomial bound is in terms of the size of the smallest representation of $\llbracket \mathcal{C} \rrbracket$. If \mathbb{C} has concise characteristic samples then \mathbb{C} is concisely distinguishable, by Lemma 3.1. In the converse direction, we have the following consequence of Corollary 3.6.

Corollary 3.7. *Assume the class \mathbb{C} satisfies (C1), (C2) and (C3) and is concisely distinguishable. Then it has concise characteristic samples for a computable teacher \mathbf{T} and a computable learner \mathbf{L} .*

Proof. Because \mathbb{C} is concisely distinguishable, two concepts \mathcal{C}_i and \mathcal{C}_j can be tested for equivalence by checking agreement on all strings of length at most a fixed polynomial in $\text{size}(\mathcal{C}_i) + \text{size}(\mathcal{C}_j)$. If they are inequivalent, this process will yield a distinguishing example

⁶See the discussion following Theorem 3.10.

of at most that length. If this method is used in the algorithm for constructing characteristic samples in the proof of Corollary 3.6, the resulting characteristic samples will be concise. \square

Efficient teachability. We say that \mathbb{C} is *efficiently teachable* if there exist a polynomial time teacher \mathbf{T} and a learner \mathbf{L} such that $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} for every $\mathcal{C} \in \mathbb{C}$.

Lemma 3.8. *If class \mathbb{C} is efficiently teachable, then \mathbb{C} has concise characteristic samples.*

Proof. Let \mathbf{T} be a polynomial time teacher and \mathbf{L} a learner witnessing the fact that \mathbb{C} is efficiently teachable. To see that \mathbb{C} has concise characteristic samples, we define the (not necessarily polynomial time) teacher \mathbf{T}' as follows. On input \mathcal{C} , let \mathcal{C}' minimize $size(\mathcal{C}')$ subject to $\llbracket \mathcal{C}' \rrbracket = \llbracket \mathcal{C} \rrbracket$. Then \mathbf{T}' outputs $\mathbf{T}(\mathcal{C}')$, which is a characteristic sample for \mathcal{C} and \mathbf{L} and is of size bounded by a polynomial in $size(\mathcal{C})$. \square

Efficient learnability. We say that \mathbb{C} is *efficiently learnable* if there exist a teacher \mathbf{T} and a polynomial time learner \mathbf{L} such that $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} for every $\mathcal{C} \in \mathbb{C}$.

As has been observed before [Pit89], if (P1) and (P2) are satisfied, this definition is superfluous because the learning algorithm of identification by enumeration used in the proof of Theorem 3.2 can be modified to run in polynomial time and the teacher adjusted appropriately to create a (potentially ridiculously large) characteristic sample. This observation motivates the introduction of the criterion of identification in the limit with polynomial time and data (described in Section 3.4). The argument for this observation follows.

Lemma 3.9. *If class \mathbb{C} satisfies (P1) and (P2) then \mathbb{C} is efficiently learnable.*

Proof. The learner \mathbf{L} on input T of length n simulates the algorithm of identification by enumeration for n^2 steps. It checks whether the last hypothesis \mathcal{C} output by the simulation is consistent with T and outputs \mathcal{C} if so, using (P1). Otherwise, it returns the default hypothesis consistent with T , if any, using (P2). Given \mathcal{C} in \mathbb{C} , the teacher \mathbf{T} takes the characteristic sample $T_{\mathcal{C}}$ from Theorem 3.2 and adds enough additional examples of \mathcal{C} that there is time for \mathbf{L} 's simulation of identification by enumeration to converge to its final answer. Note that since the domain \mathcal{X} is infinite it is always possible to add more labeled words to the sample. \square

Requiring both teacher and learner to be efficient. We say that \mathbb{C} is *efficiently teachable/learnable* if there exist a polynomial time teacher \mathbf{T} and a polynomial time learner \mathbf{L} such that $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} for every $\mathcal{C} \in \mathbb{C}$. This corresponds to the definition of *polynomially T/L teachable* of Goldman and Mathias [GM96].

It turns out that if (P1) and (P2) are satisfied, efficiently teachable/learnable is no stronger than requiring a polynomial time teacher.

Theorem 3.10. *Assume that class \mathbb{C} satisfies (P1) and (P2) and is efficiently teachable. Then \mathbb{C} is efficiently teachable/learnable.*

Proof. Let \mathbf{T} be polynomial time teacher witnessing the efficient teachability of \mathbb{C} . We define another polynomial time teacher \mathbf{T}' that on input \mathcal{C} outputs $\mathbf{T}(\mathcal{C})$ together with one pair (x, b) such that x is the string \mathcal{C} and (x, b) is consistent with \mathcal{C} . Note that the teacher directly provides the *text* of the target concept.

We define a learner \mathbf{L} as follows. On input T , \mathbf{L} finds the least x (if any) such that (x, b) is in T , the concept representation $\mathcal{C} = x$ is such that \mathcal{C} is consistent with T (using (P1)), and the sample $\mathbf{T}(\mathcal{C})$ is a subset of T (using the polynomial time algorithm \mathbf{T}). If such a \mathcal{C} is found, \mathbf{L} outputs it. Otherwise (using (P2)) it either outputs a default concept \mathcal{C} consistent with T or an arbitrary concept. Then \mathbf{L} runs in time polynomial in $|T|$.

The argument that for every \mathcal{C} in \mathbb{C} , $\mathbf{T}'(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} is the same as that in the proof of Corollary 3.6. \square

In this proof, the teacher provides the text of the target concept itself in the sample and the learner depends on this fact and considers as potential candidates only those concepts whose strings are in the sample. This may seem like a kind of unacceptable “collusion” between the teacher and the learner. In the proof of Corollary 3.6, the teacher provides the learner with less specific syntactic information about the target concept in the form of a bound on its length. Gold’s founding positive result in the area of characteristic samples may be stated as follows.

Theorem 3.11 [Gol78, Theorem 4]. *DFAs are efficiently teachable.*

In Gold’s original proof, we see a high degree of coordination between the teacher and the learner to enable the teacher to provide all the examples that the learner will consult in constructing its hypothesis, even when additional correct examples may be present.

The issue of what kinds of coordination and communication should be permitted in a model of teaching and learning is complex, and has been considered by a number of researchers. One paradigm requires the teacher to be able to teach any learner (or more precisely, any *consistent learner*, i.e., any learner that never hypothesizes a concept inconsistent with the information it has). Goldman and Mathias, in their seminal paper [GM96], show that this requirement is too strong in the sense that it makes even rather simple classes of concepts very hard to teach. Moreover, some applications call for a learning paradigm in which a teacher is required to teach a particular learner rather than an arbitrary one.

The paradigm introduced by Goldman and Mathias [GM96], which is the one we follow here, addresses the problem of collusion by allowing an adversary to add an arbitrary set of correctly labeled examples to the sample generated by the teacher before it is given to the learner. For a class \mathbb{C} and a deterministic teacher \mathbf{T} and a deterministic learner \mathbf{L} , they define \mathbf{T} and \mathbf{L} to be a *colluding pair* if there exist \mathcal{C}_i and \mathcal{C}_j such that $\llbracket \mathcal{C}_i \rrbracket = \llbracket \mathcal{C}_j \rrbracket$ and for all samples $T_i \supseteq \mathbf{T}(\mathcal{C}_i)$ and $T_j \supseteq \mathbf{T}(\mathcal{C}_j)$ consistent with \mathcal{C}_i and \mathcal{C}_j , we have $\mathbf{L}(T_i) \neq \mathbf{L}(T_j)$. That is, in a colluding pair, the teacher is able to communicate some distinguishing information about which of the two concepts, \mathcal{C}_i or \mathcal{C}_j , was its input. They prove that under their paradigm there is no colluding pair. Of course, one could consider other notions of collusion.

To see that the possibility of including the text of the target concept among the examples does not trivialize the problem of creating a characteristic sample, note that the learner cannot a priori know which examples in the sample correspond to information from the teacher and which were added by the adversary. Hence, to determine the true concept the learner must rule out possibly spurious examples that the adversary has added. The only information available to the learner, other than the strings in the sample, are their

associated labels, indicating their membership in the target concept. Hence, the learner must use semantic information about the target concept to eliminate spurious candidates, which is what we would like from such a learning paradigm.

The efficient teachability results we prove starting in Section 5 in fact include descriptions of algorithms for both the teacher and the learner, and do *not* include the text of the target concept in the sample. Like Gold’s founding result, they do exhibit a high degree of coordination between the teacher and learner.

3.4. Identification in the limit with polynomial time and data. We relate the definitions above to the definition of identification in the limit using polynomial time and data introduced by Gold [Gol78] and refined by de la Higuera [dlH97], who also showed that it is closely related to the definition of *semi-poly T/L teachable* introduced by Goldman and Mathias [GM96].

Using the terminology of this paper, a class \mathbb{C} is *identifiable in the limit using polynomial time and data* if there exist a polynomial $p(n)$, a teacher \mathbf{T} , and a polynomial time learner \mathbf{L} such that for every $\mathcal{C} \in \mathbb{C}$, the sample $\mathbf{T}(\mathcal{C})$ is a characteristic sample for \mathcal{C} and \mathbf{L} , and $|\mathbf{T}(\mathcal{C})| \leq p(\text{size}(\llbracket \mathcal{C} \rrbracket))$. This definition can be viewed as correcting the deficiency of the concept of efficiently learnable (indicated by Lemma 3.9) by requiring a polynomial bound on the size of the characteristic sample. In this terminology, “polynomial time” refers to the polynomial running time of \mathbf{L} , and “polynomial data” refers to the polynomial bound on the size of the sample $\mathbf{T}(\mathcal{C})$. Of course, the latter is not a worst-case measure; there could be arbitrarily large finite samples for which \mathbf{L} outputs an incorrect hypothesis.

Because the definition of \mathbb{C} being identifiable in the limit using polynomial time and data simply adds the requirement that the learner be computable in polynomial time to the definition of \mathbb{C} having concise characteristic samples, we immediately have the following.

Lemma 3.12. *If \mathbb{C} is identifiable in the limit using polynomial time and data then \mathbb{C} has concise characteristic samples.*

Comparing with the concept of being efficiently teachable, we note the following differences. Identifiability in the limit using polynomial time and data only requires the *existence* of a characteristic sample, and the bound on the length of the characteristic sample is in terms of the size of the *smallest* representation of $\llbracket \mathcal{C} \rrbracket$. For the definition of efficiently teachable, the characteristic sample must not only exist, but be computable in polynomial time, and the bound on the length of the characteristic sample is in terms of the length of \mathcal{C} .

Lemma 3.13. *If \mathbb{C} satisfies (P1) and (P2) and is efficiently teachable, then \mathbb{C} is identifiable in the limit using polynomial time and data.*

Proof. By Theorem 3.10, there are a polynomial time teacher \mathbf{T} and a polynomial time learner \mathbf{L} witnessing that \mathbb{C} is efficiently teachable/learnable. Let $\mathcal{C} \in \mathbb{C}$ be given. Let $\mathcal{C}' \in \mathbb{C}$ minimize $|\mathcal{C}'|$ subject to $\llbracket \mathcal{C}' \rrbracket = \llbracket \mathcal{C} \rrbracket$. Then $\mathbf{T}(\mathcal{C}')$ is a characteristic sample for \mathcal{C} and \mathbf{L} of length polynomial in $\text{size}(\mathcal{C})$. Thus \mathbb{C} is identifiable in the limit using polynomial time and data. \square

To see that identifiability in the limit using polynomial time and data may not imply efficient teachability, we consider the following example.

Theorem 3.14. *Assume that there is no polynomial time algorithm for integer factorization. Then there exists a concept class that satisfies (P1) and (P2) and is identifiable in the limit using polynomial time and data but is not efficiently teachable.*

Proof. We assume a standard decimal representation of positive integers. For any positive integer n , let $F(n)$ be the list of primes in its prime factorization in non-decreasing order, for example, $F(54) = (2, 3, 3, 3)$. Let \mathbb{N}_+ denote the set of positive integers, and let $F(\mathbb{N}_+) = \{F(n) \mid n \in \mathbb{N}_+\}$, the set of all prime factorization lists. If ℓ is a prime factorization list (p_1, p_2, \dots, p_k) then let $M(\ell)$ denote the product of the primes in ℓ , that is, $M(\ell) = p_1 \cdot p_2 \cdots p_k$. M is computable in polynomial time.

We define a class of concepts as follows. \mathbb{C}_{fact} has domain $F(\mathbb{N}_+)$ and consists of all finite subsets of $F(\mathbb{N}_+)$. We represent a finite set $\{\ell_1, \ell_2, \dots, \ell_m\} \subset F(\mathbb{N}_+)$ as a finite list (n_1, n_2, \dots, n_m) such that $n_i = M(\ell_i)$. Thus, $\{(2, 3, 3, 3)\}$ is represented by the string (54), and $\{(3), (2, 3), (3, 5, 5)\}$ is represented by the string (3, 6, 75).

Because primality can be decided in polynomial time [AKS04], there is a polynomial time algorithm to test whether $F(n) = (p_1, p_2, \dots, p_k)$ given n and (p_1, p_2, \dots, p_k) as inputs. Therefore (P1) holds for \mathbb{C}_{fact} .

Consider the learner \mathbf{L} that on input a finite sample T first checks that every example is a prime factorization list and that no list is given two different labels. (If this check fails, the output is arbitrary.) Let T_1 denote the set of examples in T with label 1. \mathbf{L} outputs a representation of the concept $T_1 = \{\ell_1, \ell_2, \dots, \ell_m\}$ as $(M(\ell_1), M(\ell_2), \dots, M(\ell_m))$. Thus property (P2) holds and \mathbb{C}_{fact} is learnable in polynomial time and has characteristic samples of polynomial length, consisting of exactly the finitely many positive examples of each concept. Thus, \mathbb{C}_{fact} is identifiable in the limit using polynomial time and data.

Now suppose for the sake of contradiction that there is a polynomial time teacher \mathbf{T} for \mathbb{C}_{fact} . Consider the sample produced by \mathbf{T} for the concept $()$, which denotes the empty set. This consists of a finite number of negative examples, say $\ell_1, \ell_2, \dots, \ell_m$, where each ℓ_i is a prime factorization list. Let $S = \{M(\ell_i) \mid 1 \leq i \leq m\}$. Consider any positive integer $n \notin S$, and let $F(n) = (p_1, p_2, \dots, p_k)$. If the sample $\mathbf{T}((n))$ does not contain $F(n)$ as a positive example, it must consist exclusively of negative examples, and both the empty set and $\{F(n)\}$ are consistent with the union of the samples produced by \mathbf{T} for the two concepts, contradicting Lemma 3.1. Thus, for all but finitely many positive integers n , \mathbf{T} must produce the prime factorization of n . This implies that there is a polynomial time algorithm for prime factorization, contradicting our assumption. \square

3.5. Relation to learning with equivalence and membership queries. In the paradigm of *learning with membership and equivalence queries*, a learning algorithm can access an oracle that truthfully answers two types of queries about the target concept \mathcal{C} , and its goal is to halt and output a representation of $\llbracket \mathcal{C} \rrbracket$. In a *membership query*, or MQ, the learning algorithm provides a string x and the answer is 1 or 0 depending on whether $x \in \llbracket \mathcal{C} \rrbracket$ or not. In an *equivalence query*, or EQ, the learning algorithm provides a representation $\mathcal{C}' \in \mathbb{C}$, and the answer from the oracle is either “yes”, if $\llbracket \mathcal{C}' \rrbracket = \llbracket \mathcal{C} \rrbracket$, and otherwise is an arbitrarily chosen element of $\llbracket \mathcal{C}' \rrbracket \oplus \llbracket \mathcal{C} \rrbracket$ (a *counterexample* to the conjecture that \mathcal{C}' is correct). If the learning algorithm successfully learns every $\mathcal{C} \in \mathbb{C}$ and at every point its running time is bounded by a polynomial in $size(\llbracket \mathcal{C} \rrbracket)$ and the length of the longest counterexample seen to that point, we say that \mathbb{C} is polynomially learnable using membership and equivalence queries.

Goldman and Mathias [GM96] prove that any class that can be learned by a deterministic polynomial time algorithm using any of a large set of example-based queries is teachable by

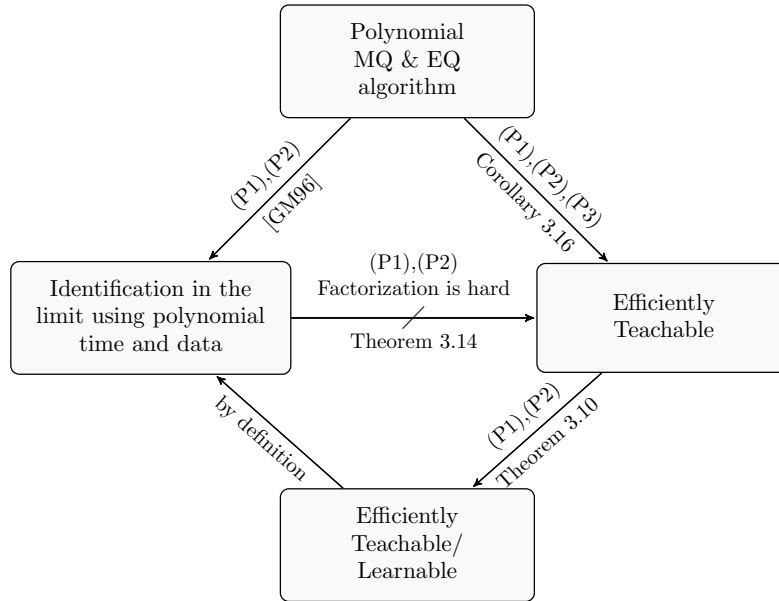


Figure 1. Summary of main general results about efficient teachability

a computationally unbounded teacher and a polynomial time learner. A corollary of their Theorem 2 is the following.

Theorem 3.15. *Suppose the class \mathbb{C} satisfies properties (P1) and (P2) and is concisely distinguishable. If \mathbb{C} is polynomially learnable using membership and equivalence queries, it is also identifiable in the limit using polynomial time and data.*

(This result is also proved by Bohn and Löding [BL21].) Can this result be strengthened to conclude that \mathbb{C} is polynomially teachable? To answer this question, we examine the proof in more detail. Let \mathbf{A} be a learning algorithm using membership and equivalence queries that learns \mathbb{C} in polynomial time. Given \mathcal{C} , \mathbf{T} constructs a sample $T_{\mathcal{C}}$ by simulating \mathbf{A} and answering its queries according to \mathcal{C} as follows. A membership query with x is answered by determining whether $x \in \llbracket \mathcal{C} \rrbracket$. For an equivalence query with \mathcal{C}' , if $\llbracket \mathcal{C}' \rrbracket \neq \llbracket \mathcal{C} \rrbracket$, then x is chosen to be the shortlex least element of $\llbracket \mathcal{C}' \rrbracket \oplus \llbracket \mathcal{C} \rrbracket$ and returned as the counterexample to the simulation of \mathbf{A} .

The counterexample x is of length polynomial in the sum of the sizes of \mathcal{C}' and \mathcal{C} by the assumption of concise distinguishability. If instead $\llbracket \mathcal{C}' \rrbracket = \llbracket \mathcal{C} \rrbracket$, then the sample $T_{\mathcal{C}}$ is constructed of all the strings x that appeared in membership queries or as counterexamples returned to equivalence queries during the simulation, labeled to be consistent with \mathcal{C} . Because of the polynomial running time of \mathbf{A} and the choice of shortest counterexamples, the length of $T_{\mathcal{C}}$ is bounded by a polynomial in $size(\llbracket \mathcal{C} \rrbracket)$.

The corresponding learning algorithm \mathbf{L} takes a sample T as input and simulates the learning algorithm \mathbf{A} , attempting to answer its queries using T as follows. For a membership query with x , if x is an example in T , the answer is its label in T . If x is not an example in T , then \mathbf{L} outputs a default \mathcal{C}_T consistent with T and halts (using property (P2)). For an equivalence query with \mathcal{C} , \mathbf{L} checks whether \mathcal{C} is consistent with T (using property (P1)). If it is consistent, then it outputs \mathcal{C} and halts. If it is not consistent, it finds the shortlex least x that is an example in T whose label is not consistent with \mathcal{C} and returns x as the

counterexample to \mathbf{A} 's equivalence query. The running time of \mathbf{L} is polynomial in the length of T . It is because of the choice of the shortlex counterexample by both \mathbf{T} and \mathbf{L} that \mathbf{T} can anticipate exactly the queries that will be made in the simulation of \mathbf{A} by \mathbf{L} , even when the sample T is a superset of the characteristic sample $T_{\mathcal{C}}$.

What must we assume in order that the function \mathbf{T} in this construction can be computed in polynomial time? Property (P1) allows \mathbf{T} to answer membership queries in polynomial time, and property (P3) ensures that \mathcal{C} is concisely distinguishable, and allows \mathbf{T} to find shortlex least counterexamples in polynomial time, so we have the following.

Corollary 3.16. *Suppose the class \mathcal{C} satisfies properties (P1), (P2) and (P3). If \mathcal{C} is polynomially learnable using membership and equivalence queries, then it is also efficiently teachable.*

Figure 1 provides a summary of the main general results about efficient teachability.

4. FOCUSING ON REGULAR ω -LANGUAGES

The rest of the paper concerns efficient teachability of regular ω -languages. This section starts by describing examples and samples for ω -languages, and continues by describing some known or immediate results on the subject.

4.1. Examples and samples for ω -languages. Because we require finite representations of examples, ω -words in our case, we work with ultimately periodic words, that is, words of the form $u(v)^\omega$ where $u \in \Sigma^*$ and $v \in \Sigma^+$. It is known that two regular ω -languages are equivalent iff they agree on the set of ultimately periodic words [Büc62, CNP93], so this choice is not limiting.

The example $u(v)^\omega$ is concretely represented by the pair (u, v) of finite strings, and its length is $|u| + |v|$. A *labeled example* is a pair $(u(v)^\omega, l)$, where the label l is either 0 or 1. A *sample* is a finite set of labeled examples such that no example is assigned two different labels. The *length* of a sample is the sum of the lengths of the examples that appear in it. A sample T and a language L are *consistent* with each other if and only if for every labeled example $(u(v)^\omega, l) \in T$, $l = 1$ iff $u(v)^\omega \in L$. A sample T and an acceptor \mathcal{A} are *consistent* with each other if and only if T is consistent with $\llbracket \mathcal{A} \rrbracket$. The following results give two useful procedures on examples that are computable in polynomial time.

Proposition 4.1. *Let $u_1, u_2 \in \Sigma^*$ and $v_1, v_2 \in \Sigma^+$. If $u_1(v_1)^\omega \neq u_2(v_2)^\omega$ then they differ in at least one of the first ℓ symbols, for $\ell = \max(|u_1|, |u_2|) + |v_1| \cdot |v_2|$.*

Let $\text{suffixes}(u(v)^\omega)$ denote the set of all ω -words that are suffixes of $u(v)^\omega$.

Proposition 4.2. *The set $\text{suffixes}(u(v)^\omega)$ consists of at most $|u| + |v|$ different examples: one of the form $u'(v)^\omega$ for every nonempty suffix u' of u , and one of the form $(v_2v_1)^\omega$ for every division of $v = v_1v_2$ into a non-empty prefix v_1 and suffix v_2 .*

4.2. Negative results for nondeterministic classes. The classes NBA, NPA, NMA, NCA do not have concise characteristic sets [AFS20]. The proof is by constructing a family of languages $\{L_n\}_{n \in \mathbb{N}}$ with an acceptor of size quadratic in n for which at least one word of length at least exponential in n must be included in any characteristic sample for L_n .⁷ Since an NBA (resp. NCA) is a special case of NRA (resp. NSA) the same is true for NRA and NSA.

Theorem 4.3. *The classes NBA, NPA, NMA, NCA, NRA and NSA do not have concise characteristic sets, and therefore are neither identifiable in the limit using polynomial time and data nor efficiently teachable.*

4.3. Consequences of membership and equivalence algorithms. In the domain of ω -automata, researchers have so far found very few polynomial time learning algorithms using membership and equivalence queries. We consider the cases of Mod 2 multiplicity automata, strongly unambiguous Büchi automata, and deterministic weak parity automata.

A Büchi automaton is *unambiguous* if no word has more than one run starting in the initial state and visiting an accepting state infinitely often. It is *strongly unambiguous* if no word has more than one run visiting an accepting state infinitely often, whether it starts at the initial state or not. For instance, consider the deterministic Büchi automaton $\mathcal{B} = (\{a, b\}, \{q_a, q_b\}, q_a, \delta, \{q_b\})$ with $\delta(q_a, a) = \delta(q_b, a) = q_a$ and $\delta(q_a, b) = \delta(q_b, b) = q_b$. Since \mathcal{B} is deterministic there is a unique run on every word, and thus it is also unambiguous. However, since the word $(b)^\omega$ is accepted from both q_a and q_b , \mathcal{B} is not strongly unambiguous.

Angluin et al. [AAF20] give a polynomial time mapping r of a strongly unambiguous Büchi automata (SUBA) \mathcal{C} to a representation $r(\mathcal{C})$ as a Modulo-2 multiplicity automaton (M2MA), and observe that there is a polynomial time algorithm for learning M2MAs using membership and equivalence queries [BBB⁺00]. (Please see [AAF20] for precise definitions.) We note that a shortlex least counterexample can be returned in case of inequivalence [DKV09, Sak09]. It follows from Corollary 3.16 that M2MAs are efficiently teachable. We then also have the following.

Corollary 4.4. *The class SUBA is efficiently teachable.*

Proof. Let \mathbf{T} be a teacher witnessing the efficient teachability of M2MAs. The teacher \mathbf{T}' with input a SUBA \mathcal{C} generates the characteristic sample $\mathbf{T}(r(\mathcal{C}))$, to which it adds one example (x, b) such that $x = \mathcal{C}$ and $b = 1$ iff $x \in \llbracket \mathcal{C} \rrbracket$. The learner \mathbf{L} with input T searches for the least x (if any) such that $(x, b) \in T$, $r(x)$ is consistent with T , and $\mathbf{T}(r(x)) \subseteq T$. If such an x is found, it is output; otherwise, \mathbf{L} constructs and outputs a SUBA consistent with T . A procedure constructing a default SUBA acceptor that agrees with a given sample T is given in the proof of Proposition 5.2. \square

Maler and Pnueli [MP95] give an algorithm that learns the class $\mathbb{D}\text{WPA}$ of deterministic weak parity automata in polynomial time using membership and equivalence queries. The weak parity condition is obtained from the parity condition using $\text{occ}(\rho)$ instead of $\text{inf}(\rho)$ where $\text{occ}(\rho)$ is the set of states visited somewhere during the run ρ . It is known that $\mathbb{D}\text{WPA} = \mathbb{D}\text{BA} \cap \mathbb{D}\text{CA}$. Membership and equivalence of $\mathbb{D}\text{WPA}$ are decidable in polynomial time, thus by Theorem 3.15, $\mathbb{D}\text{WPA}$ is identifiable in the limit using polynomial time and

⁷A negative result regarding query learning of NBA, NPA and NMA was obtained by Angluin et al. [AAF20]. That result makes a plausible assumption of cryptographic hardness, which is not required here.

data. It is not known whether there is a polynomial time algorithm for equivalence with shortlex counterexamples, so we are unable to apply Corollary 3.16 to deduce they are also efficiently teachable. However, DWPA are a special case of DPAs and the following is a corollary of Theorem 16.8 for DPAs.

Corollary 4.5. *The class DWPA is efficiently teachable.*

The learning algorithm of Maler and Pnueli for DWPA exploits the fact that $\mathbb{D}WPA = \mathbb{D}BA \cap \mathbb{D}CA = \mathbb{I}BA \cap \mathbb{I}CA$, that is, this class is fully informative and does have a one-to-one relationship between states of the minimal DWPA for a language L and the equivalence classes of the right congruence \sim_L . The class DWPA is a small sub-class of the class of the fully informative regular ω -languages — there are fully informative languages in every level of the Wagner hierarchy [AF18], whereas DWPA is one of the lowest levels in the hierarchy. The rest of the paper is dedicated to showing that fully informative languages of any of the considered ω -automata types (Büchi, coBüchi, parity, Muller, Rabin and Streett) are efficiently teachable.

5. THE INFORMATIVE CLASSES ARE EFFICIENTLY TEACHABLE

This section covers some preliminary issues and gives an overview of the milestones needed to prove that the informative classes are efficiently teachable.

5.1. Duality. There are reductions of the problem of efficient teachability between $\mathbb{I}BA$ and $\mathbb{I}CA$ and between $\mathbb{I}RA$ and $\mathbb{I}SA$, using the duality between these types of acceptors. Consequently we focus on the classes $\mathbb{I}BA$, $\mathbb{I}PA$, $\mathbb{I}RA$ and $\mathbb{I}MA$ in what follows.⁸

Proposition 5.1. *$\mathbb{I}BA$ (resp., $\mathbb{I}RA$) is efficiently teachable if and only if $\mathbb{I}CA$ (resp., $\mathbb{I}SA$) is.*

Proof. Let \mathcal{A} be an $\mathbb{I}CA$. Because \mathcal{A} is deterministic and complete, if we let \mathcal{A}' denote the $\mathbb{I}BA$ with the same components as \mathcal{A} , then \mathcal{A}' accepts the complement of the language \mathcal{A} , by Claim 2.2 (4).

We modify the characteristic sample for \mathcal{A}' by complementing all its labels to get a characteristic sample for \mathcal{A} . The algorithm to learn an $\mathbb{I}CA$ from a sample T is obtained by complementing all the labels in the sample T and calling the algorithm to learn an $\mathbb{I}BA$ from a sample. The resulting $\mathbb{I}BA$, now considered to be an $\mathbb{I}CA$, is returned as the answer.

The same conversion may be done with acceptors of types $\mathbb{I}RA$ and $\mathbb{I}SA$, by Claim 2.2 (5). \square

⁸The results regarding the classes $\mathbb{I}BA$ (and $\mathbb{I}CA$), $\mathbb{I}PA$ and $\mathbb{I}MA$ were obtained in [AFS20]; here we extend them to the classes $\mathbb{I}RA$ (and $\mathbb{I}SA$). Results for identifiability in the limit using polynomial time and data (but not efficient teachability) of the classes $\mathbb{I}RA$ (and $\mathbb{I}SA$) have also been provided in [BL21] using a different algorithm.

$$T = \{((a)^\omega, 0), (a(b)^\omega, 1), ((ab)^\omega, 0), (ab(baa)^\omega, 1), (bab(aab)^\omega, 1), ((aab)^\omega, 1), (bb(aba)^\omega, 1)\}$$




Figure 2. Top: a sample T . Middle: The distinguishing prefixes U of words in the sample T . Left: Default acceptor of type DBA for sample T (the dead state and the transitions to it are omitted). Right: Default acceptor of type SUBA for sample T .

5.2. The default acceptor. One condition of the definition of being efficiently teachable is that the learning algorithm must run in polynomial time and return an acceptor of the required type that is consistent with the input sample T , even if the sample T does not subsume a characteristic sample. To meet this condition, we use the strategy of Gold’s construction, that is, the learning algorithm optimistically assumes that the sample includes a characteristic sample, and if that assumption fails to produce an acceptor consistent with the sample, the algorithm instead produces a *default acceptor* to ensure that its hypothesis is consistent with the sample. Alternatively, one can use Bohn and Löding’s generalization of the RPNI algorithm to learning ω -automata, which has a more complex default strategy [BL21].

The construction of the default acceptor is given in the proof of the following proposition, and is accompanied by an example illustrated in Figure 2.

Proposition 5.2. *There is a polynomial time algorithm that takes a sample T and returns a DBA (resp., DCA, DPA, DRA, DSA, DMA, SUBA) consistent with T .*

Proof. Given a sample T , we initialize U to be the empty set, and for every word w_i in the sample, we find the shortest prefix u_i of w_i that distinguishes it from all other examples in T and add it to U (see Figure 2, middle). We arrange the finite words in U in a trie in the usual manner. We add self-loops on each $\sigma \in \Sigma$ to the leaves of the trie (see Figure 2, left). This deterministic automaton is termed the prefix-tree automaton [OG92]. If the automaton is incomplete, we add a new dead state with self-transitions on each $\sigma \in \Sigma$, and define all undefined transitions to go to the dead state. Recall that by Proposition 4.1 the length of a prefix distinguishing two examples $u_1(v_1)^\omega$ and $u_2(v_2)^\omega$ is polynomially bounded by the

length of the examples. It follows that the prefix-tree automaton can be constructed in time polynomial in the length of the sample T .

For a DBA, the acceptance condition F consists of all the trie-leaf states that are prefixes of positive examples in T (see Figure 2, left). For a DMA, the acceptance condition consists of $\{\{q\} \mid q \in F\}$. For a DCA, the acceptance condition consists of the dead state (if one was added) and all the trie-leaf states that are not prefixes of positive examples in T . The DBA thus constructed may be transformed to a DPA or a DRA using Claim 2.2 (1) or Claim 2.2 (2), respectively, and the DCA may be transformed to a DSA using Claim 2.2 (3).

For a SUBA the construction is different. We first construct a set V of shortest periods of positively labeled examples, by going iteratively over the examples w_1, w_2, \dots and proceeding as follows. For $w_i = u(v)^\omega$ if v_i is a shortest period of w_i , and none of the rotations of v_i is in V we add v_i to V . Then for every period $v = \sigma_1\sigma_2\dots\sigma_k$ in V we construct an automaton with k states s_0, s_1, \dots, s_{k-1} and transitions $(s_i, \sigma_{i+1}, s_{i'})$ for $0 \leq i < k$ where $i' = (i + 1) \bmod k$. We designate s_0 as an accepting state. Finally, for each positively labeled word w we look for the shortest prefix x of w such that $w = x(y)^\omega$ for some rotation $y = \sigma_i\dots\sigma_k\sigma_1\dots\sigma_{i-1}$ of a period $v = \sigma_1\sigma_2\dots\sigma_k \in V$. For a rotation y of $v \in V$, let X_y be the set of such shortest prefixes. The prefixes in X_y are arranged in a suffix-sharing trie and the trie is connected to the automata constructed for the periods by landing in the state reading σ_i in the automaton of v . Each node of this trie corresponding to a start of a prefix in X_y is added to the set of initial states (see Figure 2, right). It is easy to see that the constructed NBA is consistent with the sample. To see that it is a SUBA, consider a word $w \in \Sigma^\omega$ and assume $w = xy^\omega$ for y a shortest period of w and x the shortest prefix reaching such y . Note that there is only one such representation of w . By the construction of the SUBA, w can only be accepted via a cycle reading y and there is only one such cycle, moreover the cycle can be entered only at the position after reading x , and by the suffix-sharing trie there is only one state from which reading x gets to this position. \square

5.3. Strongly connected components. The acceptance conditions that we consider are all based on the set of states visited infinitely often in a run of the automaton on an input $w \in \Sigma^\omega$. We consider only acceptors whose automata are deterministic and complete, so for any $w \in \Sigma^\omega$ there is exactly one run, which we denote $\rho(w)$, of the automaton on input w . Thus we may define $\text{inf}(w) = \text{inf}(\rho(w))$, the set of states visited infinitely often in this unique run. In the run $\rho(w)$, there is some point after which none of the states visited finitely often is visited. Because each state in $\text{inf}(w)$ is visited infinitely often, for any states $q_1, q_2 \in \text{inf}(w)$, there exists a non-empty word $x \in \Sigma^*$ such that $\delta(q_1, x) = q_2$ and for each prefix x' of x , $\delta(q_1, x') \in \text{inf}(w)$, that is, the path from q_1 to q_2 on x does not visit any state outside the set $\text{inf}(w)$.

These properties motivate the following definition. Given an automaton \mathcal{M} , a *strongly connected component* (SCC) of \mathcal{M} is a nonempty set of states C such that for every $q_1, q_2 \in C$, there exists a nonempty string $x \in \Sigma^*$ such that $\delta(q_1, x) = q_2$ and for any prefix x' of x , $\delta(q_1, x') \in C$.

Note that an SCC need not be maximal, and that a singleton state set $\{q\}$ is an SCC if and only if the state q has a self-loop, that is, $\delta(q, \sigma) = q$ for some $\sigma \in \Sigma$. There is a close relationship between SCCs and the set of states visited infinitely often in a run.

Proposition 5.3. *Let \mathcal{M} be a complete deterministic automaton and $w \in \Sigma^\omega$. Then $\text{inf}(w)$ is an SCC of \mathcal{M} . If w is the ultimately periodic word $u(v)^\omega$, then $\text{inf}(w)$ may be computed in time polynomial in the size of \mathcal{M} and the length of $u(v)^\omega$.*

Proposition 5.4. *For any deterministic automaton $\mathcal{M} = \langle \Sigma, Q, q_\iota, \delta \rangle$ and any reachable SCC C of \mathcal{M} , there exists an ultimately periodic word $w = u(v)^\omega$ of length at most $|Q| + |C|^2$ such that $C = \text{inf}(w)$. Such a word may be found in time polynomial in $|Q|$ and $|\Sigma|$.*

Proof. Because C is reachable, a word $u \in \Sigma^*$ of minimum length such that $\delta(q_\iota, u) \in C$ may be found by breadth first search. The length of u is at most $|Q|$. If $C = \{q\}$, then there is at least one symbol $\sigma \in \Sigma$ such that $\delta(q, \sigma) = q$. Then the ω -word $w = u(\sigma)^\omega$ is such that $C = \text{inf}(w)$. The length of this ultimately periodic word is at most $|Q| + 1$.

If C contains at least two states, let q_1, \dots, q_k be the states in C that are not q . Then for each i , there exist two nonempty finite words x_i and y_i each of length at most n such that $\delta(q, x_i) = q_i$ and $\delta(q_i, y_i) = q$, and the path on x_i from q to q_i and the path on y_i from q_i to q do not visit any states outside of C . The words x_i and y_i may be found in polynomial time by breadth-first search. Then the word $w = u(x_1 y_1 \dots x_k y_k)^\omega$ is such that $\text{inf}(w) = C$. The length of this ultimately periodic word is at most $|Q| + |C|^2$. \square

We let $\mathbf{Witness}(C, \mathcal{M})$ denote the ultimately periodic word $u(v)^\omega$ returned by the algorithm described in the proof above for the reachable SCC C of automaton \mathcal{M} .

Proposition 5.5. *If C_1 and C_2 are SCCs of automaton \mathcal{M} and $C_1 \cap C_2 \neq \emptyset$, then $C_1 \cup C_2$ is also an SCC of \mathcal{M} .*

If \mathcal{M} is an automaton and S is any set of its states, define $\text{SCCs}(S)$ to be the set of all C such that $C \subseteq S$ and C is an SCC of \mathcal{M} . Also define $\text{maxSCCs}(S)$ to be the maximal elements of $\text{SCCs}(S)$ with respect to the subset ordering. The following is a consequence of Proposition 5.5.

Proposition 5.6. *If \mathcal{M} is an automaton and S is any set of its states, then the elements of $\text{maxSCCs}(S)$ are pairwise disjoint, and every set $C \in \text{SCCs}(S)$ is a subset of exactly one element of $\text{maxSCCs}(S)$.*

There are some differences in the terminology related to strong connectivity between graph theory and omega automata, which we resolve as follows. In graph theory, a *path* of length k from u to v in a directed graph (V, E) is a finite sequence of vertices v_0, v_1, \dots, v_k such that $u = v_0$, $v = v_k$ and for each i with $i \in [1..k]$, $(v_{i-1}, v_i) \in E$. Thus, for every vertex v , there is a path of length 0 from v to v . A set of vertices S is *strongly connected* if and only if for all $u, v \in S$, there is a path of some nonnegative length from u to v and all the vertices in the path are elements of S . Thus, for every vertex v , the singleton set $\{v\}$ is a strongly connected set of vertices. A *strongly connected component* of a directed graph is a maximal strongly connected set of vertices. There is a linear time algorithm to find the set of strong components of a directed graph [Tar72].

In this paper, we use the terminology SCC and maximal SCC to refer to the definitions from the theory of omega automata, and the terminology *graph theoretic strongly connected components* to refer to the definitions from graph theory. We use the term *trivial strong component* to refer to a graph theoretic strongly connected component that is a singleton vertex $\{v\}$ such that there is no edge (v, v) .

If \mathcal{M} is an automaton, we may define a related directed graph $G(\mathcal{M})$ whose vertices are the states of \mathcal{M} and whose edges (q_1, q_2) are the pairs of states such that $q_2 \in \delta(q_1, \sigma)$ for

some $\sigma \in \Sigma$. Then for any set S of states of \mathcal{M} , the maximal SCCs in S , $\text{maxSCCs}(S)$, are the graph theoretic strongly connected components of the subgraph of $G(\mathcal{M})$ induced by S , with any trivial strong components removed.

Proposition 5.7. *For automaton \mathcal{M} and any subset S of its states, $\text{maxSCCs}(S)$ can be computed in time linear in the size of \mathcal{M} .*

5.4. Proving efficient teachability of the informative classes — overview. We can show that a class is efficiently teachable by first showing that it is identifiable in the limit using polynomial time and data, and then giving a polynomial time teacher to construct the required characteristic samples. To show that a class is identifiable in the limit using polynomial time and data there are two parts: (i) defining a sample T_L of size polynomial in the size of the given acceptor \mathcal{A} for the language L at hand, and (ii) providing a polynomial time learning algorithm \mathbf{L} that for every given sample T returns an acceptor consistent with T , and, moreover, for any sample T consistent with L that subsumes T_L , returns an acceptor that accepts L .

The definition of an acceptor has two parts: (a) the definition of the automaton and (b) the definition of the acceptance condition. Correspondingly, we view the characteristic sample as a union of two parts: T_{Aut} (to specify the automaton) and T_{Acc} (to specify the acceptance condition). In Section 6 we discuss the construction of T_{Aut} , which is common to all the classes we consider, as they all are isomorphic to $\mathcal{M}_{\sim L}$ where $L = \llbracket \mathcal{A} \rrbracket$ for the target automaton \mathcal{A} . We also describe a polynomial time algorithm to construct $\mathcal{M}_{\sim L}$ using the sample T_{Aut} .

Because the acceptance conditions differ, T_{Acc} is different for each type of acceptor we consider. In Section 7 we describe the construction of T_{Acc} for acceptors of types IMA and IBA and learning algorithms for acceptors of these types, showing that $\mathbb{I}MA$, $\mathbb{I}BA$ and $\mathbb{I}CA$ are identifiable in the limit using polynomial time and data. In Section 8 we describe the construction of T_{Acc} for acceptors of type IPA and a learning algorithm for acceptors of this type, showing that $\mathbb{I}PA$ is identifiable in the limit using polynomial time and data. In Section 9 we describe the construction of T_{Acc} for acceptors of type IRA and a learning algorithm for acceptors of this type, showing that $\mathbb{I}RA$ and $\mathbb{I}SA$ are identifiable in the limit using polynomial time and data.

In Section 10 we show that the characteristic samples we have defined can be computed in polynomial time in the size of the acceptor. These results rely on polynomial time algorithms for the inclusion and equivalence problems for the acceptors. These are described in Sections 11, 12, 13, and 14.

This does not yet entail that the class $\mathbb{I}XA$ from $\{\mathbb{I}BA, \mathbb{I}CA, \mathbb{I}PA, \mathbb{I}MA, \mathbb{I}RA, \mathbb{I}SA\}$ is efficiently teachable. This is because $\mathbb{I}BA$ (for instance) includes also DBAs that are not IBAs but have equivalent IBAs. In Section 15 we show the right congruence automaton $\mathcal{M}_{\sim L}$ can be computed in polynomial time, given a DBA, DCA, DPA, DRA, DSA, or DMA accepting L , which yields a polynomial time algorithm to test whether a DBA is an IBA, and similarly for the other acceptor types. In Section 16, we consider the harder problem of deciding whether a DBA accepts a language in $\mathbb{I}BA$, and give polynomial time algorithms for DBAs, DCAs, DPAs, DRAs, DSAs and DMAs. With these results we can finally claim that the $\mathbb{I}XA$ classes are efficiently teachable.

6. THE SAMPLE T_{Aut} FOR THE AUTOMATON

In this section we describe the construction of the T_{Aut} part of the sample. We first show that if two states of the automaton are distinguishable, they are distinguishable by words of length polynomial in the number of states of the automaton.

6.1. Existence of short distinguishing words. Let A be an acceptor of one of the types DBA, DCA, DPA, DRA, DSA, or DMA over alphabet Σ . We say that states q_1 and q_2 of \mathcal{M} are *distinguishable* if there exists a word $w \in \Sigma^\omega$ that is accepted from one state but not the other, that is, $w \in \llbracket \mathcal{A}^{q_1} \rrbracket \setminus \llbracket \mathcal{A}^{q_2} \rrbracket$ or $w \in \llbracket \mathcal{A}^{q_2} \rrbracket \setminus \llbracket \mathcal{A}^{q_1} \rrbracket$. In this case we say that w is a *distinguishing word*.

Proposition 6.1. *If two states of a complete DBA, DCA, DPA, DRA, DSA, or DMA of n states are distinguishable, then they are distinguishable by an ultimately periodic ω -word of length bounded by $O(n^4)$.*

Proof. We prove the result for a DMA. Because any DBA, DCA, DPA, DRA, or DSA is equivalent to a DMA with the same automaton, this result holds for these types of acceptors as well. Let \mathcal{A} be a complete DMA of n states such that the states q_1 and q_2 are distinguishable. Then there exists an ω -word w that is accepted from exactly one of the two states, that is, w is accepted by exactly one of \mathcal{A}^{q_1} and \mathcal{A}^{q_2} .

Let \mathcal{M}_i denote the automaton of \mathcal{A} with its initial state replaced by q_i for $i = 1, 2$. Let \mathcal{M} denote the product automaton $\mathcal{M}_1 \times \mathcal{M}_2$. The number of states of \mathcal{M} is n^2 . By Proposition 5.3, $\text{inf}_{\mathcal{M}}(w)$ is a reachable SCC C of \mathcal{M} , and by Proposition 5.4 there exists an ultimately periodic word $u(v)^\omega$ of length bounded by $O(n^4)$ such that $\text{inf}_{\mathcal{M}}(u(v)^\omega) = C$. Then for $i = 1, 2$, $\text{inf}_{\mathcal{M}_i}(u(v)^\omega) = \pi_i(C) = \text{inf}_{\mathcal{M}_i}(w)$, so $u(v)^\omega$ is also accepted by exactly one of \mathcal{A}^{q_1} and \mathcal{A}^{q_2} , and $u(v)^\omega$ distinguishes q_1 and q_2 . \square

6.2. Defining the sample T_{Aut} for the automaton. We now define the T_{Aut} part of the characteristic sample, given an acceptor $\mathcal{A} = \langle \Sigma, Q, q_\iota, \delta, \alpha \rangle$ that is an IBA, ICA, IPA, IRA, ISA, or IMA. This construction is analogous to that of the corresponding part of a characteristic sample for a DFA, with distinguishing experiments that are ultimately periodic ω -words instead of finite strings.

Let \mathcal{M} be the automaton of \mathcal{A} and let n be the number of states of \mathcal{M} . Because \mathcal{A} is an IBA, ICA, IPA, IRA, ISA, or IMA, every state is reachable and every pair of states is distinguishable. We define a distinguished set of n *access strings* for the states of \mathcal{M} as follows. For each state q , $\text{access}(q)$ is the least string x in the shortlex ordering such that $\delta(q_\iota, x) = q$. Given \mathcal{A} , the access strings may be computed in polynomial time by breadth first search.

Because every pair of states is distinguishable, by Proposition 6.1, there exists a set E of at most n distinguishing experiments, each of length at most $n^2 + n^4$, that distinguish every pair of states. The issue of computing E is addressed in Section 10. The sample T_{Aut} consists of all the examples in $(S \cdot E) \cup (S \cdot \Sigma \cdot E)$, labeled to be consistent with \mathcal{A} . There are at most $(1 + |\Sigma|)n^2$ labeled examples in T_{Aut} , each of length bounded by a polynomial in n . A learner using T_{Aut} is described next.

6.3. Learning the automaton from T_{Aut} . We now describe a learning algorithm \mathbf{L}_{Aut} and prove the following.

Theorem 6.2. *The algorithm \mathbf{L}_{Aut} with a sample T as input runs in polynomial time and returns a deterministic complete automaton \mathcal{M} . Let \mathcal{A} be an acceptor of type IBA, ICA, IPA, IRA, ISA, or IMA. If T is consistent with \mathcal{A} and subsumes T_{Aut} then the returned automaton \mathcal{M} is isomorphic to the automaton of \mathcal{A} .*

Algorithm \mathbf{L}_{Aut} on input T constructs a set E of words that serve as experiments used to distinguish candidate states. For each $(u(v)^\omega, l)$ in T , all of the elements of $\text{suffixes}(u(v)^\omega)$ are placed in E . Two strings $x, y \in \Sigma^*$ are *consistent with respect to T* if and only if there does not exist any $u(v)^\omega \in E$ such that the examples $xu(v)^\omega$ and $yu(v)^\omega$ are oppositely labeled in T .

Starting with the empty string ε , the algorithm builds up a prefix-closed set S of finite strings as follows. Initially, $S_1 = \{\varepsilon\}$. After S_k has been constructed, the algorithm considers each $s \in S_k$ in shortlex order, and each symbol $\sigma \in \Sigma$ in the ordering defined on Σ . If there exists no $s' \in S_k$ such that $s\sigma$ is consistent with s' with respect to T , then S_{k+1} is set to $S_k \cup \{s\sigma\}$ and k is set to $k+1$. If no such pair s and σ is found, then the final set S is S_k .

In the second phase, the algorithm uses the strings in S as names for states and constructs a transition function δ using S and E . For each $s \in S$ and $\sigma \in \Sigma$, there is at least one $s' \in S$ such that $s\sigma$ and s' are consistent with respect to T . The algorithm selects any such s' and defines $\delta(s, \sigma) = s'$. Once S and δ are defined, the algorithm returns the automaton $\mathcal{M} = \langle \Sigma, S, \varepsilon, \delta \rangle$.

Proof of Theorem 6.2. E may be computed in time polynomial in the length of T , by Proposition 4.2. Because the default acceptor for T has a polynomial number of states and is consistent with T , the number of distinguishable states, and the number of strings added to S , is bounded by a polynomial in the length of T . The returned automaton \mathcal{M} is deterministic and complete by construction.

Assume the sample T is consistent with \mathcal{A} and subsumes T_{Aut} . For any pair of states of \mathcal{A} , the set E includes an experiment to distinguish them. Also, if x and y reach the same state of \mathcal{A} , there is no experiment in E that distinguishes them. Then the set S is precisely the access strings of \mathcal{A} . The choice of s' for $\delta(s, \sigma)$ is unique in each case, and the returned automaton \mathcal{M} is isomorphic to the automaton of \mathcal{A} . \square

Although the processes of constructing T_{Aut} and learning an automaton from it are the same for acceptors of types IBA, ICA, IPA, IRA, ISA, or IMA, different types of acceptance condition require different kinds of characteristic samples and learning algorithms.

In the following sections we describe for each type of acceptor the corresponding sample T_{Acc} and learning algorithm. Each learning algorithm takes as input an automaton \mathcal{M} and a sample T and returns in polynomial time an acceptor of the appropriate type consistent with T . We show that for each type of acceptor \mathcal{A} , if the input automaton \mathcal{M} is isomorphic to the automaton of \mathcal{A} and the sample T is consistent with \mathcal{A} and subsumes the T_{Acc} for \mathcal{A} , then the learning algorithm returns an acceptor that is equivalent to \mathcal{A} . This learning algorithm is then combined with \mathbf{L}_{Aut} to prove the relevant class of languages are identifiable in the limit using polynomial time and data.

7. THE SAMPLES T_{Acc} AND LEARNING ALGORITHMS FOR IMA AND IBA

The straightforward cases of Muller, Büchi and coBüchi acceptance conditions are covered in this section. Subsequent sections cover the cases of parity, Rabin and Street acceptance conditions, which are somewhat more involved.

7.1. Muller acceptors. Let \mathcal{A} be an IMA with acceptance condition $\alpha = \{F_1, \dots, F_k\}$. By Proposition 5.3, we may assume that each F_i is a reachable SCC of \mathcal{A} . The sample T_{Acc}^{IMA} consists of k positive examples, one for each set F_i . The example for F_i is $(u(v)^\omega, 1)$ where $\text{inf}(u(v)^\omega) = F_i$. These examples may be found in polynomial time in the size of \mathcal{A} by Proposition 5.4.

The learning algorithm \mathbf{L}_{Acc}^{IMA} takes as input a deterministic complete automaton \mathcal{M} and a sample T . It constructs an acceptance condition α' as follows. For each positive labeled example $(u(v)^\omega, 1) \in T$, it computes the set $C = \text{inf}_{\mathcal{M}}(u(v)^\omega)$ and makes C a member of α' . Once the set α' is complete, the algorithm checks whether the DMA (\mathcal{M}, α') is consistent with T . If so, it returns (\mathcal{M}, α') ; if not, it returns the default acceptor of type DMA for T .

Theorem 7.1. *Algorithm \mathbf{L}_{Acc}^{IMA} runs in time polynomial in the sizes of the inputs \mathcal{M} and T . Let \mathcal{A} be an IMA. If the input automaton \mathcal{M} is isomorphic to the automaton of \mathcal{A} , and the sample T is consistent with \mathcal{A} and subsumes T_{Acc}^{IMA} , then algorithm \mathbf{L}_{Acc}^{IMA} returns an IMA (\mathcal{M}, α') equivalent to \mathcal{A} .*

Proof. The construction of α' can be done in time polynomial in the sizes of \mathcal{M} and T by Proposition 5.3. The returned acceptor is consistent with T by construction.

Assume \mathcal{M} is isomorphic to the automaton of \mathcal{A} and that T is consistent with \mathcal{A} . For ease of notation, assume the isomorphism is the identity. Then for each positive example $(u(v)^\omega, 1)$ in T , the set $F = \text{inf}(u(v)^\omega)$ must be in α , so α' is a subset of α .

If T subsumes T_{Acc}^{IMA} , then for every set $F \in \alpha$ there is a positive example $(u(v)^\omega, 1)$ in T with $F = \text{inf}(u(v)^\omega)$. Thus the set F is added to α' , and α is a subset of α' . Thus, (\mathcal{M}, α') is equivalent to \mathcal{A} , and because T is consistent with \mathcal{A} , the IMA (\mathcal{M}, α') is returned by \mathbf{L}_{Acc}^{IMA} . \square

Theorem 7.2. *The class IMA is identifiable in the limit using polynomial time and data.*

Proof. Let \mathcal{A} be an IMA accepting a language L . The characteristic sample $T_L = T_{Aut} \cup T_{Acc}^{IMA}$ is of size polynomial in the size of \mathcal{A} .

The combined learner \mathbf{L}^{IMA} takes a sample T as input and runs \mathbf{L}_{Aut} on T to produce an automaton \mathcal{M} and then runs \mathbf{L}_{Acc}^{IMA} on \mathcal{M} and T and returns the resulting acceptor. It runs in polynomial time in the size of T because it is the composition of two polynomial time algorithms, and the acceptor it returns is guaranteed to be consistent with T .

If the sample T is consistent with \mathcal{A} and subsumes T_L , then by Theorem 6.2 the automaton \mathcal{M} returned by \mathbf{L}_{Aut} is isomorphic to the automaton of \mathcal{A} . Then by Theorem 7.1 the acceptor returned by \mathbf{L}_{Acc}^{IMA} with inputs \mathcal{M} and T is an IMA equivalent to \mathcal{A} . \square

7.2. Büchi acceptors. The case of Büchi acceptors is nearly as straightforward as that of Muller acceptors. Let \mathcal{A} be an IBA with n states and acceptance condition F . For every state q of \mathcal{A} , if there is an ω -word w such that \mathcal{A} rejects w and $q \in \text{inf}(w)$, then there is an example $u(v)^\omega$ of length $O(n^2)$ such that \mathcal{A} rejects $u(v)^\omega$ and $q \in \text{inf}(u(v)^\omega)$, by Proposition 5.4. The negative labeled example $(u(v)^\omega, 0)$ is included in T_{Acc}^{IBA} .

The learning algorithm \mathbf{L}_{Acc}^{IBA} takes as input a deterministic complete automaton \mathcal{M} and a sample T . The acceptance condition F' consists of all the states q of \mathcal{M} such that for no negative example $(u(v)^\omega, 0)$ in T do we have $q \in \text{inf}_{\mathcal{M}}(u(v)^\omega)$. Once F' has been computed, the algorithm checks whether the DBA (\mathcal{M}, F') is consistent with the sample T . If so, it returns (\mathcal{M}, F') ; if not, it returns the default acceptor of type DBA for T .

Theorem 7.3. *Algorithm \mathbf{L}_{Acc}^{IBA} runs in time polynomial in the sizes of the inputs \mathcal{M} and T . Let \mathcal{A} be an IBA. If the input automaton \mathcal{M} is isomorphic to the automaton of \mathcal{A} , and the sample T is consistent with \mathcal{A} and subsumes T_{Acc}^{IBA} , then algorithm \mathbf{L}_{Acc}^{IBA} returns an IBA (\mathcal{M}, F') equivalent to \mathcal{A} .*

Proof. The construction of F' can be done in time polynomial in the sizes of \mathcal{M} and T by Proposition 5.3. The returned acceptor is consistent with T by construction.

Assume the input \mathcal{M} is isomorphic to the automaton of \mathcal{A} , and that T is consistent with \mathcal{A} and subsumes T_{Acc}^{IBA} . For ease of notation, assume the isomorphism is the identity. We show that the DBA (\mathcal{M}, F') is equivalent to \mathcal{A} .

If \mathcal{A} rejects the word $u(v)^\omega$ then let $C = \text{inf}_{\mathcal{M}}(u(v)^\omega)$. Because T subsumes T_{Acc}^{IBA} , for each $q \in C$, there is a negative example $(u'(v')^\omega, 0)$ in T such that $q \in \text{inf}_{\mathcal{M}}(u'(v')^\omega)$. Thus no $q \in C$ is in F' and (\mathcal{M}, F') also rejects $u(v)^\omega$.

Conversely, if \mathcal{A} accepts the word $u(v)^\omega$, then there is at least one state $q \in F$ such that $q \in \text{inf}_{\mathcal{M}}(u(v)^\omega)$. Because T is consistent with \mathcal{A} , there is no negative example $(u(v)^\omega, 0)$ in T such that $q \in \text{inf}_{\mathcal{M}}(u(v)^\omega)$, so $q \in F'$ and (\mathcal{M}, F') also accepts $u(v)^\omega$. Thus (\mathcal{M}, F') is equivalent to \mathcal{A} . Because T is consistent with \mathcal{A} , the IBA (\mathcal{M}, F') is returned by \mathbf{L}_{Acc}^{IBA} . \square

Theorem 7.4. *The classes \mathbb{IBA} and \mathbb{ICA} are identifiable in the limit using polynomial time and data.*

Proof. The result for \mathbb{ICA} follows from that for \mathbb{IBA} by Proposition 5.1. Let \mathcal{A} be an IBA accepting language L . The characteristic sample $T_L = T_{Aut} \cup T_{Acc}^{IBA}$ is of size polynomial in $\text{size}(\mathcal{A})$.

The combined learning algorithm \mathbf{L}^{IBA} takes a sample T as input and runs \mathbf{L}_{Aut} to get a deterministic complete automaton \mathcal{M} . It then runs \mathbf{L}_{Acc}^{IBA} on inputs \mathcal{M} and T , and returns the resulting acceptor. \mathbf{L}^{IBA} runs in polynomial time in the length of T and returns a DBA consistent with T .

If the sample T is consistent with \mathcal{A} and subsumes T_L then \mathbf{L}_{Aut} returns an automaton \mathcal{M} isomorphic to the automaton of \mathcal{A} by Theorem 6.2. Then the acceptor returned by \mathbf{L}_{Acc}^{IBA} on inputs \mathcal{M} and T is an IBA equivalent to \mathcal{A} by Theorem 7.3. \square

8. THE SAMPLE T_{Acc} AND THE LEARNING ALGORITHM FOR IPA

The construction of T_{Acc}^{IPA} for an IPA \mathcal{P} builds on the construction of the canonical forest of SCCs for \mathcal{P} , whose construction and properties are described next. Roughly speaking, the purpose of the canonical forest for a given parity automaton \mathcal{P} is to expose a set of words that if placed in the sample will lead a smart learner to correctly determine a coloring function for the constructed automaton. It is thus not surprising, that while developed for a different motivation, it has similarities with Carton and Maceiras's algorithm to compute the minimal number of colors for a given parity automaton [CM99]. We begin with the definition and properties of a decreasing forest of SCCs of an ω -automaton.

8.1. A decreasing forest of SCCs of an automaton. Let \mathcal{M} be a deterministic automaton, and let S be a subset of its states. A *decreasing forest of SCCs of \mathcal{M} rooted in S* is a finite rooted forest \mathcal{F} in which every node C is an SCC of \mathcal{M} that is contained in S , and the following properties are satisfied.

- (1) The roots of \mathcal{F} are the elements of $\text{maxSCCs}(S)$.
- (2) Whenever D_1, \dots, D_k are the children of node C , we have $D_1 \cup \dots \cup D_k \subsetneq C$. Also, letting $\Delta(C) = C \setminus (D_1 \cup \dots \cup D_k)$, the children D_1, \dots, D_k are exactly the elements of $\text{maxSCCs}(C \setminus \Delta(C))$.

Proposition 8.1. *Let \mathcal{M} be a deterministic automaton, S a subset of its states, and \mathcal{F} a decreasing forest of SCCs of \mathcal{M} rooted in S . Then the following are true.*

- (1) *The roots of \mathcal{F} are pairwise disjoint.*
- (2) *The children of any node are pairwise disjoint.*
- (3) *\mathcal{F} has at most $|S|$ nodes.*
- (4) *For any $D \subseteq S$ that is an SCC of \mathcal{M} , there is a unique node C in \mathcal{F} such that $D \subseteq C$ and D is not a subset of any of the children of C .*

Proof. The roots of \mathcal{F} are the elements of $\text{maxSCCs}(S)$, which are pairwise disjoint. The children of a node C are the elements of $\text{maxSCCs}(C \setminus \Delta(C))$, which are pairwise disjoint. The sets $\Delta(C)$ for nodes C in \mathcal{F} are contained in S , nonempty, and pairwise disjoint, so the number of nodes is at most $|S|$. If $D \subseteq S$ is an SCC of \mathcal{M} , then D is a subset of exactly one of the roots of \mathcal{F} , say C_1 . If $D \cap \Delta(C_1) \neq \emptyset$, then D is not a subset of any of the children of C_1 . Otherwise, D must be a subset of exactly one of the children of C_1 , say C_2 . If $D \cap \Delta(C_2) \neq \emptyset$, then D is not a subset of any of the children of C_2 . Continuing in this way, we eventually arrive at the required node C . \square

Given a decreasing forest \mathcal{F} of SCCs of automaton \mathcal{M} rooted in S , and an SCC $D \subseteq S$, we denote by $\mathbf{Node}(D, \mathcal{F})$ the unique node C of \mathcal{F} such that $D \subseteq C$ and D is not a subset of any of the children of C . We note that if $C = \mathbf{Node}(D, \mathcal{F})$ then $D \cap \Delta(C) \neq \emptyset$. If D is a child of some C in \mathcal{F} , we define *merging D into C* as the operation of removing D from \mathcal{F} and making the children of D (if any) direct children of C .

Proposition 8.2. *Let \mathcal{M} be a deterministic automaton and S a subset of its states. Let \mathcal{F} be a decreasing forest of SCCs of \mathcal{M} rooted in S . Let D be a child of C in \mathcal{F} and let \mathcal{F}' be obtained from \mathcal{F} by merging D into C . Then \mathcal{F}' is also a decreasing forest of SCCs of \mathcal{M} rooted in S .*

Proof. After the merge, the roots of \mathcal{F} remain the elements of $\text{maxSCCs}(S)$. Let D_1, \dots, D_k be the children of C in \mathcal{F} , where $D = D_k$, and let E_1, \dots, E_ℓ be the children of D in \mathcal{F} . Because the union of E_1, \dots, E_ℓ is a proper subset of $D = D_k$ and the union of D_1, \dots, D_k is a proper subset of C , the union of D_1, \dots, D_{k-1} with the union of E_1, \dots, E_ℓ is a proper subset of C , therefore the union of the children of C in \mathcal{F}' is a proper subset of C . Also, the children of C in \mathcal{F} are the maximum SCCs of $C \setminus \Delta_{\mathcal{F}}(C)$, and no other nodes are affected, so \mathcal{F}' is a decreasing forest of SCCs of \mathcal{M} rooted in S . \square

8.2. Constructing the Canonical Forest and Coloring of a DPA. Let $\mathcal{P} = \langle \Sigma, Q, q_i, \delta, \kappa \rangle$ be a complete DPA. We extend the coloring function κ to nonempty sets of states by $\kappa(S) = \min\{\kappa(q) \mid q \in S\}$, the minimum color of any state in S . We define the κ -parity of S to be 1 if $\kappa(S)$ is odd, and 0 if $\kappa(S)$ is even. A word $w \in \Sigma^\omega$ is accepted by \mathcal{P} iff the

κ -parity of $\text{inf}(w)$ is 1. Note that the union of two sets of κ -parity b is also of κ -parity b . For any nonempty $S \subseteq Q$, we define $\text{minStates}(S) = \{q \in S \mid \kappa(q) = \kappa(S)\}$, the states of S that are assigned the minimum color among all states of S .

8.2.1. *The minStates-Forest.* We describe an algorithm to construct the *minStates-forest* of \mathcal{P} . The roots of the *minStates-forest* are the elements of $\text{maxSCCs}(Q)$, each marked as unprocessed. The following procedure is repeated until all elements are marked as processed. If C is unprocessed, then $\mathcal{D} = \text{maxSCCs}(C \setminus \text{minStates}(C))$ is computed. If \mathcal{D} is empty, C becomes a leaf in the forest, and is marked as processed. Otherwise, C is marked as processed and the elements of \mathcal{D} are made the children of C and are marked as unprocessed.

Proposition 8.3. *Let $\mathcal{P} = \langle \Sigma, Q, q_\iota, \delta, \kappa \rangle$ be a complete DPA with automaton \mathcal{M} . Let \mathcal{F} be the *minStates-forest* of \mathcal{P} . Then \mathcal{F} is a decreasing forest of SCCs of \mathcal{M} rooted in Q , and can be computed in polynomial time. For any SCC D , $\kappa(D) = \kappa(\mathbf{Node}(D, \mathcal{F}))$.*

Proof. Referring to the construction of the *minStates-forest* \mathcal{F} , its roots are the elements of $\text{maxSCCs}(Q)$. When a node C is processed, the nonempty set $\text{minStates}(C)$ is removed and the maximum SCCs (if any) of the result become the children of C , so the union of the children of C is a proper subset of C , and the children are the maximum SCCs of $C \setminus \Delta(C)$. Let D be an SCC. Then $D \subseteq Q$ and for the node $C = \mathbf{Node}(D, \mathcal{F})$, we have that $D \subseteq C$ and D is not a subset of any child of C . Thus $D \cap \text{minStates}(C) \neq \emptyset$, because otherwise D would be a subset of some child of C . This implies that $\kappa(D) = \kappa(C)$. The *minStates-forest* of \mathcal{P} can be computed in polynomial time because it has at most $|Q|$ nodes, and each set $\text{maxSCCs}(S)$ can be computed in polynomial time by Proposition 5.7. \square

8.2.2. *The Canonical Forest and Coloring.* The canonical forest of \mathcal{P} is constructed as follows, starting with the *minStates-forest* of \mathcal{P} . While there exist in the forest a node D and its parent C of the same κ -parity, one such pair D and C is selected, and the child node D is merged into the parent node C . When no such pair remains, the result is the *canonical forest* of \mathcal{P} , denoted $\mathcal{F}^*(\mathcal{P})$. The canonical forest $\mathcal{F}^*(\mathcal{P})$ can be computed from \mathcal{P} in polynomial time.

From the canonical forest $\mathcal{F}^*(\mathcal{P})$, we define the *canonical coloring* κ^* . The states in $(Q \setminus \bigcup \text{maxSCCs}(Q))$ are not contained in any SCC of \mathcal{P} and do not affect the acceptance or rejection of any ω -word. For definiteness, we assign them $\kappa^*(q) = 0$. For a root node C of κ -parity b , we define $\kappa^*(q) = b$ for all $q \in \Delta(C)$. Let C be an arbitrary node of $\mathcal{F}^*(\mathcal{P})$. If the states of $\Delta(C)$ have been assigned color k by κ^* and D is a child of C , then the states of $\Delta(D)$ are assigned color $k + 1$ by κ^* . Clearly κ^* can be computed from \mathcal{P} in polynomial time.

Example 8.4. Figure 3 (a) shows the graph $G(\mathcal{P})$ of a DPA \mathcal{P} with states a through m , labeled by the colors assigned by κ . Figure 3 (b) shows the *minStates-forest* of \mathcal{P} , with the nodes labeled by their κ -parities. Figure 3 (c) shows the canonical forest $\mathcal{F}^*(\mathcal{P})$ of \mathcal{P} , with the nodes labeled by their κ -parities. Figure 3 (d) shows the graph $G(\mathcal{P})$ re-colored using the canonical coloring κ^* .

Theorem 8.5. *Let $\mathcal{P} = \langle \Sigma, Q, q_\iota, \delta, \kappa \rangle$ be a complete DPA with automaton \mathcal{M} . The canonical forest $\mathcal{F}^*(\mathcal{P})$ is a decreasing forest of SCCs of \mathcal{M} rooted in Q and has the following properties.*

(1) *For any SCC D , both D and $\mathbf{Node}(D, \mathcal{F}^*(\mathcal{P}))$ have the same κ -parity.*

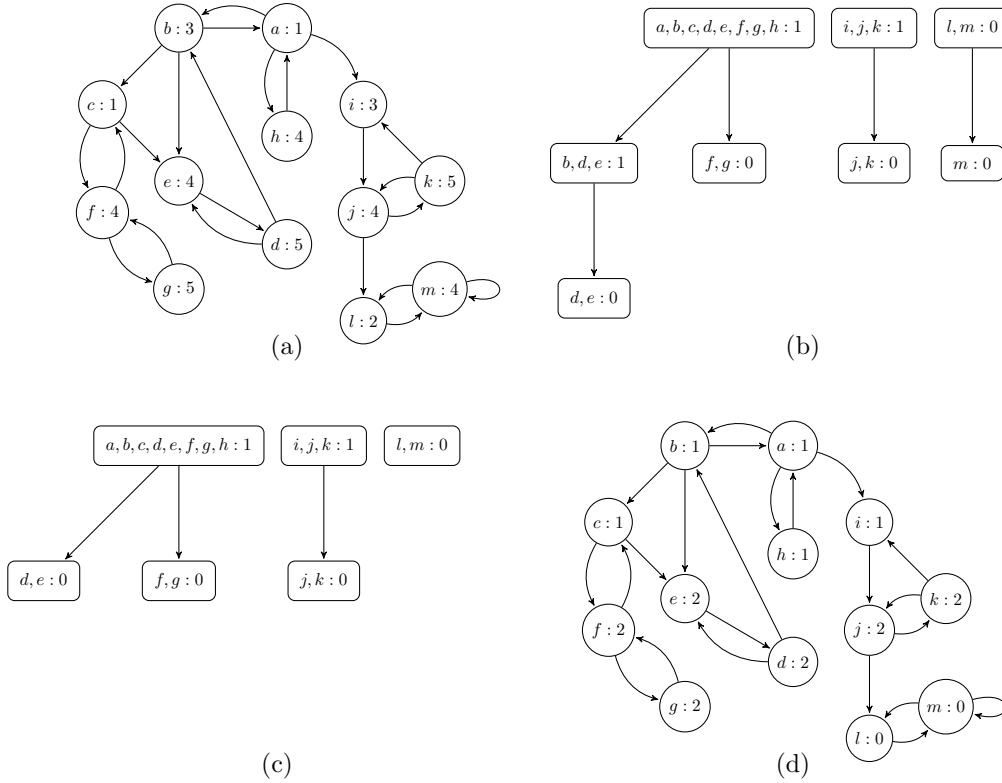


Figure 3. (a) Graph $G(\mathcal{P})$ with states colored by κ . (b) The minStates -forest of \mathcal{P} , with κ -parities of nodes. (c) Canonical forest $\mathcal{F}^*(\mathcal{P})$, with κ -parities of nodes. (d) Graph $G(\mathcal{P})$ with the canonical coloring κ^* .

- (2) For every node C of $\mathcal{F}^*(\mathcal{P})$, the κ -parity of C is the same as the κ^* -parity of C .
 (3) The children in $\mathcal{F}^*(\mathcal{P})$ of a node C of κ -parity b are the maximal SCCs $D \subseteq C$ of κ -parity $1 - b$.

Proof. Because $\mathcal{F}^*(\mathcal{P})$ is obtained from the minStates -forest of \mathcal{P} by a sequence of merges, $\mathcal{F}^*(\mathcal{P})$ is a decreasing forest of SCCs of \mathcal{M} rooted in Q by Proposition 8.2. Let \mathcal{F}_0 denote the minStates -forest of \mathcal{P} , and let \mathcal{F}_i denote the forest after i merges have been performed in the computation to produce the canonical acceptor $\mathcal{F}^*(\mathcal{P})$.

By Proposition 8.3, for any SCC D , $\kappa(D) = \kappa(\mathbf{Node}(D, \mathcal{F}_0))$, so property (1) holds for \mathcal{F}_0 . We show by induction that it holds for each \mathcal{F}_i and therefore for $\mathcal{F}^*(\mathcal{P})$. Assume that property (1) holds of \mathcal{F}_i , and \mathcal{F}_{i+1} is obtained from \mathcal{F}_i by merging child node D into parent node C . Let D' be any SCC. If $\mathbf{Node}(D', \mathcal{F}_i) = \mathbf{Node}(D', \mathcal{F}_{i+1})$, then because property (1) holds in \mathcal{F}_i , we have that the κ -parity of D' is the same as the κ -parity of $\mathbf{Node}(D', \mathcal{F}_{i+1})$. Otherwise, it must be that $D = \mathbf{Node}(D', \mathcal{F}_i)$ and $C = \mathbf{Node}(D', \mathcal{F}_{i+1})$. Because D is only merged to C if they are of the same κ -parity, this implies that the κ -parity of D' is the same as the κ -parity of C . Thus, property (1) holds also in \mathcal{F}_{i+1} .

For property (2), we note that the κ -parity and the κ^* -parity of each root node of $\mathcal{F}^*(\mathcal{P})$ is the same. Suppose C is a node of $\mathcal{F}^*(\mathcal{P})$ whose κ -parity and κ^* -parity are equal to b , and D is a child of C . Then by the construction of $\mathcal{F}^*(\mathcal{P})$, the κ -parity of D is $1 - b$. And by

the definition of κ^* , the κ^* parity of the elements of $\Delta(D)$ is the opposite of the κ^* -parity of the elements of $\Delta(C)$. Because the κ^* -parity of C is b , the κ^* -parity of D is also $1 - b$.

For property (3), let D be any maximal SCC of \mathcal{P} that is contained in C and has κ -parity $1 - b$. Then $\mathbf{Node}(D, \mathcal{F}_0)$ is contained in C and has the same κ -parity as D . Since D is maximal, we must have $D = \mathbf{Node}(D, \mathcal{F}_0)$, and any nodes on the path between C and D must have κ -parity b and must be merged into C to form $\mathcal{F}^*(\mathcal{P})$. Thus, D is a child of C in $\mathcal{F}^*(\mathcal{P})$.

Conversely, if D is a child of C in $\mathcal{F}^*(\mathcal{P})$ then $D \subseteq C$ and the κ -parity of D is $1 - b$. Assume D' is an SCC such that $D' \subseteq C$, the κ -parity of D' is $1 - b$ and $D \subsetneq D'$. Then $D'' = \mathbf{Node}(D', \mathcal{F}_0)$ is a descendant of C in \mathcal{F}_0 that has κ -parity $1 - b$, and $D \subsetneq D''$, so because there is another node of parity $1 - b$ on the path between D and C in \mathcal{F}_0 , D cannot be a child of C in $\mathcal{F}^*(\mathcal{P})$, a contradiction. \square

Replacing the coloring function of \mathcal{P} by the canonical coloring does not change the ω -language accepted.

Theorem 8.6. *Let $\mathcal{P} = \langle \Sigma, Q, q_i, \delta, \kappa \rangle$ be a complete DPA, and \mathcal{P}^* be \mathcal{P} with the canonical coloring κ^* for \mathcal{P} in place of κ . Then \mathcal{P} and \mathcal{P}^* recognize the same ω -language.*

Proof. Let w be an ω -word and let $D = \mathit{inf}(w)$. This is an SCC of the (common) automaton of \mathcal{P} and \mathcal{P}^* . Let $C = \mathbf{Node}(D, \mathcal{F}^*(\mathcal{P}))$. Then $D \cap \Delta(C) \neq \emptyset$, and $\kappa^*(D) = \kappa^*(C)$, by the definition of κ^* . The κ^* parity of C is the same as the κ -parity of C , by property (2) of Theorem 8.5. The κ -parity of C is the same as the κ -parity of D , by property (1) of Theorem 8.5. Thus, the κ^* -parity of D is the same as the κ -parity of D , and $w \in \llbracket \mathcal{P} \rrbracket$ iff $w \in \llbracket \mathcal{P}^* \rrbracket$. \square

8.3. Constructing T_{Acc}^{IPA} . We now describe the construction of T_{Acc}^{IPA} , the second part of the characteristic sample for an IPA \mathcal{P} with the automaton \mathcal{M} of n states. The sample T_{Acc}^{IPA} consists of one example $u(v)^\omega = \mathbf{Witness}(C, \mathcal{M})$ of length $O(n^2)$ for each reachable SCC C in the canonical forest $\mathcal{F}^*(\mathcal{P})$. The example $u(v)^\omega$ is labeled 1 if it is accepted by \mathcal{P} and 0 otherwise. Thus T_{Acc}^{IPA} contains at most n labeled examples, each of length $O(n^2)$.

8.4. The learning algorithm \mathbf{L}_{Acc}^{IPA} . Given a complete deterministic automaton $\mathcal{M} = \langle \Sigma, Q, q_i, \delta \rangle$ and a sample T as input, the learning algorithm \mathbf{L}_{Acc}^{IPA} attempts to construct a coloring of the states of \mathcal{M} consistent with T .

The algorithm first constructs the set Z of all $C \subseteq Q$ such that for some labeled example $(u(v)^\omega, l)$ in T we have $C = \mathit{inf}_{\mathcal{M}}(u(v)^\omega)$. If two examples with different labels are found to yield the same set C , this is evidence that the automaton \mathcal{M} is not correct, and the learning algorithm returns the default acceptor of type DPA for T .

Otherwise, each set C in Z is associated with the label of the one or more examples that yield C . The set Z is partially ordered by the subset relation. The learning algorithm then attempts to construct a rooted forest \mathcal{F}' with nodes that are elements of Z , corresponding to the canonical forest of the target acceptor. Initially, \mathcal{F}' contains as roots all the maximal elements of Z . If these are not pairwise disjoint, it returns the default acceptor of type DPA for T . Otherwise, the root nodes are all marked as unprocessed.

For each unprocessed node C in \mathcal{F}' , it computes the set of all $D \in Z$ such that $D \subseteq C$, D has the opposite label to C , and D is maximal with these properties, and makes D a child

of C and marks D as unprocessed. When all the children of a node C have been determined, the algorithm checks two conditions: (1) that the children of C are pairwise disjoint, and (2) there is at least one $q \in C$ that is not in any child of C . If either of these conditions fail, then it returns the default acceptor of type DPA for T . If both conditions are satisfied, then the node C is marked as processed. When there are no more unprocessed nodes, the construction of \mathcal{F}' is complete. Note that \mathcal{F}' has at most $|Q|$ nodes.

When the construction of \mathcal{F}' is complete, for each node C in \mathcal{F}' let $\Delta(C)$ denote the elements of C that do not appear in any of its children. Then the learning algorithm assigns colors to the elements of Q starting from the roots of \mathcal{F}' , as follows. If C is a root with label l , then $\kappa'(q) = l$ for all $q \in \Delta(C)$. If the elements of $\Delta(C)$ have been assigned color k and D is a child of C , then $\kappa'(q) = k + 1$ for all $s \in \Delta(D)$. When this process is complete, any uncolored states q are assigned $\kappa'(q) = 0$.

If the resulting DPA (\mathcal{M}, κ') is consistent with the sample T , the algorithm \mathbf{L}_{Acc}^{IPA} returns (\mathcal{M}, κ') . If not, it returns the default acceptor of type DPA for T .

Theorem 8.7. *Algorithm \mathbf{L}_{Acc}^{IPA} runs in time polynomial in the sizes of the inputs \mathcal{M} and T . Let \mathcal{P} be an IPA. If the input automaton \mathcal{M} is isomorphic to the automaton of \mathcal{P} , and the sample T is consistent with \mathcal{P} and subsumes T_{Acc}^{IPA} , then algorithm \mathbf{L}_{Acc}^{IPA} returns an IPA (\mathcal{M}, κ') equivalent to \mathcal{P} .*

Proof. The construction of κ' can be done in time polynomial in the sizes of \mathcal{M} and T . The returned acceptor is consistent with T by construction.

Assume the input $\mathcal{M} = \langle \Sigma, Q, q_\iota, \delta \rangle$ is isomorphic to the automaton of \mathcal{P} , and that T is consistent with \mathcal{P} and subsumes T_{Acc}^{IPA} . For ease of notation, assume the isomorphism is the identity. We show that the forest \mathcal{F}' constructed by the learning algorithm is equal to the canonical forest of $\mathcal{F}^*(\mathcal{P})$, the coloring κ' is equal to the canonical coloring κ^* , and therefore the acceptor (\mathcal{M}, κ') is equivalent to \mathcal{P} .

The roots of $\mathcal{F}^*(\mathcal{P})$ are the maximal SCCs contained in Q , and for each such root C , T_{Acc}^{IPA} contains an example $(u(v)^\omega, l)$ such that $C = \text{inf}(u(v)^\omega)$. Thus, the set of maximal elements of Z is equal to the set of roots of $\mathcal{F}^*(\mathcal{P})$.

Let C be any node of $\mathcal{F}^*(\mathcal{P})$, and let D be a child of C in $\mathcal{F}^*(\mathcal{P})$. Then D is an SCC, $D \subseteq C$, the parity of D is opposite to the parity of C , and D is maximal in the subset ordering with these properties, by property (3) of Theorem 8.5. In the sample T_{Acc}^{IPA} there is an example $(u(v)^\omega, l)$ with $D = \text{inf}(u(v)^\omega)$, so D is an element of Z , and will be made a child of C in \mathcal{F}' because $D \subseteq C$, the label l is the opposite of the label of C , and D is maximal in Z with these properties. Conversely, if D is made a child of C in \mathcal{F}' , then $D \subseteq C$, the label of D is opposite to the label of C (that is, they are of opposite κ -parity), and D is maximal in Z with these properties. This implies D is a child of C in $\mathcal{F}^*(\mathcal{P})$, by property (3) of Theorem 8.5.

By induction, \mathcal{F}' is equal to $\mathcal{F}^*(\mathcal{P})$, and therefore κ' is equal to the canonical coloring κ^* . Then the IPA (\mathcal{M}, κ') is equivalent to \mathcal{P} , by Theorem 8.6. Because T is consistent with \mathcal{P} , the IPA (\mathcal{M}, κ') is returned by \mathbf{L}_{Acc}^{IPA} . \square

Theorem 8.8. *The class IPA is identifiable in the limit using polynomial time and data.*

Proof. Let \mathcal{P} be an IPA accepting the language L . The characteristic sample $T_L = T_{Aut} \cup T_{Acc}^{IPA}$ for \mathcal{P} is of size polynomial in the size of \mathcal{P} .

The combined learning algorithm \mathbf{L}^{IPA} with a sample T as input first runs \mathbf{L}_{Aut} on T to get a complete deterministic automaton \mathcal{M} and then runs \mathbf{L}_{Acc}^{IPA} on inputs \mathcal{M} and T and

returns the resulting acceptor. The running time of \mathbf{L}^{IPA} is polynomial in the length of T and the returned acceptor is consistent with T .

Assume the sample T is consistent with \mathcal{P} and subsumes T_L . By Theorem 6.2, the automaton \mathcal{M} is isomorphic to the automaton of \mathcal{P} . By Theorem 8.7, the acceptor (\mathcal{M}, κ') is equivalent to \mathcal{P} . Because it is consistent with T , the IPA (\mathcal{M}, κ') is the acceptor returned by \mathbf{L}^{IPA} . \square

9. THE SAMPLE T_{Acc} AND LEARNING ALGORITHM FOR IRA

In this section we introduce some terminology, establish a normal form for Rabin acceptors, define an ordering on sets of states of an automaton, and then describe the learning algorithm $\mathbf{L}_{\text{Acc}}^{\text{IRA}}$ and sample $T_{\text{Acc}}^{\text{IRA}}$ for IRAs. We then prove that the classes $\mathbb{I}\text{RA}$ and $\mathbb{I}\text{SA}$ are identifiable in the limit using polynomial time and data.

Let $\mathcal{R} = \langle \Sigma, Q, q_i, \delta, \alpha \rangle$ be a Rabin acceptor, where the acceptance condition $\alpha = \{(G_1, B_1), \dots, (G_k, B_k)\}$ is a set of ordered pairs of states. We say that an ω -word w *satisfies a pair of state sets* (G, B) iff $\text{inf}(w) \cap G \neq \emptyset$ and $\text{inf}(w) \cap B = \emptyset$. Also, w *satisfies the acceptance condition* α iff there exists $i \in [1..k]$ such that w satisfies (G_i, B_i) . Then an ω -word w is accepted by \mathcal{R} iff w satisfies the acceptance condition α of \mathcal{R} .

9.1. Singleton normal form for a Rabin acceptor. We say that a Rabin acceptor \mathcal{R} is in *singleton normal form* iff for every pair (G_i, B_i) in its acceptance condition we have $|G_i| = 1$, that is, every G_i is a singleton set. To avoid extra braces, we abbreviate the pair $(\{q\}, B)$ by (q, B) . Every Rabin acceptor may be put into singleton normal form by a polynomial time algorithm.

Proposition 9.1. *Let $\mathcal{R} = \langle \Sigma, Q, q_i, \delta, \alpha \rangle$ be a Rabin acceptor where $\alpha = \{(G_1, B_1), \dots, (G_k, B_k)\}$. Define the acceptance condition α' to contain (q, B_i) for every $i \in [1..k]$ and $q \in G_i$, and let \mathcal{R}' be \mathcal{R} with α replaced by α' . Then \mathcal{R}' is in singleton normal form and accepts the same language as \mathcal{R} . Also, \mathcal{R}' is of size at most $|Q|$ times the size of \mathcal{R} .*

Proof. If an ω -word w satisfies a pair (q, B) of α' , then there exists a pair (G_i, B_i) of α with $q \in G_i$ and $B_i = B$, so w also satisfies the pair (G_i, B_i) in α . Conversely, if w satisfies a pair (G_i, B_i) in α , then there exists $q \in G_i$ such that $q \in \text{inf}(w)$, so w also satisfies (q, B_i) in α' . Each pair (G_i, B_i) in α is replaced by at most $|Q|$ pairs in α' . \square

9.2. An ordering on sets of states. Given an automaton, we define an ordering \preceq on sets of its states that is used to coordinate between the characteristic sample and the learning algorithm for an IRA.

Let $\mathcal{M} = \langle \Sigma, Q, q_i, \delta \rangle$ be a deterministic complete automaton in which every state is reachable from the initial state q_i . Recall from Section 6.2 that $\text{access}(q)$ is the shortlex least string $s \in \Sigma^*$ such that $\delta(q_i, s) = q$.

For a set of states S , we define $\text{access}(S)$ to be the sequence of values $\text{access}(q)$ for $q \in S$, sorted into increasing shortlex order. We define the total ordering \preceq on sets of states of \mathcal{M} as follows. If $S_1, S_2 \subseteq Q$, then $S_1 \preceq S_2$ iff either $|S_1| < |S_2|$ or $|S_1| = |S_2|$ and the sequence $\text{access}(S_1)$ is less than or equal to the sequence $\text{access}(S_2)$ in the lexicographic ordering, using the shortlex ordering on Σ^* to compare the component entries. For example, if

$access(S_1) = \langle \epsilon, a, baa \rangle$, $access(S_2) = \langle a, ba \rangle$, and $access(S_3) = \langle \epsilon, ab, ba \rangle$ then $S_2 \preceq S_1 \preceq S_3$.

9.3. The learning algorithm \mathbf{L}_{Acc}^{IRA} . We begin with the description of the learning algorithm \mathbf{L}_{Acc}^{IRA} , which is used in the definition of the sample T_{Acc}^{IRA} in the next section. The inputs to \mathbf{L}_{Acc}^{IRA} are a deterministic complete automaton \mathcal{M} and a sample T . The algorithm attempts to construct a singleton normal form Rabin acceptance condition β to produce an acceptor (\mathcal{M}, β) consistent with T .

The processing of an example w from T depends only on $inf_{\mathcal{M}}(w)$. The algorithm first computes the set $\{inf_{\mathcal{M}}(w) \mid (w, 1) \in T\}$, sorts its elements into decreasing order C_1, C_2, \dots, C_ℓ using \preceq , and for each $i \in [1..\ell]$ chooses a positive example $(z_i, 1) \in T$ with $C_i = inf_{\mathcal{M}}(z_i)$.

At each stage k of the learning algorithm, β_k is a singleton normal form acceptance condition. Initially, $\beta_0 = \emptyset$ (which is satisfied by no words) and $k = 0$.

The main loop processes the positive examples z_i for $i = 1, 2, \dots, \ell$. If z_i is accepted by (\mathcal{M}, β_k) , then the algorithm goes on to the next positive example. Otherwise, we say that the example z_i *causes the update of β_k* . Let $G = inf_{\mathcal{M}}(z_i)$ and $B = Q \setminus G$ and let S_k be the set of pairs (q, B) such that $q \in G$ and there is no negative example $(w, 0)$ in T such that w satisfies (q, B) . Then β_{k+1} is set to $\beta_k \cup S_k$ and k is set to $k + 1$.

When the positive examples z_1, z_2, \dots, z_ℓ have been processed, let $\beta = \beta_k$ for the final value of k . If the Rabin acceptor (\mathcal{M}, β) is consistent with T , then it is returned. If not, the learning algorithm returns the default acceptor of type DRA for T .

Proposition 9.2. *Let \mathcal{R} be an IRA in singleton normal form with acceptance condition α and assume that the input \mathcal{M} is an automaton isomorphic to the automaton of \mathcal{R} . Assume the sample T is consistent with \mathcal{R} . For each k , if z_i is the example that causes the update of β_k , then (\mathcal{M}, β_k) accepts z_j for all $j < i$. Moreover, (\mathcal{M}, β) is consistent with T .*

Proof. No pair (q, B) is added to the acceptance condition if there is a negative example in T that satisfies it, so (\mathcal{M}, β) is consistent with all the negative examples in T .

Consider any positive example $(w, 1)$ from T . There exists i with $inf_{\mathcal{M}}(w) = inf_{\mathcal{M}}(z_i)$. When z_i is processed by the algorithm, if it is already accepted by the current (\mathcal{M}, β_k) , then it (and the word w) is also accepted by every subsequent hypothesis, including (\mathcal{M}, β) , because pairs are not removed from β_k .

Otherwise, z_i causes the update to β_k , and among the pairs in S_k that are added to β_k there is at least one that z_i satisfies. To see this, note that z_i must satisfy some pair in α , say (q, B_j) . Thus, $q \in inf(z_i)$ and $B_j \cap inf(z_i) = \emptyset$. The pair (q, B) where $B = Q \setminus inf(z_i)$ has $B_j \subseteq B$. Thus any ω -word that satisfies (q, B) will also satisfy (q, B_j) . Because T is consistent with \mathcal{R} , there can be no negative example in T satisfying (q, B) , so pair (q, B) is part of S_k and is added to β_k . The word z_i (and the word w) satisfies (q, B) and is therefore accepted by $(\mathcal{M}, \beta_{k+1})$ and every subsequent hypothesis, including (\mathcal{M}, β) . \square

9.4. Constructing T_{Acc}^{IRA} . In this section we describe the construction of the sample T_{Acc}^{IRA} , which conveys the acceptance condition of an IRA. Let $\mathcal{R} = \langle \Sigma, Q, q_\iota, \delta, \alpha \rangle$ be a deterministic complete IRA of n states in singleton normal form, and let \mathcal{M} be its automaton. The construction of the sample T_{Acc}^{IRA} proceeds in stages, simulating the learning algorithm \mathbf{L}_{Acc}^{IRA}

on the portion of the sample constructed so far to determine what examples still need to be added.

Initially, $\gamma_0 = \emptyset$ and $k = 0$. The acceptance condition γ_k tracks the learning algorithm's β_k . The set of words accepted by (\mathcal{M}, γ_k) is always a subset of the set of words accepted by (\mathcal{M}, α) . The main loop is as follows. If (\mathcal{M}, γ_k) is equivalent to (\mathcal{M}, α) then the construction of T_{Acc}^{IRA} is complete.

Otherwise, let D_k be the set of ω -words that satisfy α but not γ_k . Let C be the \preceq -largest set in $\{\text{inf}(w) \mid w \in D_k\}$, and let $w_{k+1} = \mathbf{Witness}(C, \mathcal{M})$, an ultimately periodic word of length $O(n^2)$. Then w_{k+1} is added as a positive example to T_{Acc}^{IRA} .

Let $B = Q \setminus \text{inf}(w_{k+1})$. Define P_k to be the set of all (q, B) such that $q \in \text{inf}(w_{k+1})$ and there is no ω -word w' that satisfies (q, B) but not α . Set $\gamma_{k+1} = \gamma_k \cup P_k$.

For each $q \in \text{inf}(w_{k+1})$ such that there is some ω -word w' that satisfies (q, B) but not α , let $u(v)^\omega = \mathbf{Witness}(\text{inf}(w'), \mathcal{M})$ for some such w' and include $(u(v)^\omega, 0)$ as a negative example in T_{Acc} . The example $u(v)^\omega$ is of length $O(n^2)$. Then set k to $k + 1$ and continue with the main loop.

We prove a polynomial bound on the number of examples added to the sample T_{Acc}^{IRA} , thus showing that its length is bounded by a polynomial in the size of \mathcal{R} .

Proposition 9.3. *If the acceptance condition α is in singleton normal form and has m pairs, then at most m positive examples and at most $m|Q|$ negative examples are added to T_{Acc} .*

Proof. We say an acceptance condition γ covers a pair (q, B) iff every ω -word w that satisfies (q, B) also satisfies γ . We will show that after each positive example w_{k+1} is added to T_{Acc} , the condition γ_{k+1} covers at least one pair in α that was not covered by γ_k .

Suppose not, and let $k + 1$ be the least index for which γ_{k+1} does not cover a pair of α that was not covered by γ_k . Because w_{k+1} is an example that satisfies α but not γ_k , there must be a pair (q, B_j) of α that is satisfied by w_{k+1} . Note that γ_k does not cover the pair (q, B_j) . Then $q \in \text{inf}(w_{k+1})$ and letting $B = Q \setminus \text{inf}(w_{k+1})$, $B_j \subseteq B$. The pair (q, B) will be added to γ_k in constructing γ_{k+1} because every word that satisfies (q, B) also satisfies (q, B_j) .

If γ_{k+1} does not cover (q, B_j) , there must be a word w' that satisfies (q, B_j) but not (q, B) . So $q \in \text{inf}(w')$ and $B_j \cap \text{inf}(w') = \emptyset$ but $B \cap \text{inf}(w') \neq \emptyset$. Let $B' = Q \setminus \text{inf}(w')$, so $B_j \subseteq B'$. We have $B \cap B' \subsetneq B$. Because $\text{inf}(w_{k+1})$ and $\text{inf}(w')$ are SCCs that overlap in q , their union is an SCC as well. Let w'' be an ω -word such that $\text{inf}(w'')$ is the union of $\text{inf}(w_{k+1})$ and $\text{inf}(w')$. Then $Q \setminus \text{inf}(w'') = B \cap B'$. Note that w'' satisfies (q, B_j) because $B_j \subseteq B \cap B'$ and thus is a positive example of α .

Because $B \cap B'$ is a proper subset of B , $\text{inf}(w'')$ is a proper superset of $\text{inf}(w_{k+1})$ and the positive example w'' would have been considered before w_{k+1} in the construction of T_{Acc} . (We can imagine all the positive examples of α being considered in order to find a maximum positive counterexample at each stage.) At that time, it was either passed over because (1) the current γ_r already covered it, or (2) it contributed a new pair to the current γ_r to yield γ_{r+1} .

In case (1), there is some pair (q'', B'') in γ_r that is satisfied by w'' . Then $q'' \in \text{inf}(w'')$ and $B'' \cap \text{inf}(w'') = \emptyset$. Recall $\text{inf}(w'')$ is the union of $\text{inf}(w_{k+1})$ and $\text{inf}(w')$. Thus, $B'' \cap \text{inf}(w_{k+1}) = \emptyset$ and $B'' \cap \text{inf}(w') = \emptyset$. Note that $q'' \in \text{inf}(w_{k+1})$ or $q'' \in \text{inf}(w')$. If $q'' \in \text{inf}(w_{k+1})$, w_{k+1} satisfies the pair (q'', B'') in γ_j , a contradiction, because w_{k+1} is not

accepted by γ_k and $r \leq k$. And if $q'' \in \text{inf}(w')$ then w' satisfies the pair (q'', B'') in γ_r , a contradiction, because w' is not accepted by γ_{k+1} and $r \leq k$.

In case (2), the positive example w'' contributes at least one term (q'', B'') to γ_{r+1} . In this case $B'' = B \cap B'$ and $q'' \in \text{inf}(w'')$. Thus, $q'' \in \text{inf}(w_{k+1})$ or $q'' \in \text{inf}(w')$, so w_{k+1} or w' satisfies the term (q'', B'') of γ_{r+1} , a contradiction because $r + 1 \leq k$ and neither w_{k+1} nor w' is covered by γ_k .

Thus, each positive example added to T_{Acc} covers a new pair of α , and at most m positive examples can be added. Each positive example added requires at most $|Q|$ negative examples to avoid adding incorrect pairs, so at most $m|Q|$ negative examples are added. \square

9.5. Correctness of \mathbf{L}_{Acc}^{IRA} . We prove the correctness of the learning algorithm \mathbf{L}_{Acc}^{IRA} and show that the classes \mathbb{IRA} and \mathbb{ISA} are identifiable in the limit using polynomial time and data.

Theorem 9.4. *Algorithm \mathbf{L}_{Acc}^{IRA} runs in time polynomial in the sizes of the inputs \mathcal{M} and T . Let \mathcal{R} be an IRA. If the input automaton \mathcal{M} is isomorphic to the automaton of \mathcal{R} , and the sample T is consistent with \mathcal{R} and subsumes T_{Acc}^{IRA} , then algorithm \mathbf{L}_{Acc}^{IRA} returns an IRA (\mathcal{M}, β) equivalent to \mathcal{R} .*

Proof. By Proposition 5.3, \mathbf{L}_{Acc}^{IRA} can construct the sequence z_1, z_2, \dots, z_ℓ and the successive acceptance conditions β_k in time polynomial in the size of \mathcal{M} and the length of T .

Assume \mathcal{R} is an IRA, that \mathcal{M} is isomorphic to the automaton of \mathcal{R} , and that the sample T is consistent with \mathcal{R} and subsumes T_{Acc}^{IRA} . For ease of notation, we assume that the isomorphism is the identity.

We show by induction that for each k , the acceptance condition β_k in the learning algorithm \mathbf{L}_{Acc}^{IRA} is the same as the acceptance condition γ_k in the construction of T_{Acc}^{IRA} . This is true for $k = 0$ because $\beta_0 = \gamma_0 = \emptyset$.

Assume that $\beta_k = \gamma_k$ for some $k \geq 0$. If (\mathcal{M}, γ_k) is equivalent to $\mathcal{R} = (\mathcal{M}, \alpha)$, then also (\mathcal{M}, β_k) is equivalent to \mathcal{R} , and none of the remaining positive examples cause any additions to β_k . Thus this is the final value of k , so $\beta = \beta_k$ and (\mathcal{M}, β) is equivalent to \mathcal{R} .

If (\mathcal{M}, γ_k) accepts a proper subset of the language accepted by (\mathcal{M}, α) , then in the construction of sample T_{Acc}^{IRA} , D_k is equal to the ω -words accepted by (\mathcal{M}, α) but not by (\mathcal{M}, γ_k) . This causes the positive example $(w_{k+1}, 1)$ to be added to T_{Acc}^{IRA} , where $\text{inf}_{\mathcal{M}}(w_{k+1})$ is \preceq -largest in the set $\{\text{inf}_{\mathcal{M}}(w) \mid w \in D_k\}$.

In the learning algorithm, because (\mathcal{M}, β_k) does not accept w_{k+1} , Proposition 9.2 implies that there must be an example z_i that causes the update to β_k , and all of the examples z_1, \dots, z_{i-1} are accepted by (\mathcal{M}, β_k) . Because for every positive example $(w, 1)$ in T there exists j such that $\text{inf}_{\mathcal{M}}(w) = \text{inf}_{\mathcal{M}}(z_j)$, there must be some r such that $\text{inf}_{\mathcal{M}}(w_{k+1}) = \text{inf}_{\mathcal{M}}(z_r)$. Moreover, $i \leq r$.

If $i < r$, then $\text{inf}_{\mathcal{M}}(w_i)$ is strictly \preceq -larger than $\text{inf}_{\mathcal{M}}(w_r)$, which contradicts the choice of w_{k+1} by the sample construction procedure, because z_i is accepted by (\mathcal{M}, α) but not (\mathcal{M}, γ_k) . Thus $i = r$ and the example $z_i = w_{k+1}$ is the element that causes the update to β_k . The negative examples included in T_{Acc}^{IRA} for the positive example w_{k+1} ensure that the update to β_k is the same as the update to γ_k , and $\beta_{k+1} = \gamma_{k+1}$.

Because β_k and γ_k are equal for all k , for the final value of k , $\beta = \beta_k = \gamma_k$, and therefore (\mathcal{M}, β) is equivalent to \mathcal{R} . Because the IRA (\mathcal{M}, β) is consistent with T , it is the acceptor returned by \mathbf{L}_{Acc}^{IRA} . \square

Theorem 9.5. *The classes $\mathbb{I}\mathbb{R}\mathbb{A}$ and $\mathbb{I}\mathbb{S}\mathbb{A}$ are identifiable in the limit using polynomial time and data.*

Proof. By Proposition 5.1 it suffices to prove this for $\mathbb{I}\mathbb{R}\mathbb{A}$. Let \mathcal{R} be an IRA in singleton normal form accepting the language L . The characteristic sample $T_L = T_{Aut} \cup T_{Acc}^{\mathbb{I}\mathbb{R}\mathbb{A}}$ is of size polynomial in the size of \mathcal{R} .

The combined learning algorithm $\mathbf{L}^{\mathbb{I}\mathbb{R}\mathbb{A}}$ with a sample T as input first runs \mathbf{L}_{Aut} on T to get a deterministic complete automaton \mathcal{M} and then runs $\mathbf{L}_{Acc}^{\mathbb{I}\mathbb{R}\mathbb{A}}$ on inputs \mathcal{M} and T and returns the resulting acceptor. $\mathbf{L}^{\mathbb{I}\mathbb{R}\mathbb{A}}$ runs in time polynomial in the length of T and returns a DRA consistent with T .

Now assume that the sample T is consistent with \mathcal{R} and subsumes T_L . Then by Theorem 6.2, the automaton \mathcal{M} is isomorphic to the automaton of \mathcal{R} . By Theorem 9.4, the acceptor returned by $\mathbf{T}^{\mathbb{I}\mathbb{R}\mathbb{A}}$ is an IRA (\mathcal{M}, β) that is equivalent to \mathcal{R} , and this is the acceptor also returned by $\mathbf{L}^{\mathbb{I}\mathbb{R}\mathbb{A}}$. \square

10. CONSTRUCTING CHARACTERISTIC SAMPLES IN POLYNOMIAL TIME

The definition of identification in the limit using polynomial time and data requires that a characteristic sample exist and be of polynomial size, but says nothing about the cost of computing it. An additional desirable property is that a characteristic sample be computable in polynomial time given an acceptor \mathcal{A} as input. Recall that when this holds, we say that the class is efficiently teachable. We now show that given an acceptor that is fully informative we can design efficient teachers, i.e. algorithms that run in polynomial time and compute the characteristic samples we have defined. This is conditioned on having polynomial time algorithms for equivalence (that are given in Sections 11-14). To claim the class $\mathbb{I}\mathbb{X}\mathbb{A}$ is efficiently teachable we also need to show that we can construct such sets when starting with an acceptor that is not, say an IBA, but has an equivalent IBA acceptor (and similarly for the other classes). This is done in Sections 15-16.

10.1. Computing T_{Aut} . For T_{Aut} , we need to be able to decide for two states q_1 and q_2 of an acceptor \mathcal{A} whether there exists an ω -word that distinguishes them, and if so, to return one such word. We are thus led to consider the problems of inclusion and equivalence.

The problems of inclusion and equivalence. The *inclusion problem* is the following. Given as input two ω -acceptors \mathcal{A}_1 and \mathcal{A}_2 over the same alphabet, determine whether the language accepted by \mathcal{A}_1 is a subset of the language accepted by \mathcal{A}_2 , that is, whether $\llbracket \mathcal{A}_1 \rrbracket \subseteq \llbracket \mathcal{A}_2 \rrbracket$. If so, the answer should be “yes”; if not, the answer should be “no” and a *witness*, that is, an ultimately periodic ω -word $u(v)^\omega$ accepted by \mathcal{A}_1 but rejected by \mathcal{A}_2 .

The *equivalence problem* is similar: the input is two ω -acceptors \mathcal{A}_1 and \mathcal{A}_2 over the same alphabet, and the problem is to determine whether they are equivalent, that is, whether $\llbracket \mathcal{A}_1 \rrbracket = \llbracket \mathcal{A}_2 \rrbracket$. If so, the answer should be “yes”; if not, the answer should be “no” and a witness, that is, an ultimately periodic ω -word $u(v)^\omega$ that is accepted by exactly one of \mathcal{A}_1 and \mathcal{A}_2 .

If we have a procedure to solve the inclusion problem, at most two calls to it will solve the equivalence problem. We describe polynomial time algorithms to solve the inclusion problem for DBAs, DCAs and DPAs in Section 11 and 12, for DRAs and DSAs in Section 13, and for DMAs in Section 14. Referring to those sections, we obtain polynomial time algorithms

to solve the equivalence problem for DBAs, DCAs and DPAs from Theorem 11.1 and Theorem 12.2, for DRAs and DSAs from Theorem 13.2, and for DMAs from Theorem 14.6. Thus, we have the following.

Theorem 10.1. *Given an acceptor \mathcal{A} of type IBA, ICA, IPA, IRA, ISA, or IMA, the sample T_{Aut} for the automaton portion of \mathcal{A} can be computed in polynomial time.*

Proof. Given an acceptor \mathcal{A} and two states q_1 and q_2 , to determine whether there is an ω -word that distinguishes them, we call the relevant polynomial time equivalence algorithm on the acceptors \mathcal{A}^{q_1} and \mathcal{A}^{q_2} , which returns a distinguishing word $u(v)^\omega$ if they are not equivalent. \square

10.2. **Computing T_{Acc} .** For T_{Acc} , the requirements depend on the type of acceptor.

Proposition 10.2. *Given an IBA \mathcal{B} , the sample T_{Acc}^{IBA} can be computed in time polynomial in the size of \mathcal{B} .*

Proof. Given an IBA $\mathcal{B} = \langle \Sigma, Q, q_\nu, \delta, F \rangle$ with automaton \mathcal{M} , the sample T_{Acc}^{IBA} described in Section 7.2 is computed as follows. For each SCC $C \in \text{maxSCCs}(Q \setminus F)$, let $u(v)^\omega = \mathbf{Witness}(C, \mathcal{M})$ and include the negative example $(u(v)^\omega, 0)$ in T_{Acc}^{IBA} .

To see that these examples are sufficient, suppose that $q \in Q$ and $w \in \Sigma^\omega$ are such that \mathcal{B} rejects w and $q \in \text{inf}(w)$. Then $D = \text{inf}(w)$ is an SCC of \mathcal{B} contained in $Q \setminus F$, so it is contained in some $C \in \text{maxSCCs}(Q \setminus F)$, and there is a negative example $(u(v)^\omega, 0)$ in T_{Acc}^{IBA} such that $\text{inf}(u(v)^\omega) = C$. Because $D \subseteq C$, we have $q \in C$. \square

Proposition 10.3. *Given an IPA \mathcal{P} , the sample T_{Acc}^{IPA} can be computed in time polynomial in the size of \mathcal{P} .*

Proof. Given an IPA \mathcal{P} with automaton \mathcal{M} , the computation of T_{Acc}^{IPA} proceeds as described in Section 8.3. That is, the canonical forest $\mathcal{F}^*(\mathcal{P})$ is computed in polynomial time, and for each node C in the forest, $u(v)^\omega = \mathbf{Witness}(C, \mathcal{M})$ is computed and $(u(v)^\omega, l)$ is added to T_{Acc}^{IPA} , where l is the label of node C in the canonical forest. \square

Proposition 10.4. *Given an IRA \mathcal{R} , the sample T_{Acc}^{IRA} can be computed in time polynomial in the size of \mathcal{R} .*

Proof. Given an IRA $\mathcal{R} = (\mathcal{M}, \alpha)$, the computation of T_{Acc} proceeds as described in Section 9.4. At each stage of the computation, it is necessary to find an ω -word $u(v)^\omega$ with the \preceq -largest $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ that is accepted by (\mathcal{M}, α) and rejected by (\mathcal{M}, γ_k) . Theorem 13.2 gives a polynomial time algorithm that not only tests the inclusion of two DRAs, but returns a witness $u(v)^\omega$ with the \preceq -largest $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ in the case of non-inclusion, because \mathcal{M} is isomorphic to $\mathcal{M} \times \mathcal{M}$. \square

Proposition 10.5. *Given an IMA \mathcal{A} , the sample T_{Acc}^{IMA} can be computed in time polynomial in the size of \mathcal{A} .*

Proof. Given an IMA $\mathcal{A} = (\mathcal{M}, \mathcal{F})$, the sample T_{Acc}^{IMA} described in Section 7.1 is computed as follows. For each $F \in \mathcal{F}$ determine whether F is a reachable SCC of \mathcal{M} , and if so, compute $u(v)^\omega = \mathbf{Witness}(F, \mathcal{M})$ and add $(u(v)^\omega, 1)$ to the sample T_{Acc}^{IMA} . \square

Theorem 10.6. *Let \mathcal{A} be an IBA, IPA, IRA, or IMA accepting the ω -language L . Then the characteristic sample T_L for \mathcal{A} can be computed in polynomial time in the size of \mathcal{A} .*

Proof. By Theorem 10.1, the sample T_{Aut} can be computed in polynomial time in the size of \mathcal{A} , and by Proposition 10.2, 10.3, 10.4, or 10.5 the sample T_{Acc}^{IBA} , T_{Acc}^{IPA} , T_{Acc}^{IRA} , or T_{Acc}^{IMA} can also be computed in polynomial time in the size of \mathcal{A} . \square

Note that Theorem 10.6 does not imply that the class $\mathbb{X}\mathbb{A}$ for $\mathbb{X} \in \{\mathbb{B}, \mathbb{P}, \mathbb{R}, \mathbb{M}\}$ is efficiently teachable, since this class also has representations by non-isomorphic automata.

11. INCLUSION ALGORITHMS

We show that there are polynomial time algorithms for the inclusion problem for DBAs, DCAs, DPAs, DRAs, DSAs and DMAs. Recall that two calls to an inclusion algorithm suffice to solve the equivalence problem. By Claim 2.2 (4), the inclusion and equivalence problems for DCAs are efficiently reducible to those for DBAs, and vice versa. Also, by Claim 2.2 (1), the inclusion and equivalence problems for DBAs are efficiently reducible to those for DPAs. Thus it suffices to consider the inclusion problem for DPAs, DRAs and DMAs.

Remark. In the case of DFAs, a polynomial algorithm for the inclusion problem can be obtained using polynomial algorithms for complementation, intersection and emptiness (since for any two languages $L_1 \subseteq L_2$ if and only if $L_1 \cap \overline{L_2} = \emptyset$). However, a similar approach does not work in the case of DPAs; although complementation and emptiness for DPAs can be computed in polynomial time, intersection cannot [Bok18, Theorem 9].

For the inclusion problem for DBAs, DCAs and DPAs, Schewe [Sch10, Sch11] gives the following result.

Theorem 11.1 [Sch10]. *The inclusion problems for DBAs, DCAs and DPAs are in NL.*

Because NL (nondeterministic logarithmic space) is contained in polynomial time, this implies the existence of polynomial time inclusion and equivalence algorithms for DBAs, DCAs and DPAs. For the sake of completeness, and to address the problem of returning a witness we include a proof sketch.

Proof sketch. For $i = 1, 2$, let $\mathcal{P}_i = \langle \Sigma, Q_i, (q_l)_i, \delta_i, \kappa_i \rangle$ be a DPA. It suffices to guess two states $q_1 \in Q_1$ and $q_2 \in Q_2$, and two words $u \in \Sigma^*$ and $v \in \Sigma^+$, and to check that for $i = 1, 2$, $\delta_i((q_l)_i, u) = q_i$ and $\delta_i(q_i, v) = q_i$, and also, that the smallest value of $\kappa_1(q)$ in the loop in \mathcal{P}_1 from q_1 to q_1 on input v is odd, while the smallest value of $\kappa_2(q)$ in the loop in \mathcal{P}_2 from q_2 to q_2 on input v is even. If these checks succeed, then $\llbracket \mathcal{P}_1 \rrbracket$ is not a subset of $\llbracket \mathcal{P}_2 \rrbracket$, and the ultimately periodic word $u(v)^\omega$ is a witness.

Logarithmic space is enough to record the two guessed states q_1 and q_2 as well as the current minimum values of κ_1 and κ_2 as the loops on v are traversed in the two automata. The words u and v need only be guessed symbol-by-symbol, using a pointer in each automaton to keep track of its current state. \square

This approach does not seem to work in the case of testing DRA or DMA inclusion, because the acceptance conditions would seem to require keeping track of more information than would fit in logarithmic space. To supplement the proof sketch for Schewe's theorem, in the next section (Section 12) we give an explicit polynomial time algorithm for testing DPA inclusion.

For inclusion of DRAs and DMAs, [CDK93] provides a reduction to the problem of model checking a formula in the temporal logic CTL*. While the complexity of model checking

CTL* formulas is in general PSPACE-hard, for the fair-CTL fragment, [EL87] provides a model checking algorithm that runs in polynomial time. It is further shown in [CDK93] that the CTL* formulas they reduce to can be modified to formulas in a fragment slightly extending fair-CTL that can still be handled by the model checking algorithm of [EL87] for fair-CTL. Thus overall this gives a polynomial time algorithm for inclusion of DRAs and DMAs. Since our learning algorithm relies on the sample including shortlex examples and the algorithm above does not guarantee shortlex counterexamples, we give in Section 13 and Section 14 polynomial time automata-theoretic algorithms for testing inclusion for DRAs and DMAs that provide shortlex counterexamples, which are novel results.

12. INCLUSION AND EQUIVALENCE FOR DPAs, DBAs, DCAs

In this section we describe an explicit polynomial time algorithm for the inclusion problem for two DPAs, which yields algorithms for DBAs and DCAs. If $\mathcal{P} = \langle \Sigma, Q, q_t, \delta, \kappa \rangle$ is a complete DPA and $w \in \Sigma^\omega$, we let $\mathcal{P}(w)$ denote the minimum color visited by \mathcal{P} infinitely often on input w , that is, $\mathcal{P}(w) = \kappa(\inf(w))$.

12.1. Searching for w with given minimum colors in two acceptors. We first describe an algorithm that searches for an ω -word that yields specified minimum colors in two different DPAs over the same alphabet.

For $i = 1, 2$, let $\mathcal{P}_i = \langle \Sigma, Q_i, (q_t)_i, \delta_i, \kappa_i \rangle$ be a DPA, and let \mathcal{M}_i be the automaton of \mathcal{P}_i . Given inputs of \mathcal{P}_1 and \mathcal{P}_2 and two nonnegative integers k_1 and k_2 , the **Colors** algorithm constructs the product automaton $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ and the set $Q' = \{(q_1, q_2) \in Q_1 \times Q_2 \mid \kappa_1(q_1) \geq k_1 \wedge \kappa_2(q_2) \geq k_2\}$.

The algorithm then computes $S = \text{maxSCCs}(Q')$ for the automaton \mathcal{M} , and loops through the SCCs $C \in S$ checking whether C is reachable in \mathcal{M} , $\min(\kappa_1(\pi_1(C))) = k_1$, and $\min(\kappa_2(\pi_2(C))) = k_2$. If so, it returns the ultimately periodic word $u(v)^\omega = \mathbf{Witness}(C, \mathcal{M})$. If none of the elements $C \in S$ satisfies this condition, then the answer “no” is returned.

Theorem 12.1. *The algorithm **Colors** takes as input two DPAs \mathcal{P}_1 and \mathcal{P}_2 over the same alphabet and two nonnegative integers k_1 and k_2 , runs in polynomial time, and determines whether there exists an ω -word w such that $\mathcal{P}_1(w) = k_1$ and $\mathcal{P}_2(w) = k_2$. If not, it returns the answer “no”. If so, it returns an ultimately periodic ω -word $u(v)^\omega$ such that $\mathcal{P}_1(u(v)^\omega) = k_1$ and $\mathcal{P}_2(u(v)^\omega) = k_2$.*

Proof. The polynomial running time of the algorithm follows from Props. 5.4 and 5.7. To see the correctness of the algorithm, suppose first that it returns an ultimately periodic word $u(v)^\omega$. This occurs only if it finds an SCC C of \mathcal{M} such that C is reachable in \mathcal{M} , $\kappa_1(\pi_1(C)) = k_1$, and $\kappa_2(\pi_2(C)) = k_2$. Then for $i = 1, 2$, $\pi_i(C)$ is the set of states visited infinitely often by \mathcal{M}_i on the input $u(v)^\omega$, which has minimum color k_i .

To see that the algorithm does not incorrectly answer “no”, suppose w is an ω -word such that for $i = 1, 2$, $\mathcal{P}_i(w) = k_i$. Let $D_i = \inf_{\mathcal{M}_i}(w)$ be an SCC of \mathcal{M}_i . No state in D_i has a color less than k_i , so if $D = \inf_{\mathcal{M}}(w)$, then $D \subseteq Q'$. Also, D is a reachable SCC in \mathcal{M} .

Then D is contained in some element C of $\text{maxSCCs}(Q')$. Because there are no states (q_1, q_2) in C with $\kappa_1(q_1) < k_1$ or $\kappa_2(q_2) < k_2$, we must have $\kappa_i(\pi_i(C)) = k_i$ for $i = 1, 2$. Also, C is reachable in \mathcal{M} because D is. Thus, the algorithm will find at least one such C and return $u(v)^\omega$ such that $\inf_{\mathcal{M}}(u(v)^\omega) = C$. \square

12.2. An inclusion algorithm for DPAs. The inclusion problem for DPAs \mathcal{P}_1 and \mathcal{P}_2 over the same alphabet can be solved by looping over all odd k_1 in the range of κ_1 and all even k_2 in the range of κ_2 , calling the **Colors** algorithm with inputs \mathcal{P}_1 , \mathcal{P}_2 , k_1 , and k_2 . If the **Colors** algorithm returns any witness $u(v)^\omega$, then $u(v)^\omega \in \llbracket \mathcal{P}_1 \rrbracket \setminus \llbracket \mathcal{P}_2 \rrbracket$, and $u(v)^\omega$ is returned as a witness of non-inclusion. Otherwise, by Theorem 12.1, there is no ω -word w accepted by \mathcal{P}_1 and not accepted by \mathcal{P}_2 , and the answer “yes” is returned for the inclusion problem. Note that for $i = 1, 2$, the range of κ_i has at most $|Q_i|$ distinct elements. Thus we have the following.

Theorem 12.2. *There are polynomial time algorithms for the inclusion and equivalence problems for two DPAs over the same alphabet.*

From Claim 2.2 (1 and 4), we have the following.

Theorem 12.3. *There are polynomial time algorithms for the inclusion and equivalence problems for two DBAs (or DCAs) over the same alphabet.*

13. AN INCLUSION ALGORITHM FOR DRAs

In this section we describe a polynomial time algorithm to solve the inclusion problem for two DRAs. The algorithm returns a \preceq -largest witness in the case of non-inclusion.

Algorithm 1 **SubInc**^{DRA}

Input: Two DRAs $\mathcal{R}_1 = (\mathcal{M}_1, \alpha_1)$ and $\mathcal{R}_2 = (\mathcal{M}_2, \alpha_2)$ in singleton normal form, where $\alpha_1 = \{(q', B')\}$ and $\alpha_2 = \{(q''_1, B''_1), \dots, (q''_k, B''_k)\}$, and a set S of states of $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$.

Output: $u(v)^\omega \in \llbracket \mathcal{R}_1 \rrbracket \setminus \llbracket \mathcal{R}_2 \rrbracket$ with $\text{inf}_{\mathcal{M}}(u(v)^\omega) \subseteq S$ if such exists, else “none”.

$\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$

$W \leftarrow \emptyset$

$S' \leftarrow S \setminus \{(q_1, q_2) \in S \mid q_1 \in B'\}$

$\mathcal{C} \leftarrow \text{maxSCCs}(S')$

for each reachable $C \in \mathcal{C}$ such that $q' \in \pi_1(C)$ **do**

if for no $j \in [1..k]$ is $q''_j \in \pi_2(C)$ and $B''_j \cap \pi_2(C) = \emptyset$ **then**

$W \leftarrow W \cup \{\mathbf{Witness}(C, \mathcal{M})\}$

 ▷ A new candidate witness

else

$J = \{q''_j \mid j \in [1..k], B''_j \cap \pi_2(C) = \emptyset\}$

$S'' \leftarrow C \setminus \{(q_1, q_2) \in C \mid q_2 \in J\}$

 Call **SubInc**^{DRA} recursively with \mathcal{R}_1 , \mathcal{R}_2 , and S''

if the returned value is $u(v)^\omega$ **then**

$W \leftarrow W \cup \{u(v)^\omega\}$

if W is \emptyset **then**

return “none”

else

 Let $u(v)^\omega \in W$ have the \preceq -largest value of $\text{inf}_{\mathcal{M}}(u(v)^\omega)$

return $u(v)^\omega$

The algorithm **SubInc**^{DRA} takes as input two DRAs $\mathcal{R}_1 = (\mathcal{M}_1, \alpha_1)$ and $\mathcal{R}_2 = (\mathcal{M}_2, \alpha_2)$ in singleton normal form, where α_1 consists of a single pair (q', B') . It also takes as input a subset S of the state set of the product automaton $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$. The problem it solves is to determine whether there exists an ω -word $u(v)^\omega$ with $\text{inf}_{\mathcal{M}}(u(v)^\omega) \subseteq S$ such that

$u(v)^\omega \in \llbracket \mathcal{R}_1 \rrbracket \setminus \llbracket \mathcal{R}_2 \rrbracket$. If there is such a word, the algorithm returns one with the \preceq -largest value of $\text{inf}_{\mathcal{M}}(u(v)^\omega)$, and otherwise, it returns “none”.

Proposition 13.1. *For $i = 1, 2$ let $\mathcal{R}_i = (\mathcal{M}_i, \alpha_i)$ be a DRA in singleton normal form. Assume $\alpha_1 = \{(q', B')\}$ and $\alpha_2 = \{(q''_1, B''_1), \dots, (q''_k, B''_k)\}$. Let $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$, and let S be a subset of the states of \mathcal{M} . Then with inputs $\mathcal{R}_1, \mathcal{R}_2$, and S , the algorithm **SubInc**^{DRA} runs in polynomial time and returns $u(v)^\omega \in \llbracket \mathcal{R}_1 \rrbracket \setminus \llbracket \mathcal{R}_2 \rrbracket$ with the \preceq -largest value of $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ contained in S , if such exists, else it returns “none”.*

Proof. When the element $u(v)^\omega = \mathbf{Witness}(\mathcal{M}, C)$ is added to W , we have that C is a reachable SCC of \mathcal{M} contained in S , $q' \in \pi_1(C)$, and $B' \cap \pi_1(C) = \emptyset$ (because S' contains no elements (q_1, q_2) with $q_1 \in B'$), so $(\mathcal{M}_1, \{(q', B')\})$ accepts $u(v)^\omega$. Also, we have that for no $j \in [1..k]$ do we have $q''_j \in \pi_2(C)$ and $B''_j \cap \pi_2(C) = \emptyset$, so \mathcal{R}_2 rejects $u(v)^\omega$. Thus, any returned $u(v)^\omega$ is a witness to the non-inclusion of $\llbracket (\mathcal{M}_1, \{(q', B')\}) \rrbracket$ in $\llbracket \mathcal{R}_2 \rrbracket$ with $\text{inf}_{\mathcal{M}}(u(v)^\omega) \subseteq S$.

We now show by induction on the recursive calls that if w is any ω -word such that $\text{inf}_{\mathcal{M}}(w) \subseteq S$, $(\mathcal{M}_1, \{(q', B')\})$ accepts w , and \mathcal{R}_2 rejects w , then **SubInc**^{DRA} returns a witness $u(v)^\omega$ such that $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ is at least as large as $\text{inf}_{\mathcal{M}}(w)$ in the \preceq -ordering. Let $D = \text{inf}_{\mathcal{M}}(w)$. Then D is a reachable SCC of \mathcal{M} such that $q' \in \pi_1(D)$, $B' \cap \pi_1(D) = \emptyset$, and for no $j \in [1..k]$ do we have $q''_j \in \pi_2(D)$ and $B''_j \cap \pi_2(D) = \emptyset$. Then $D \subseteq S'$ because $B' \cap \pi_1(D) = \emptyset$. Thus, D must be a subset of exactly one of the elements C of $\text{maxSCCs}(S')$. Then C is reachable, $q' \in \pi_1(C)$, and $B' \cap \pi_1(C) = \emptyset$ (because C is a subset of S').

If C is such that for no $j \in [1..k]$ do we have $q''_j \in \pi_2(C)$ and $B''_j \cap \pi_2(C) = \emptyset$, then a witness $u(v)^\omega = \mathbf{Witness}(C, \mathcal{M})$ is added to W , and we have that $C = \text{inf}_{\mathcal{M}}(u(v)^\omega)$ is at least as large in the \preceq -ordering as $D = \text{inf}_{\mathcal{M}}(w)$, because $D \subseteq C$.

Otherwise, the set $J = \{q''_j \mid j \in [1..k], B''_j \cap \pi_2(C) = \emptyset\}$ is non-empty, and the algorithm removes from C all the states (q_1, q_2) such that $q_2 \in J$ to form the set S'' . Because $D \subseteq C$, if $B''_j \cap \pi_2(C) = \emptyset$, then also $B''_j \cap \pi_2(D) = \emptyset$. Thus, if for any $q''_j \in J$ we have $q''_j \in \pi_2(D)$, this would violate the assumption that \mathcal{R}_2 rejects w . Hence, $D \subseteq S''$, and by the inductive assumption on the recursive calls, the recursive call to **SubInc**^{DRA} returns a witness $u(v)^\omega$ such that $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ is at least as large in the \preceq -ordering as $\text{inf}_{\mathcal{M}}(w)$. Because the top-level algorithm returns $u(v)^\omega$ to maximize $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ with respect to \preceq , it will be at least as large as $\text{inf}_{\mathcal{M}}(w)$.

For the polynomial running time, we note that all the SCCs C considered are distinct elements of a decreasing forest of SCCs for the automaton \mathcal{M} , and so there can be at most as many as the number of states of \mathcal{M} . \square

Theorem 13.2. *There are polynomial time algorithms to solve the inclusion and equivalence problems for two DRAs (resp. DSAs) \mathcal{R}_1 and \mathcal{R}_2 . In the case of non-inclusion or non-equivalence, these algorithms return a witness $u(v)^\omega$ with the \preceq -largest value of $\text{inf}_{\mathcal{M}}(u(v)^\omega)$, where \mathcal{M} is the product of the automata \mathcal{R}_1 and \mathcal{R}_2 .*

Proof. It suffices to consider just DRAs, by Claim 2.2 (5). Given two DRAs $\mathcal{R}_1 = (\mathcal{M}_1, \alpha_1)$ and $\mathcal{R}_2 = (\mathcal{M}_2, \alpha_2)$, we may assume they are in singleton normal form. Then for each pair (q_i, B_i) in α_1 , we call **SubInc**^{DRA} with inputs $(\mathcal{M}_1, \{(q_i, B_i)\})$, \mathcal{R}_2 , with S equal to the whole state set of \mathcal{M} . If all of these calls return “none”, then $\llbracket \mathcal{R}_1 \rrbracket$ is a subset of $\llbracket \mathcal{R}_2 \rrbracket$, and the answer returned is “yes”. Otherwise, one or more calls return a witness, and $u(v)^\omega$ is returned such that $\text{inf}_{\mathcal{M}}(u(v)^\omega)$ is \preceq -largest among the witnesses returned by the calls. The

running time and correctness follow from the running time and correctness guarantees of **SubInc**^{DRA}. \square

14. AN INCLUSION ALGORITHM FOR DMAS

In this section we develop a polynomial time algorithm to solve the inclusion problem for two DMAs over the same alphabet. The proof proceeds in two parts: (1) a polynomial time reduction of the inclusion problem for two DMAs to the inclusion problem for a DBA and a DMA, and (2) a polynomial time algorithm for the inclusion problem for a DBA and a DMA.

14.1. Reduction of DMA inclusion to DBA/DMA inclusion. We first reduce the problem of inclusion for two arbitrary DMAs to the inclusion problem for two DMAs where the first one has just a single final state set. For $i = 1, 2$, define the DMA $\mathcal{U}_i = \langle Q_i, \Sigma, (q_i)_i, \delta_i, \mathcal{F}_i \rangle$, where \mathcal{F}_i is the set of final state sets for \mathcal{U}_i . Let the elements of \mathcal{F}_1 be $\{F_1, \dots, F_k\}$, and for each $j \in [1..k]$, let

$$\mathcal{U}_{1,j} = \langle Q_1, \Sigma, (q_1)_1, \delta_1, \{F_j\} \rangle,$$

that is, $\mathcal{U}_{1,j}$ is \mathcal{U}_1 with F_j as its only final state set. Then by the definition of DMA acceptance,

$$\llbracket \mathcal{U}_1 \rrbracket = \bigcup_{j=1}^k \llbracket \mathcal{U}_{1,j} \rrbracket,$$

which implies that to test whether $\llbracket \mathcal{U}_1 \rrbracket \subseteq \llbracket \mathcal{U}_2 \rrbracket$, it suffices to test for all $j \in [1..k]$ that $\llbracket \mathcal{U}_{1,j} \rrbracket \subseteq \llbracket \mathcal{U}_2 \rrbracket$.

Proposition 14.1. *Suppose \mathbf{L} is a procedure that solves the inclusion problem for two DMAs over the same alphabet, assuming that the first DMA has a single final state set. Then there is an algorithm that solves the inclusion problem for two arbitrary DMAs over the same alphabet, say \mathcal{U}_1 and \mathcal{U}_2 , which simply makes $|\mathcal{F}_1|$ calls to \mathbf{L} , where \mathcal{F}_1 is the family of final state sets of \mathcal{U}_1 .*

Next we describe a procedure **SCCtoDBA** that takes as inputs a deterministic automaton \mathcal{M} , an SCC F of \mathcal{M} , and a state $q \in F$, and returns a DBA $B(\mathcal{M}, F, q)$ that accepts exactly $L(\mathcal{M}, F, q)$, where $L(\mathcal{M}, F, q)$ is the set of ω -words w that visit only the states of F when processed by \mathcal{M} starting at state q , and visits each of them infinitely many times.

Assume the states in F are $\{q_0, q_1, \dots, q_{m-1}\}$, where $q_0 = q$. The DBA $B(\mathcal{M}, F, q)$ is $\langle Q', \Sigma, q_0, \delta', \{q_0\} \rangle$, where we define Q' and δ' as follows. We create new states $r_{i,j}$ for $i, j \in [0..m-1]$ such that $i \neq j$, and denote the set of these by R . We also create a new dead state d_0 . Then the set of states Q' is $Q \cup R \cup \{d_0\}$.

For δ' , the dead state d_0 behaves as expected: for all $\sigma \in \Sigma$, $\delta'(d_0, \sigma) = d_0$. For the other states in Q' , let $\sigma \in \Sigma$ and $i \in [0..m-1]$. If $\delta(q_i, \sigma)$ is not in F , then in order to deal with runs that would visit states outside of F , we define $\delta'(q_i, \sigma) = d_0$ and, for all $j \neq i$, $\delta'(r_{i,j}, \sigma) = d_0$.

Otherwise, for some $k \in [0..m-1]$ we have $q_k = \delta(q_i, \sigma)$. If $k = (i+1) \bmod m$, then we define $\delta'(q_i, \sigma) = q_k$, and otherwise we define $\delta'(q_i, \sigma) = r_{k, (i+1) \bmod m}$. For all $j \in [0..m-1]$ with $j \neq i$, if $k = j$, we define $\delta'(r_{i,j}, \sigma) = q_k$, and otherwise we define $\delta'(r_{i,j}, \sigma) = r_{k,j}$.




Figure 4. Example of the construction of $B(\mathcal{M}, F, q)$ with $F = \{q_0, q_1, q_2\}$ and $q = q_0$.

Intuitively, for an input from $L(\mathcal{M}, F, q)$, in $B(\mathcal{M}, F, q)$ the states q_i are visited in a repeating cyclic order: q_0, q_1, \dots, q_{m-1} , and the meaning of the state $r_{i,j}$ is that at this point in the input, \mathcal{M} would be in state q_i , and the machine $B(\mathcal{M}, F, q)$ is waiting for a transition that would arrive at state q_j in \mathcal{M} , in order to proceed to state q_j in $B(\mathcal{M}, F, q)$.⁹ An example of the construction is shown in Figure 4; the dead state and unreachable states are omitted for clarity.

Lemma 14.2. *Let \mathcal{M} be a deterministic automaton with alphabet Σ and states Q , and let F be an SCC of \mathcal{M} and $q \in F$. With these inputs, the procedure **SCCtoDBA** runs in polynomial time and returns the DBA $B(\mathcal{M}, F, q)$, which accepts the language $L(\mathcal{M}, F, q)$ and has $|F|^2 + 1$ states.*

Proof. Suppose w is in $L(\mathcal{M}, F, q)$. Let $q = s_0, s_1, s_2, \dots$ be the sequence of states in the run of \mathcal{M} from state q on input w . This run visits only states in F and visits each one of them infinitely many times. We next define a particular increasing sequence $i_{k,\ell}$ of indices in s , where k is a positive integer and $\ell \in [0, m-1]$. These indices mark particular visits to the states q_0, q_1, \dots, q_{m-1} in repeating cyclic order. The initial value is $i_{1,0} = 0$, marking the initial visit to q_0 . If $i_{k,\ell}$ has been defined and $\ell < m-1$, then $i_{k,\ell+1}$ is defined as the least natural number j such that $j > i_{k,\ell}$ and $s_j = q_{\ell+1}$, marking the next visit to $q_{\ell+1}$. If $\ell = m-1$, then $i_{k+1,0}$ is defined as the least natural number j such that $j > i_{k,\ell}$ and $s_j = q_0$, marking the next visit to q_0 .

There is a corresponding division of w into a concatenation of finite segments $w_{1,1}, w_{1,2}, \dots, w_{1,m-1}, w_{2,0}, \dots$ between consecutive elements in the increasing sequence of indices. An inductive argument shows that in $B(\mathcal{M}, F, q)$, the prefix of w up through $w_{k,\ell}$ arrives at the state q_ℓ , so that w visits q_0 infinitely often and is therefore accepted by $B(\mathcal{M}, F, q)$.

Conversely, suppose $B(\mathcal{M}, F, q)$ accepts the ω -word w . Let s_0, s_1, s_2, \dots be the run of $B(\mathcal{M}, F, q)$ on w , and let t_0, t_1, t_2, \dots be the run of \mathcal{M} starting from q on input w . An inductive argument shows that if $s_n = q_i$ then $t_n = q_i$, and if $s_n = r_{i,j}$ then $t_n = q_i$. Because the only way the run s_0, s_1, \dots can visit the final state q_0 infinitely often is to progress through the states q_0, q_1, \dots, q_{m-1} in repeating cyclic order, the run t_0, t_1, \dots must visit only states in F and visit each of them infinitely often, so $w \in L(\mathcal{M}, F, q)$.

⁹This construction is reminiscent of the construction transforming a generalized Büchi into a Büchi automaton [Var08, Cho74], by considering each state in F as a singleton set of a generalized Büchi, but here we need to send transitions to states outside F to a sink state.

The DBA $B(\mathcal{M}, F, q)$ has a dead state, and $|F|$ states for each element of F , for a total of $|F|^2 + 1$ states. The running time of the procedure **SCCtoDBA** is linear in the size of \mathcal{M} and the size of the resulting DBA, which is polynomial in the size of \mathcal{M} . \square

We now show that this construction may be used to reduce the inclusion of two DMAs to the inclusion of a DBA and a DMA. Recall that if \mathcal{A} is an acceptor and q is a state of \mathcal{A} , then \mathcal{A}^q denotes the acceptor \mathcal{A} with the initial state changed to q .

Lemma 14.3. *Let \mathcal{U}_1 be a DMA with automaton \mathcal{M}_1 and a single final state set F_1 . Let \mathcal{U}_2 be an arbitrary DMA over the same alphabet as \mathcal{U}_1 , with automaton \mathcal{M}_2 and family of final state sets \mathcal{F}_2 . Let \mathcal{M} denote the product automaton $\mathcal{M}_1 \times \mathcal{M}_2$ with unreachable states removed. Then $\llbracket \mathcal{U}_1 \rrbracket \subseteq \llbracket \mathcal{U}_2 \rrbracket$ iff for every state (q_1, q_2) of \mathcal{M} with $q_1 \in F_1$ we have $\llbracket B(\mathcal{M}_1, F_1, q_1) \rrbracket \subseteq \llbracket \mathcal{U}_2^{q_2} \rrbracket$.*

Proof. Suppose that for some state (q_1, q_2) of \mathcal{M} with $q_1 \in F_1$, we have $w \in \llbracket B(\mathcal{M}_1, F_1, q_1) \rrbracket \setminus \llbracket \mathcal{U}_2^{q_2} \rrbracket$. Let C_1 be the set of states visited infinitely often in $B(\mathcal{M}_1, F_1, q_1)$ on input w , and let C_2 be the set of states visited infinitely often in $\mathcal{U}_2^{q_2}$ on input w . Then $Q_1 \cap C_1 = F_1$ and $C_2 \notin \mathcal{F}_2$. Let u be a finite word such that $\mathcal{M}(u) = (q_1, q_2)$. Then $\text{inf}_{\mathcal{M}_1}(uw) = Q_1 \cap C_1 = F_1$ and $\text{inf}_{\mathcal{M}_2}(uw) = C_2$, so $uw \in \llbracket \mathcal{U}_1 \rrbracket \setminus \llbracket \mathcal{U}_2 \rrbracket$.

Conversely, suppose that $w \in \llbracket \mathcal{U}_1 \rrbracket \setminus \llbracket \mathcal{U}_2 \rrbracket$. For $i = 1, 2$ let $C_i = \text{inf}_{\mathcal{M}_i}(w)$. Note that $C_1 = F_1$ and $C_2 \notin \mathcal{F}_2$. Let $w = xw'$, where x is a finite prefix of w that is sufficiently long that the run of \mathcal{M}_1 on w does not visit any state outside C_1 after x has been processed, and for $i = 1, 2$ let $q_i = \mathcal{M}_i(x)$. Then (q_1, q_2) is a (reachable) state of \mathcal{M} , $q_1 \in F_1$, and the ω -word w' , when processed by \mathcal{M}_1 starting at state q_1 visits only states of $C_1 = F_1$ and visits each of them infinitely many times, that is, $w' \in \llbracket B(\mathcal{M}_1, F_1, q_1) \rrbracket$. Moreover, when w' is processed by \mathcal{M}_2 starting at state q_2 , the set of states visited infinitely often is C_2 , which is not in \mathcal{F}_2 . Thus, $w' \in \llbracket B(\mathcal{M}_1, F_1, q_1) \rrbracket \setminus \llbracket \mathcal{U}_2^{q_2} \rrbracket$. \square

To turn this into an algorithm to test inclusion for two DMAs, \mathcal{U}_1 with automaton \mathcal{M}_1 and a single final state set F_1 that is an SCC of \mathcal{M}_1 and \mathcal{U}_2 with automaton \mathcal{M}_2 , we proceed as follows. Construct the product automaton $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ with unreachable states removed, and for each state (q_1, q_2) of \mathcal{M} , if $q_1 \in F_1$, construct the DBA $B(\mathcal{M}_1, F_1, q_1)$ and the DMA $\mathcal{U}_2^{q_2}$ and test the inclusion of language accepted by the DBA in the language accepted by the DMA. If all of these tests return “yes”, then the algorithm returns “yes” for the inclusion question for \mathcal{U}_1 and \mathcal{U}_2 . Otherwise, for the first test that returns “no” and a witness $u(v)^\omega$, the algorithm finds by breadth-first search a minimum length finite word u' such that $\mathcal{M}(u') = (q_1, q_2)$, and returns the witness $u'u(v)^\omega$.

Combining this with Proposition 14.1, we have the following.

Theorem 14.4. *Let \mathbf{L} be an algorithm to test inclusion for an arbitrary DBA and an arbitrary DMA over the same alphabet. There is an algorithm to test inclusion for an arbitrary pair of DMAs \mathcal{U}_1 and \mathcal{U}_2 over the same alphabet whose running time is linear in the sizes of \mathcal{U}_1 and \mathcal{U}_2 plus the time for at most $k \cdot |Q_1| \cdot |Q_2|$ calls to the procedure \mathbf{L} , where k is the number of final state sets in \mathcal{U}_1 , and Q_i is the state set of \mathcal{U}_i for $i = 1, 2$.*

14.2. A DBA/DMA inclusion algorithm. In this section, we give a polynomial time algorithm **DBAinDMA** to test inclusion for an arbitrary DBA and an arbitrary DMA over the same alphabet.

Assume the inputs are a DBA $\mathcal{B} = (\mathcal{M}_1, F_1)$ and a DMA $\mathcal{U} = (\mathcal{M}_2, \mathcal{F})$. The overall strategy of the algorithm is to seek an SCC C of $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ such that $\pi_1(C) \cap F_1 \neq \emptyset$ and $\pi_2(C) \notin \mathcal{F}$. If such a C is found, the algorithm calls **Witness**(C, \mathcal{M}), which returns $u(v)^\omega$ such that $\text{inf}_{\mathcal{M}}(u(v)^\omega) = C$. Because $\text{inf}_{\mathcal{M}_1}(u(v)^\omega) = \pi_1(C)$ and $\pi_1(C) \cap F_1 \neq \emptyset$, $u(v)^\omega \in \llbracket \mathcal{B} \rrbracket$, and because $\text{inf}_{\mathcal{M}_2}(u(v)^\omega) = \pi_2(C)$ and $\pi_2(C) \notin \mathcal{F}$, $u(v)^\omega \notin \llbracket \mathcal{U} \rrbracket$. The details are given in Algorithm 2.

Algorithm 2 DBAinDMA

Input: A DBA $\mathcal{B} = (\mathcal{M}_1, F_1)$ and a DMA $\mathcal{U} = (\mathcal{M}_2, \mathcal{F})$, where Q_i is the state set of \mathcal{M}_i for $i = 1, 2$.

Output: $u(v)^\omega \in \llbracket \mathcal{B} \rrbracket \setminus \llbracket \mathcal{U} \rrbracket$ if such exists, else “yes”.

$\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$

$\mathcal{C} \leftarrow \text{maxSCCs}(Q_1 \times Q_2)$

for each reachable $C \in \mathcal{C}$ such that $\pi_1(C) \cap F_1 \neq \emptyset$ **do**

if $\pi_2(C) \notin \mathcal{F}$ **then**

return **Witness**(C, \mathcal{M})

else

for each $F \in \mathcal{F}$ such that $F \subseteq \pi_2(C)$ and each $q \in F$ **do**

$S \leftarrow \{(q_1, q_2) \in Q_1 \times Q_2 \mid q_1 \in \pi_1(C) \wedge q_2 \in F \setminus \{q\}\}$

$\mathcal{D} \leftarrow \text{maxSCCs}(S)$

for each $D \in \mathcal{D}$ **do**

if $\pi_1(D) \cap F_1 \neq \emptyset$ and $\pi_2(D) \notin \mathcal{F}$ **then**

return **Witness**(D, \mathcal{M})

return “yes”

Theorem 14.5. *The DBAinDMA algorithm runs in polynomial time and solves the inclusion problem for an arbitrary DBA \mathcal{B} and an arbitrary DMA \mathcal{U} over the same alphabet.*

Proof. Suppose the returned value is a witness $u(v)^\omega$. Then the algorithm found an SCC E with $\pi_1(E) \cap F_1 \neq \emptyset$ and $\pi_2(E) \notin \mathcal{F}$ and returned **Witness**(E, \mathcal{M}). In this case, the returned value is correct.

Suppose for the sake of contradiction that the algorithm incorrectly returns the answer “yes”, that is, there exists an ω -word w such that $w \in \llbracket \mathcal{B} \rrbracket$ and $w \notin \llbracket \mathcal{U} \rrbracket$. Let C' denote $\text{inf}_{\mathcal{M}}(w)$. Then because $w \in \llbracket \mathcal{B} \rrbracket$, $\pi_1(C') \cap F_1 \neq \emptyset$, and because $w \notin \llbracket \mathcal{U} \rrbracket$, $\pi_2(C') \notin \mathcal{F}$.

Then C' is a subset of a unique SCC $C \in \text{maxSCCs}(Q_1 \times Q_2)$ and $\pi_1(C) \cap F_1 \neq \emptyset$. It must be that $\pi_2(C) \in \mathcal{F}$, because otherwise the algorithm would have returned **Witness**(C, \mathcal{M}). Consider the collection

$$R = \{F \in \mathcal{F} \mid \pi_2(C') \subseteq F \subseteq \pi_2(C)\},$$

of all the $F \in \mathcal{F}$ contained in $\pi_2(C)$ that contain $\pi_2(C')$. The collection R is nonempty because $C' \subseteq C$, and therefore $\pi_2(C') \subseteq \pi_2(C)$, and $\pi_2(C) \in \mathcal{F}$, so at least $\pi_2(C)$ is in R . Let F' denote a minimal element of R in the subset ordering.

Then $\pi_2(C') \subseteq F'$ but because $\pi_2(C') \notin \mathcal{F}$, it must be that $\pi_2(C') \neq F'$. Thus, there exists some $q \in F'$ that is not in $\pi_2(C')$. When the algorithm considers this F' and q , then because $\pi_2(C') \subseteq F' \setminus \{q\}$, C' is contained in R and therefore is a subset of a unique SCC D in $\text{maxSCCs}(R)$.

Because $C' \subseteq D$, and $\pi_1(C') \cap F_1 \neq \emptyset$, we have $\pi_1(D) \cap F_1 \neq \emptyset$. Also, $\pi_2(C') \subseteq \pi_2(D) \subseteq F'$, but because $q \notin \pi_2(D)$, $\pi_2(D)$ is a proper subset of F' . When the algorithm considers this D , because $\pi_1(D) \cap F_1 \neq \emptyset$, it must find that $\pi_2(D) \in \mathcal{F}$, or else it will return

Witness(D, \mathcal{M}). But then $\pi_2(D)$ is in R and is a proper subset of F' , contradicting our choice of F' as a minimal element of R . Thus, if the algorithm outputs “yes”, this is a correct answer. \square

Combining Theorem 14.4, Theorem 14.5, and the reduction of equivalence to inclusion, we have the following.

Theorem 14.6. *There are polynomial time algorithms to solve the inclusion and equivalence problems for two arbitrary DMAs over the same alphabet.*

15. COMPUTING THE AUTOMATON \mathcal{M}_{\sim_L}

In this section we use polynomial time algorithms to construct the automaton \mathcal{M}_{\sim_L} of the right congruence relation \sim_L of the language L accepted by an acceptor \mathcal{A} of one of the types DBA, DCA, DPA, DRA, DSA, or DMA. This gives a polynomial time algorithm to test whether a given DBA (resp., DCA, DPA, DRA, DSA, DMA) is of type IBA (resp., ICA, IPA, IRA, ISA, IMA).

Recall that \mathcal{A}^q is the acceptor \mathcal{A} with the initial state changed to q . If q_1 and q_2 are two states of \mathcal{A} , testing the equivalence of \mathcal{A}^{q_1} to \mathcal{A}^{q_2} determines whether these two states have the same right congruence class, and, if not, returns a witness $u(v)^\omega$ that is accepted from exactly one of the two states. The following is a consequence of Theorems 12.2, 13.2, and 14.6.

Proposition 15.1. *There is a polynomial time procedure to test whether two states of an arbitrary DBA, DCA, DPA, DRA, DSA or DMA \mathcal{A} have the same right congruence class, returning the answer “yes” if they do, and returning “no” and a witness $u(v)^\omega$ accepted from exactly one of the states if they do not.*

We now describe an algorithm **RightCon** that takes as input a DBA (or DCA, DPA, DRA, DSA, or DMA) \mathcal{A} accepting a language L and returns a deterministic automaton \mathcal{M} isomorphic to the right congruence automaton of L , i.e., \mathcal{M}_{\sim_L} .

Algorithm 3 RightCon

Input: An acceptor $\mathcal{A} = \langle \Sigma, Q, q_\iota, \delta, \alpha \rangle$ of type DBA, DCA, DPA, DRA, DSA, or DMA.

Output: A deterministic automaton \mathcal{M} isomorphic to \mathcal{M}_{\sim_L} , where $L = \llbracket \mathcal{A} \rrbracket$.

$Q' \leftarrow \{\varepsilon\}$

$q'_\iota \leftarrow \varepsilon$

δ' is initially undefined

while there exists $x \in Q'$ and $\sigma \in \Sigma$ such that $\delta'(x, \sigma)$ is undefined **do**

$q_1 \leftarrow \delta(q_\iota, x\sigma)$

if there exists $y \in Q'$ such that $\llbracket \mathcal{A}^{q_1} \rrbracket = \llbracket \mathcal{A}^{q_2} \rrbracket$ for $q_2 = \delta(q_\iota, y)$ **then**

Define $\delta'(x, \sigma) = y$

else

$Q' \leftarrow Q' \cup \{x\sigma\}$

Define $\delta'(x, \sigma) = x\sigma$

return $\mathcal{M} = \langle \Sigma, Q', q'_\iota, \delta' \rangle$

Assume the input acceptor is $\mathcal{A} = \langle \Sigma, Q, q_\iota, \delta, \alpha \rangle$. The **RightCon** algorithm constructs a deterministic automaton $\mathcal{M} = \langle \Sigma, Q', q'_\iota, \delta' \rangle$ in which the states are elements of Σ^* and $q'_\iota = \varepsilon$. The set Q' initially contains just ε , and δ' is completely undefined.

While there exists a word $x \in Q'$ and a symbol $\sigma \in \Sigma$ such that $\delta'(x, \sigma)$ has not yet been defined, loop through the words $y \in Q'$ and ask whether the states $\delta(q_i, x\sigma)$ and $\delta(q_i, y)$ have the same right congruence class in \mathcal{A} . If so, then define $\delta'(x, \sigma)$ to be y . If no such y is found, then the word $x\sigma$ is added as a new state to Q' , and the transition $\delta'(x, \sigma)$ is defined to be $x\sigma$.

This process must terminate because the elements of Q' represent distinct right congruence classes of L , and \mathcal{M}_{\sim_L} cannot have more than $|Q|$ states. When it terminates, the automaton $\mathcal{M} = \langle \Sigma, Q', q'_i, \delta' \rangle$ is isomorphic to the right congruence automaton of \mathcal{A} , \mathcal{M}_{\sim_L} .

Theorem 15.2. *The **RightCon** algorithm with input an acceptor \mathcal{A} (a DBA, DCA, DPA, DRA, DSA, or DMA) accepting L , runs in polynomial time and returns \mathcal{M} , a deterministic automaton isomorphic to \mathcal{M}_{\sim_L} ,*

To test whether a given DBA (resp., DCA, DPA, DRA, DSA, DMA) \mathcal{A} is an IBA (resp., ICA, IPA, IRA, ISA, IMA), we run the **RightCon** algorithm on \mathcal{A} and test the returned automaton \mathcal{M} for isomorphism with the automaton of \mathcal{A} . (Note that isomorphism can be checked in polynomial time by gradually constructing a map h between states of \mathcal{A} to states of \mathcal{M} . Initially $h(q_i) = q'_i$. Assume $h(q) = q'$ and $\delta(q, \sigma) = p$, $\delta(q', \sigma) = p'$. If $h(p)$ is defined and is different from p' return “non-isomorphic”, otherwise set $h(p) = p'$. This is repeated until all states are mapped and have been tested with respect to each letter of the alphabet.) If they are isomorphic, then \mathcal{A} is an IBA (resp. ICA, IPA, IRA, ISA, IMA), otherwise it is not. This proves the following.

Theorem 15.3. *There is a polynomial time algorithm to test whether a given DBA (resp., DCA, DPA, DRA, DSA, DMA) is an IBA (resp., ICA, IPA, IRA, ISA, IMA).*

16. TESTING MEMBERSHIP IN $\mathbb{I}\mathbb{X}\mathbb{A}$

In the previous section we showed that there is a polynomial time algorithm to test whether a given DBA \mathcal{B} is an IBA. However, we can also ask the following harder question. Given a DBA \mathcal{B} that is not an IBA, is $\llbracket \mathcal{B} \rrbracket \in \mathbb{I}\mathbb{B}\mathbb{A}$, that is, does there exist an IBA \mathcal{B}' such that $\llbracket \mathcal{B}' \rrbracket = \llbracket \mathcal{B} \rrbracket$? This section shows that there are such polynomial time algorithms for DBAs, DCAs, DPAs, DRAs, DSAs and DMAs. The algorithms first compute the right congruence automaton $\mathcal{M} = \mathcal{M}_{\sim_L}$, where $L = \llbracket \mathcal{A} \rrbracket$, and then attempt to construct an acceptance condition α of the appropriate type such that $\llbracket (\mathcal{M}, \alpha) \rrbracket = L$.

16.1. Testing membership in $\mathbb{I}\mathbb{B}\mathbb{A}$. We describe the algorithm **TestInIB** that takes as input a DBA \mathcal{B} and returns an IBA accepting $\llbracket \mathcal{B} \rrbracket$ if $\llbracket \mathcal{B} \rrbracket \in \mathbb{I}\mathbb{B}\mathbb{A}$, and otherwise returns “no”. By Claim 2.2 (4), the case of a DCA is reduced to that of a DBA.

Theorem 16.1. *The algorithm **TestInIB** takes a DBA \mathcal{B} as input, runs in polynomial time, and returns an IBA accepting $\llbracket \mathcal{B} \rrbracket$ if $\llbracket \mathcal{B} \rrbracket \in \mathbb{I}\mathbb{B}\mathbb{A}$, and otherwise returns “no”.*

Proof. The algorithm calls the **RightCon** algorithm, and also the inclusion and equivalence algorithms from Theorem 12.3, which run in polynomial time in the size of \mathcal{B} . If the algorithm returns an acceptor, it is an IBA accepting $\llbracket \mathcal{B} \rrbracket$.

To see that the algorithm does not incorrectly return the answer “no”, suppose \mathcal{B}' is an IBA accepting $\llbracket \mathcal{B} \rrbracket$. Then because \mathcal{M} is isomorphic to \mathcal{M}_{\sim_L} , we may assume that $\mathcal{B}' = (\mathcal{M}, F')$. For every state $q \in F'$, the inclusion query with $(\mathcal{M}, \{q\})$ will answer “yes”,

Algorithm 4 TestInIB

Input: A DBA \mathcal{B} .

Output: If $\llbracket \mathcal{B} \rrbracket \in \mathbb{IBA}$ then return an IBA accepting $\llbracket \mathcal{B} \rrbracket$, else return “no”.

$\mathcal{M} \leftarrow \mathbf{RightCon}(\mathcal{B})$

$F \leftarrow \emptyset$

for each state q of \mathcal{M} **do**

if $\llbracket (\mathcal{M}, \{q\}) \rrbracket \subseteq \llbracket \mathcal{B} \rrbracket$ **then**

$F \leftarrow F \cup \{q\}$

if $\llbracket (\mathcal{M}, F) \rrbracket = \llbracket \mathcal{B} \rrbracket$ **then**

return (\mathcal{M}, F)

else

return “no”

so q will be added to F . Thus, $F' \subseteq F$, and $\llbracket (\mathcal{M}, F) \rrbracket$ subsumes $\llbracket (\mathcal{M}, F') \rrbracket$. Every state q added to F preserves the condition that $\llbracket (\mathcal{M}, F) \rrbracket$ is a subset of $\llbracket \mathcal{B} \rrbracket$, so the final equivalence check will pass, and (\mathcal{M}, F) will be returned. \square

16.2. Testing membership in \mathbb{IPA} . We describe the algorithm **TestInIP** that takes as input a DPA \mathcal{P} and returns an IPA accepting $\llbracket \mathcal{P} \rrbracket$ if $\llbracket \mathcal{P} \rrbracket \in \mathbb{IPA}$, and otherwise returns “no”.

Algorithm 5 TestInIP

Input: A DPA \mathcal{P} .

Output: If $\llbracket \mathcal{P} \rrbracket \in \mathbb{IPA}$ then return an IPA accepting $\llbracket \mathcal{P} \rrbracket$, else return “no”.

$\mathcal{M} = \langle \Sigma, Q, q_i, \delta \rangle \leftarrow \mathbf{RightCon}(\mathcal{P})$

Define $\kappa(q) = 0$ for all states $q \in Q$

for $k = 1$ to $|Q|$ **do**

if $\llbracket (\mathcal{M}, \kappa) \rrbracket = \llbracket \mathcal{P} \rrbracket$ **then**

return (\mathcal{M}, κ)

else if k is odd **then**

while $\llbracket \mathcal{P} \rrbracket$ is not a subset of $\llbracket (\mathcal{M}, \kappa) \rrbracket$ **do**

 Let $u(v)^\omega$ be the returned witness

 Define $\kappa(q) = k$ for all $q \in \text{inf}_{\mathcal{M}}(u(v)^\omega)$

else

while $\llbracket \mathcal{P} \rrbracket$ is not a superset of $\llbracket (\mathcal{M}, \kappa) \rrbracket$ **do**

 Let $u(v)^\omega$ be the returned witness

 Define $\kappa(q) = k$ for all $q \in \text{inf}_{\mathcal{M}}(u(v)^\omega)$

return “no”

Theorem 16.2. *The algorithm **TestInIP** takes a DPA \mathcal{P} as input, runs in polynomial time, and returns an IPA accepting $\llbracket \mathcal{P} \rrbracket$ if $\llbracket \mathcal{P} \rrbracket \in \mathbb{IPA}$, and otherwise returns “no”.*

Proof. The algorithm calls the **RightCon** algorithm and the inclusion and equivalence algorithms for DPAs from Theorem 12.2, which run in polynomial time in the size of \mathcal{P} . Below we show that each **while** loop terminates after at most $|Q|$ iterations. If the algorithm returns an acceptor, then the acceptor is an IPA accepting $\llbracket \mathcal{P} \rrbracket$.

To see that the algorithm does not incorrectly return the answer “no”, suppose \mathcal{P}' is an IPA accepting $\llbracket \mathcal{P} \rrbracket$. We may assume that $\mathcal{P}' = (\mathcal{M}, \kappa^*)$, where κ^* is the canonical

coloring of \mathcal{P}' . We prove inductively that the final coloring κ is equal to κ^* . To do so, we consider the conditions after the **for** loop has been completed ℓ times: (1) if ℓ is even then $\llbracket(\mathcal{M}, \kappa)\rrbracket \subseteq \llbracket\mathcal{P}\rrbracket$, and if ℓ is odd, then $\llbracket\mathcal{P}\rrbracket \subseteq \llbracket(\mathcal{M}, \kappa)\rrbracket$, and (2) for all $q \in Q$, if $\kappa^*(q) \leq \ell$ then $\kappa(q) = \kappa^*(q)$, and if $\kappa^*(q) > \ell$ then $\kappa(q) = \ell$.

The initialization of $\kappa(q) = 0$ for all $q \in Q$ implies that these two conditions hold for $\ell = 0$. Suppose the conditions hold for some $\ell \geq 0$. If the equivalence check at the start of the next iteration returns “yes” then the correct IPA (\mathcal{M}, κ) is returned. Otherwise, $k = \ell + 1$; we consider the cases of odd and even k .

If k is odd, then by condition (1), $\llbracket(\mathcal{M}, \kappa)\rrbracket \subsetneq \llbracket\mathcal{P}\rrbracket$ and at least one witness $u(v)^\omega$ accepted by \mathcal{P} and rejected by $\llbracket(\mathcal{M}, \kappa)\rrbracket$ will be processed in the **while** loop. Consider such a witness $u(v)^\omega$ and let $C = \text{inf}_{\mathcal{M}}(u(v)^\omega)$. Then because $\kappa(q) = \kappa^*(q)$ if $\kappa(q) \leq \ell$, it must be that $\kappa^*(C) > \ell$ and $\kappa(C) = \ell$. For all $q \in C$, $\kappa(q)$ is set to $k = \ell + 1$, so at least one state changes κ -color from ℓ to $\ell + 1$. This can happen at most $|Q|$ times, so the **while** loop for this k must terminate after at most $|Q|$ iterations. No state q with $\kappa^*(q) \leq \ell$ has its κ -value changed, so when the **while** loop is terminated, we have that $\kappa^*(q) \leq \ell$ implies $\kappa(q) = \kappa^*(q)$.

Consider any state q with $\kappa^*(q) \geq \ell + 1$. By property (2), at the start of this iteration of the **for** loop, $\kappa(q) = \ell$. Referring to the canonical forest \mathcal{F}^* for \mathcal{P}' , the state q is in $\Delta(D)$ for some node D of \mathcal{F}^* . The node D is a descendant (or possibly equal to) some node C for which the states $q \in \Delta(C)$ all have $\kappa^*(q) = \ell + 1$. Thus, as long as the value of $\kappa(q)$ remains ℓ , the SCC C will have $\kappa(C)$ even and $\kappa^*(C)$ odd, and the **while** loop cannot terminate. But we have shown that it does terminate, so after termination we must have $\kappa(q) = k = \ell + 1$. Thus, after this iteration of the **for** loop, property (2) holds for $\ell + 1$.

The case of even k is dual to the case of odd k . Because the range of κ^* is $[0..j]$ for some $j \leq |Q|$, the equivalence test must return “yes” before the **for** loop completes, at which point the IPA (\mathcal{M}, κ) is returned. \square

16.3. Testing membership in \mathbb{IRA} . We describe the algorithm **TestInIR** that takes as input a DRA \mathcal{R} and returns an IRA accepting $\llbracket\mathcal{R}\rrbracket$ if $\llbracket\mathcal{R}\rrbracket \in \mathbb{IRA}$, and otherwise returns “no”. By Claim 2.2 (5), the case of a DSA is reduced to that of a DRA.

We first show that given a DRA \mathcal{R} such that $\llbracket\mathcal{R}\rrbracket \in \mathbb{IRA}$, there is an IRA equivalent to \mathcal{R} whose size is bounded by a polynomial in the size of \mathcal{R} .

Lemma 16.3. *Let \mathcal{R} be a DRA in singleton normal form whose acceptance condition has m pairs, and assume $\llbracket\mathcal{R}\rrbracket \in \mathbb{IRA}$. Let \mathcal{M} be the right congruence automaton of $\llbracket\mathcal{R}\rrbracket$ with state set Q and assume $|Q| = n$. Then there exists an acceptance condition α in singleton normal form with at most mn pairs such that (\mathcal{M}, α) accepts $\llbracket\mathcal{R}\rrbracket$.*

Proof. Let $\mathcal{R} = (\mathcal{M}_1, \alpha_1)$, where all the states of \mathcal{M}_1 are reachable, and let the function f map each state of \mathcal{M}_1 to the state of its right congruence class in \mathcal{M} . It suffices to show that for each $(q, B) \in \alpha_1$ there exists an acceptance condition α' of \mathcal{M} containing at most n pairs such that $\llbracket(\mathcal{M}_1, \{(q, B)\})\rrbracket \subseteq \llbracket(\mathcal{M}, \alpha')\rrbracket \subseteq \llbracket\mathcal{R}\rrbracket$. Taking the union of these α' conditions for all m pairs $(q, B) \in \alpha_1$ yields the desired acceptance condition α for \mathcal{M} .

Because we assume $\llbracket\mathcal{R}\rrbracket \in \mathbb{IRA}$, there exists an IRA (\mathcal{M}, α_2) in singleton normal form that accepts $\llbracket\mathcal{R}\rrbracket$. Given any $u(v)^\omega$ in $\llbracket\mathcal{R}\rrbracket$, let $C = \text{inf}_{\mathcal{M}_1}(u(v)^\omega)$. Then $f(C) = \text{inf}_{\mathcal{M}}(u(v)^\omega)$ and there exists $(q', B') \in \alpha_2$ such that $q' \in f(C)$, and $f(C) \cap B' = \emptyset$. Then also $\llbracket(\mathcal{M}, \{(q', Q \setminus f(C))\})\rrbracket \subseteq \llbracket\mathcal{R}\rrbracket$. To see this, consider any $u'(v')^\omega$ with $D = \text{inf}_{\mathcal{M}}(u'(v')^\omega)$

and $q' \in D$ and $D \cap (Q \setminus f(C)) = \emptyset$. Then D is a subset of $f(C)$, $D \cap B' = \emptyset$, $u'(v')^\omega$ satisfies (q', B') , and $u'(v')^\omega \in \llbracket \mathcal{R} \rrbracket$.

Given a pair $(q, B) \in \alpha_1$ the construction of the initially empty acceptance condition α' proceeds as follows. Let C_0 be the maximum SCC of \mathcal{M}_1 that contains q and contains no element of B . If C_0 is empty, then $(\mathcal{M}_1, \{(q, B)\})$ does not accept any words, and the empty condition α' suffices. If C_0 is nonempty, then there is an element $u(v)^\omega$ of $\llbracket (\mathcal{M}_1, \{(q, B)\}) \rrbracket$ such that $C_0 = \text{inf}_{\mathcal{M}_1}(u(v)^\omega)$ and there is a pair (q_0, B_0) in α_2 such that $q_0 \in f(C_0)$ and $B_0 \cap f(C_0) = \emptyset$. We add the pair $(q_0, Q \setminus f(C_0))$ to α' and note that by the argument in the preceding paragraph, $\llbracket (\mathcal{M}, \alpha') \rrbracket \subseteq \llbracket \mathcal{R} \rrbracket$.

If $\llbracket (\mathcal{M}_1, \{(q, B)\}) \rrbracket \subseteq \llbracket (\mathcal{M}, \alpha') \rrbracket$ then α' is the desired acceptance condition. If not, there exists a word $u(v)^\omega$ such that for $C = \text{inf}_{\mathcal{M}_1}(u(v)^\omega)$ we have $q \in C$ and $C \cap B = \emptyset$, but either $q_0 \notin f(C)$ or $f(C) \cap (Q \setminus f(C_0)) \neq \emptyset$. Because C_0 is the maximum SCC of \mathcal{M}_1 containing q and containing no element of B , we have $C \subseteq C_0$, so $f(C) \subseteq f(C_0)$ and therefore $q_0 \notin f(C)$. Let C_1 be the maximum SCC C' of \mathcal{M}_1 such that $C' \subseteq C_0$, $q \in C'$, and $q_0 \notin f(C')$. This is not empty, so there is a word $u'(v')^\omega$ such that $C_1 = \text{inf}_{\mathcal{M}_1}(u'(v')^\omega)$, which is in $\llbracket \mathcal{R} \rrbracket$ because it satisfies (q, B) . Thus there exists a pair (q_1, B_1) in α_2 that is satisfied by $u'(v')^\omega$, and we add the pair $(q_1, Q \setminus f(C_1))$ to the acceptance condition α' . As above, we have $\llbracket (\mathcal{M}, \alpha') \rrbracket \subseteq \llbracket \mathcal{R} \rrbracket$.

If now $\llbracket (\mathcal{M}_1, \{(q, B)\}) \rrbracket \subseteq \llbracket (\mathcal{M}, \alpha') \rrbracket$, then α' is the desired acceptance condition. If not, we repeat this step again. In general, after k steps of this kind, α' consists of k pairs of the form $(q_i, Q \setminus f(C_i))$ for $i \in [0..k-1]$, where all of the states q_i are distinct and $C_{i+1} \subseteq C_i$ for $i \in [0, k-2]$. Because Q has n states, there can only be n repetitions of this step before α' satisfies the required condition, and thus α' has at most n pairs. \square

The algorithm **TestInIR** is based on the algorithm to learn Horn sentences by Angluin, Frazier, and Pitt [AFP92], using the analogy between singleton normal form for Rabin automata and propositional Horn clauses. **TestInIR** maintains for each state q of the right congruence automaton \mathcal{M} an ordered sequence S_q of SCCs of \mathcal{M} , each of which corresponds to a positive example of $\llbracket \mathcal{R} \rrbracket$. At each iteration, the algorithm uses these sequences and inclusion queries with $\llbracket \mathcal{R} \rrbracket$ to construct an acceptance condition α for a hypothesis (\mathcal{M}, α) , which it tests for equivalence to \mathcal{R} . In the case of non-equivalence, the witness is a positive example of $\llbracket \mathcal{R} \rrbracket$ that is used to update the sequences S_q .

In **TestInIR** the test of whether $C \cup C_i$ is positive is implemented by calling **Witness** $(C \cup C_i)$ and testing the resulting word $u(v)^\omega$ for membership in $\llbracket \mathcal{R} \rrbracket$.

Because pairs are only added to α that preserve inclusion in $\llbracket \mathcal{R} \rrbracket$, it is clear that any witness $u(v)^\omega$ returned in response to the test of equivalence of (\mathcal{M}, α) and \mathcal{R} is a positive example of $\llbracket \mathcal{R} \rrbracket$. Note also that all elements of S_q are SCCs of \mathcal{M} that contain q . The proof of correctness and running time of **TestInIR** depends on the following two lemmas.

Lemma 16.4. *Assume that \mathcal{R}' is an IRA equivalent to the target DRA \mathcal{R} . Consider a positive example $u(v)^\omega$ of $\llbracket \mathcal{R} \rrbracket$ returned in response to the test of equivalence of (\mathcal{M}, α) and \mathcal{R} , and let $C = \text{inf}_{\mathcal{M}}(u(v)^\omega)$. Let (q', B) be a pair of \mathcal{R}' such that $q' \in C$ and $C \cap B = \emptyset$. If for some $q \in C$ and some C_i in S_q we have $C_i \cap B = \emptyset$ then for some $j \leq i$, the element C_j of S_q will be replaced by $C_j \cup C$.*

Proof. Assume that there is no such replacement for $j < i$. When i is considered, we have $q \in C_i$ and $q \in C$, so $C_i \cup C$ is the union of overlapping SCCs and therefore an SCC. Then $C_i \cup C$ is positive because $q' \in C$, so $q' \in C_i \cup C$, and $C \cap B = \emptyset$ and $C_i \cap B = \emptyset$ by hypothesis, so $(C_i \cup C) \cap B = \emptyset$ and the word **Witness** $(C_i \cup C)$ is a positive example of $\llbracket \mathcal{R} \rrbracket$ because it satisfies (q', B) .

Algorithm 6 TestInIR

Input: A DRA $\mathcal{R} = (\mathcal{M}_1, \alpha_1)$ in singleton normal form with $|\alpha_1| = m$.

Output: If $\llbracket \mathcal{R} \rrbracket \in \mathbb{IRA}$ then return an IRA accepting $\llbracket \mathcal{R} \rrbracket$, else return “no”.

$\mathcal{M} \leftarrow \mathbf{RightCon}(\mathcal{R})$

Let Q be the states of \mathcal{M} and $n = |Q|$

For each $q \in Q$ initialize a sequence S_q to be empty

for $k = 1$ to mn^3 **do**

for all $q \in Q$ **do**

$\alpha_q = \emptyset$

for all $C \in S_q$ **do**

for all $q' \in C$ **do**

if $\llbracket (\mathcal{M}, \{(q', Q \setminus C)\}) \rrbracket \subseteq \llbracket \mathcal{R} \rrbracket$ **then**

$\alpha_q = \alpha_q \cup \{(q', Q \setminus C)\}$

$\alpha \leftarrow \bigcup_{q \in Q} \alpha_q$

if $\llbracket (\mathcal{M}, \alpha) \rrbracket = \llbracket \mathcal{R} \rrbracket$ **then**

return (\mathcal{M}, α)

else

 Let $u(v)^\omega$ be the witness returned

 Let $C = \mathit{inf}_{\mathcal{M}}(u(v)^\omega)$

for all $q \in C$ **do**

if there is some $C_i \in S_q$ such that $C \not\subseteq C_i$ and $C \cup C_i$ is positive **then**

 Let i be the least such i and replace C_i by $C_i \cup C$

else

 Add C to the end of the sequence S_q

return “no”

To see that $C \not\subseteq C_i$, we assume to the contrary. Then $(q', Q \setminus C_i)$ is an element of α . To see this, we show that $\llbracket (\mathcal{M}, \{(q', Q \setminus C_i)\}) \rrbracket \subseteq \llbracket \mathcal{R} \rrbracket$. Let $u'(v')^\omega$ with $D = \mathit{inf}_{\mathcal{M}}(u'(v')^\omega)$ satisfy $(q', Q \setminus C_i)$. Then $q' \in D$ and $D \cap (Q \setminus C_i) = \emptyset$, and therefore $D \subseteq C_i$ and $C_i \cap B = \emptyset$ by hypothesis. Thus $u'(v')^\omega$ satisfies (q', B) and is in $\llbracket \mathcal{R} \rrbracket$. Because $(q', Q \setminus C_i)$ is an element of α , $u(v)^\omega$ is accepted by (\mathcal{M}, α) because $q' \in C$ and $C \cap (Q \setminus C_i) = \emptyset$ (because we assume $C \subseteq C_i$). But this means that $u(v)^\omega$ cannot be a witness to the non-equivalence of (\mathcal{M}, α) and \mathcal{R} , a contradiction.

Thus, the conditions for C_i to be replaced by $C_i \cup C$ are satisfied. \square

Lemma 16.5. *Assume that $\mathcal{R}' = (\mathcal{M}, \alpha')$ is an IRA in singleton normal form equivalent to the target DRA \mathcal{R} . The following two conditions hold throughout the algorithm **TestInIR**.*

- (1) *For all $q \in Q$, elements C_i of S_q , and $(q', B) \in \alpha'$, if C_i satisfies (q', B) then for no $j < i$ do we have $C_j \cap B = \emptyset$.*
- (2) *For all $q \in Q$, elements C_i and C_j of S_q with $j < i$, and $(q', B) \in \alpha'$, if C_i satisfies (q', B) then C_j does not satisfy (q', B) .*

Proof. We first show that condition (1) implies condition (2). Let $q \in Q$, C_i and C_j be in S_q with $j < i$ and $(q', B) \in \alpha'$. If C_i satisfies (q', B) then by condition (1), we have $C_j \cap B \neq \emptyset$, so C_j does not satisfy (q', B) .

We now prove condition (1) by induction on the number of witnesses to non-equivalence. The condition holds of the empty sequences S_q . Suppose conditions (1) and (2) hold of the sequences S_q , and the witness to non-equivalence is $u(v)^\omega$ with $C = \mathit{inf}_{\mathcal{M}}(u(v)^\omega)$. If $q \notin C$

then S_q is not modified, so assume $q \in C$. We consider two cases, depending on whether C is added to the end of S_q or causes some C_ℓ to be replaced by $C_\ell \cup C$.

Assume that C is added to the end of S_q and property (1) fails to hold. Then it must be that for $C_i = C$, some C_j in S_q with $j < i$ and some $(q', B) \in \alpha'$, C_i satisfies (q', B) and $C_j \cap B = \emptyset$. By Lemma 16.4, because $C_j \cap B = \emptyset$, C should cause C_ℓ for some $\ell \leq j$ to be replaced by $C \cup C_\ell$ rather than being added to the end of S_q , a contradiction.

Assume that C causes C_ℓ in S_q to be replaced by $C_\ell \cup C$ and property (1) fails to hold. Then it must be for some C_i in S_q and pair (q', B) in α' , either (i) $i > \ell$ and C_i satisfies (q', B) and $(C_\ell \cup C) \cap B = \emptyset$, or (ii) $i < \ell$ and $C_\ell \cup C$ satisfies (q', B) and $C_i \cap B = \emptyset$. In case (i), it must be that C_i satisfies (q', B) and $C_\ell \cap B = \emptyset$, which contradicts the assumption that property (1) holds before C is processed, because $\ell < i$. In case (ii), $q' \in C_\ell \cup C$ and $(C_\ell \cup C) \cap B = \emptyset$. Thus $C \cap B = \emptyset$ and $C_\ell \cap B = \emptyset$. If $q' \in C_\ell$, then C_ℓ satisfies (q', B) , violating the assumption that property (2) holds before C is processed. If $q' \in C$ then C satisfies (q', B) and because $C_i \cap B = \emptyset$, by Lemma 16.4, for some $j \leq i$ we have C_j replaced by $C_j \cup C$, a contradiction because $\ell > i$. Thus, in either case property (1) holds after C_ℓ is replaced by $C_\ell \cup C$. \square

Theorem 16.6. *The algorithm **TestInIR** takes as input a DRA \mathcal{R} , runs in polynomial time, and returns an IRA accepting $\llbracket \mathcal{R} \rrbracket$ if $\llbracket \mathcal{R} \rrbracket \in \text{IRA}$, and otherwise returns “no”.*

Proof. Assume the input is a DRA $\mathcal{R} = (\mathcal{M}_1, \alpha_1)$ in singleton normal form with $|\alpha_1| = m$. The algorithm **TestInIR** computes the right congruence automaton \mathcal{M} of $\llbracket \mathcal{R} \rrbracket$, which has n states, at most the number of states of \mathcal{R} . The main loop of the algorithm is executed at most mn^3 times, and each execution makes calls to the inclusion and equivalence algorithms for DRAs, and runs in time polynomial in the size of \mathcal{R} , so the overall running time of **TestInIR** is polynomial in the size of \mathcal{R} .

Clearly, if $\llbracket \mathcal{R} \rrbracket \notin \text{IRA}$, then the test of equivalence between (\mathcal{M}, α) and \mathcal{R} will not succeed, and the value returned will be “no”. Assume that $\llbracket \mathcal{R} \rrbracket \in \text{IRA}$. Then by Lemma 16.3, there is an IRA $\mathcal{R}' = (\mathcal{M}, \alpha')$ in singleton normal form equivalent to \mathcal{R} such that $|\alpha'| \leq mn$. For every state $q \in Q$, each member of the sequence S_q satisfies some pair (q', B) in α' , and by Lemma 16.5, no two members of S_q can satisfy the same pair, so the length of each S_q is bounded by mn . Each positive counterexample must either add another member to at least one sequence S_q or cause at least one member of some sequence S_q to increase in cardinality by 1. The maximum cardinality of any member of any S_q is n , and the total number of sequences S_q is n , so no more than mn^3 positive counterexamples can be processed before the test of equivalence between (\mathcal{M}, α) and \mathcal{R} succeeds and (\mathcal{M}, α) is returned. \square

16.4. Testing membership in IMA. We describe the algorithm **TestInIM** that takes as input a DMA \mathcal{U} and returns an IMA accepting $\llbracket \mathcal{U} \rrbracket$ if $\llbracket \mathcal{U} \rrbracket \in \text{IMA}$, and otherwise returns “no”.

Theorem 16.7. *The algorithm **TestInIM** takes as input a DMA \mathcal{U} , runs in polynomial time, and returns an IMA accepting $\llbracket \mathcal{U} \rrbracket$ if $\llbracket \mathcal{U} \rrbracket \in \text{IMA}$, and otherwise returns “no”.*

Proof. The algorithm calls the **RightCon** algorithm and also the DMA equivalence algorithm from Theorem 14.6, which run in polynomial time. When there is a witness $u(v)^\omega$ accepted by \mathcal{U} , there is a set $F \in \mathcal{F}$ whose image in \mathcal{M} is added to \mathcal{F}' , so there can be no more such witnesses than the number of sets in \mathcal{F} . After this, there must be a successful equivalence

Algorithm 7 TestInIM

Input: A DMA $\mathcal{U} = \langle \Sigma, Q, q_i, \delta, \mathcal{F} \rangle$.**Output:** If $\llbracket \mathcal{U} \rrbracket \in \mathbb{IMA}$ then return an IMA accepting $\llbracket \mathcal{U} \rrbracket$, else return “no”. $\mathcal{M} \leftarrow \mathbf{RightCon}(\mathcal{U})$ $\mathcal{F}' \leftarrow \emptyset$ **while** $\llbracket (\mathcal{M}, \mathcal{F}') \rrbracket \neq \llbracket \mathcal{U} \rrbracket$ **do** Let $u(v)^\omega$ be the witness returned Let $C = \mathit{inf}_{\mathcal{M}}(u(v)^\omega)$ **if** $u(v)^\omega \in \llbracket \mathcal{U} \rrbracket$ **then** $\mathcal{F}' \leftarrow \mathcal{F}' \cup \{C\}$ **else** **return** “no”**return** $(\mathcal{M}, \mathcal{F}')$

test or a witness rejected by \mathcal{U} , either of which terminates the while loop. Thus, the overall running time is polynomial in the size of \mathcal{U} .

If the algorithm returns an acceptor $(\mathcal{M}, \mathcal{F}')$, then the acceptor is an IBA that accepts $\llbracket \mathcal{U} \rrbracket$. To see that the algorithm does not incorrectly return the answer “no”, assume that \mathcal{U}' is an IBA accepting $\llbracket \mathcal{U} \rrbracket$. We may assume that $\mathcal{U}' = (\mathcal{M}, \mathcal{F}'')$, where \mathcal{F}'' contains no redundant sets. Then the first witness will be a word accepted by \mathcal{U} that will add an element of \mathcal{F}'' to \mathcal{F}' . This continues until all the elements of \mathcal{F}'' have been added to \mathcal{F}' , at which point the while loop terminates with equivalence. \square

16.5. Variants of the testing algorithms. A variant of the task considered above is the following. Given the right congruence automaton \mathcal{M} of a language $\llbracket \mathcal{B} \rrbracket$ in \mathbb{IBA} , and access to information from certain queries about $\llbracket \mathcal{B} \rrbracket$, learn an acceptance condition α such that (\mathcal{M}, α) accepts $\llbracket \mathcal{B} \rrbracket$. In the case of \mathbb{IBA} , the algorithm **TestInIB** could be modified to perform this task using just equivalence queries with respect to $\llbracket \mathcal{B} \rrbracket$. Similarly, equivalence queries would suffice in the case of \mathbb{IMA} . For \mathbb{IPA} , subset and superset queries with respect to the target language would suffice. And for \mathbb{IRA} , **TestInIR** could be modified to use membership and equivalence queries with respect to the target language, relying on negative examples to remove incorrect pairs rather than using subset queries.

16.6. Efficient teachability of the informative classes. We can finally claim that the the informative classes are efficiently teachable.

Theorem 16.8. *The classes \mathbb{IBA} , \mathbb{ICA} , \mathbb{IPA} , \mathbb{IMA} , \mathbb{IRA} and \mathbb{ISA} are efficiently teachable.*

Proof. By Theorems 7.2, 7.4, 8.8, and 9.5 the classes \mathbb{IMA} , \mathbb{IBA} , \mathbb{ICA} , \mathbb{IPA} , \mathbb{IRA} and \mathbb{ISA} are identifiable in the limit using polynomial time and data. It remains to show that the characteristic samples can be computed in polynomial time for any acceptor in the class.

Let \mathcal{A} be an acceptor of type DXA for $X \in \{B, C, P, R, S, M\}$. By Theorems 16.1, 16.2, 16.6, 16.7, there are polynomial time algorithms that return an equivalent acceptor \mathcal{A}' in IXA if such exists and “no” otherwise. By Theorem 10.6, given an acceptor \mathcal{A}' of type IBA, ICA, IPA, IRA, ISA, or IMA, the characteristic sample T_L for \mathcal{A} may be computed in polynomial time in the size of \mathcal{A} . It follows that the classes \mathbb{IMA} , \mathbb{IBA} , \mathbb{ICA} , \mathbb{IPA} , \mathbb{IRA} and \mathbb{ISA} are efficiently teachable. \square

17. DISCUSSION

We have provided general definitions and comparisons for characteristic samples, efficient learnability, efficient teachability, efficient teachability/learnability, and identifiability in the limit using polynomial time and data. While efficient teachability implies identifiability in the limit using polynomial time and data, we have shown that the converse is not true if there is no polynomial time algorithm for integer factorization. We have shown that if a class is efficiently teachable, then it is it efficiently teachable/learnable.

We then asked which classes of representations of regular ω -languages are efficiently teachable. The non-deterministic acceptors NBA, NCA, NPA, NRA, NSA, and NMA do not have polynomial size characteristic sets, and thus are neither efficiently identifiable in the limit with polynomial time and data nor efficiently teachable. We have shown that the classes M2MA and SUBA are efficiently teachable.

Focusing on the classes of informative languages, IBA, ICA, IPA, IRA, ISA and IMA, we have shown that they are efficiently teachable. To obtain these results we have given new polynomial time algorithms to test inclusion and equivalence for DBAs, DCAs, DPAs, DRAs, DSAs and DMAs. We have given a polynomial time algorithm to compute the right congruence automaton \mathcal{M}_{\sim_L} for a language L specified by a DBA, DCA, DPA, DRA, DSA, or DMA. This yields a polynomial time algorithm to test whether an acceptor \mathcal{A} of type DBA is of type IBA, and similarly for acceptors of types DCA, DPA, DRA, DSA and DMA. Moreover, we have given a polynomial time algorithm to test whether an acceptor \mathcal{A} of type DBA accepts a language in the class IBA, and similarly for acceptors of types DCA, DPA, DRA, DSA and DMA.

The questions of whether the full deterministic classes DBA, DCA, DPA, DRA, DSA and DMA are efficiently teachable or identifiable in the limit using polynomial time and data remain open. We note that Bohn and Löding [BL22, BL23] have recently obtained interesting results for passive learning of DBA and DPA, but from characteristic samples of cardinality that may be exponential in the size of the minimal representations in the worst case. Thus these results do not settle the question of whether these classes are efficiently teachable or identifiable in the limit using polynomial time and data. Another intriguing open question is whether the classes IBA, ICA, IPA, IRA, ISA and IMA can be learned by polynomial time algorithms using membership and equivalence queries. However, as shown by Bohn and Löding [BL21], this question is not easier than whether the corresponding deterministic classes can be learned by polynomial time algorithms using membership and equivalence queries.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for very constructive comments that greatly improved the quality of this paper. This work was partially supported by ISF grant 2507/21.

REFERENCES

- [AAF20] D. Angluin, T. Antonopoulos, and D. Fisman. Strongly unambiguous Büchi automata are polynomially predictable with membership queries. In *28th EACSL Annual Conference on Computer Science Logic, CSL*, pages 8:1–8:17, 2020.

- [AAFG22] D. Angluin, T. Antonopoulos, D. Fisman, and N. George. Representing regular languages of infinite words using mod 2 multiplicity automata. In *Foundations of Software Science and Computation Structures - 25th Inter. Conference, FOSSACS*, pages 1–20, 2022.
- [ABF16] D. Angluin, U. Boker, and D. Fisman. Families of DFAs as acceptors of omega-regular languages. In *41st Inter. Symp. on Mathematical Foundations of Computer Science, MFCS*, pages 11:1–11:14, 2016.
- [ABL02] G. Ammons, R. Bodík, and J. R. Larus. Mining specifications. In *The 29th SIGPLAN-SIGACT Symp. on Principles of Programming Languages, POPL.*, pages 4–16, 2002.
- [AF14] D. Angluin and D. Fisman. Learning regular omega languages. In *Algorithmic Learning Theory - 25th Inter. Conference, ALT 2014*, pages 125–139, 2014.
- [AF16] D. Angluin and D. Fisman. Learning regular omega languages. *Theor. Comput. Sci.*, 650:57–72, 2016.
- [AF18] D. Angluin and D. Fisman. Regular omega-languages with an informative right congruence. In *9th Inter. Symp. on Games, Automata, Logics, and Formal Verification, GandALF*, volume 277 of *EPTCS*, pages 265–279, 2018.
- [AFP92] D. Angluin, M. Frazier, and L. Pitt. Learning conjunctions of Horn clauses. *Machine Learning*, 9:147–164, 1992.
- [AFS20] D. Angluin, D. Fisman, and Y. Shoval. Polynomial identification of omega-automata. In *Proc. 26th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, pages 325–343, 2020.
- [AKS04] M. Agrawal, N. Kayal, and N. Saxena. Primes is in P. *Annals of Mathematics*, 160:781–793, June 2004. URL: <https://www.microsoft.com/en-us/research/publication/primes-is-in-p/>.
- [Ang87] D. Angluin. Learning regular sets from queries and counterexamples. *Inf. Comput.*, 75(2):87–106, 1987.
- [AV10] F. Aarts and F. Vaandrager. Learning I/O automata. In *21th Inter. Conference, CONCUR 2010*, pages 71–85, 2010.
- [BBB⁺00] A. Beimel, F. Bergadano, N. H. Bshouty, E. Kushilevitz, and S. Varricchio. Learning functions represented as multiplicity automata. *J. ACM*, 47(3):506–530, May 2000.
- [BF72] J.M. Barzdin and R. V. Freivalds. On the prediction of total recursive functions. *Doklady Akademii Nauk SSSR*, 206(3):521–524, 1972. (In Russian).
- [BL21] L. Bohn and C. Löding. Constructing deterministic ω -automata from examples by an extension of the RPNI algorithm. In *46th Int. Symp. on Mathematical Foundations of Computer Science*, pages 20:1–20:18, 2021.
- [BL22] L. Bohn and C. Löding. Passive learning of deterministic Büchi automata by combinations of DFAs. In *49th Inter. Colloquium on Automata, Languages, and Programming, ICALP 2022*, pages 114:1–114:20, 2022. doi:10.4230/LIPIcs.ICALP.2022.114.
- [BL23] L. Bohn and C. Löding. Constructing deterministic parity automata from positive and negative examples. *CoRR*, abs/2302.11043, 2023. arXiv:2302.11043, doi:10.48550/arXiv.2302.11043.
- [Bok18] U. Boker. Why these automata types? In *LPAR-22. 22nd Inter. Conference on Logic for Programming, Artificial Intelligence and Reasoning*, pages 143–163, 2018.
- [Büc62] J.R. Büchi. On a decision method in restricted second order arithmetic. In *Inter. Congress on Logic, Methodology and Philosophy*, pages 1–11, 1962.
- [CBSS10] C. Y. Cho, D. Babic, E. C.R. Shin, and D. Song. Inference and analysis of formal models of botnet command and control protocols. In *Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS*, pages 426–439, 2010.
- [CCK⁺15] M. Chapman, H. Chockler, P. Kesseli, D. Kroening, O. Strichman, and M. Tautschnig. Learning the language of error. In *Automated Technology for Verification and Analysis - 13th Inter. Symp., ATVA*, pages 114–130, 2015.
- [CDK93] E.M. Clarke, I. A. Draghicescu, and R. P. Kurshan. A unified approach for showing language inclusion and equivalence between various types of omega-automata. *Inf. Process. Lett.*, 46(6):301–308, 1993.
- [CGP03] J.M. Cobleigh, D. Giannakopoulou, and C.S. Păsăreanu. Learning assumptions for compositional verification. In *Proceedings of the 9th Inter. Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 331–346, 2003.

- [Cho74] Y. Choueka. Theories of automata on omega-tapes: A simplified approach. *J. Comput. Syst. Sci.*, 8(2):117–141, 1974.
- [CM99] O. Carton and R. Maceiras. Computing the Rabin index of a parity automaton. *RAIRO Theor. Informatics Appl.*, 33(6):495–506, 1999.
- [CNP93] H. Calbrix, M. Nivat, and A. Podelski. Ultimately periodic words of rational w -languages. In *9th Inter. Conf. on Mathematical Foundations of Programming Semantics (MFPS)*, pages 554–566, 1993.
- [CPPdR14] G. Chalupar, S. Peherstorfer, E. Poll, and J. de Ruitter. Automated reverse engineering using Lego®. In *8th USENIX Workshop on Offensive Technologies (WOOT 14)*, August 2014.
- [DD17] D. Drews and L. D’Antoni. Learning symbolic automata. In *Tools and Algorithms for the Construction and Analysis of Systems - 23rd Inter. Conference, TACAS*, pages 173–189, 2017.
- [DKV09] M. Droste, W. Kuich, and H. Vogler, editors. *Handbook of Weighted Automata*, chapter 4: Rational and Recognizable Series, by Jaques Sakarovitch, pages 105–174. 2009.
- [dlH97] C. de la Higuera. Characteristic sets for polynomial grammatical inference. *Machine Learning*, 27(2):125–138, 1997.
- [dlH10] C. de la Higuera. *Grammatical inference: learning automata and grammars*. 2010.
- [EL87] E.A. Emerson and C-L Lei. Modalities for model checking: Branching time logic strikes back. *Sci. Comput. Program.*, 8(3):275–306, 1987.
- [FBP91] R. Freivalds, J. Bārzdīņš, and K. Podnieks. Inductive inference of recursive functions: Complexity bounds. In J. Bārzdīņš and D. Bjørner, editors, *Baltic Computer Science*, pages 111–155, 1991.
- [FCC⁺08] A. Farzan, Y-F. Chen, E.M. Clarke, Y-K. Tsay, and B-Y. Wang. Extending automated compositional verification to the full class of omega-regular languages. In *TACAS*, pages 2–17, 2008.
- [FFZ23] D. Fisman, H. Frenkel, and S. Zilles. Inferring symbolic automata. *Log. Methods Comput. Sci.*, 19(2), 2023. doi:10.46298/lmcs-19(2:5)2023.
- [FS22] D. Fisman and S. Saadon. Learning and characterizing fully-ordered lattice automata. In *Automated Technology for Verification and Analysis - 20th Inter. Symp., ATVA 2022*, pages 266–282, 2022. doi:10.1007/978-3-031-19992-9_17.
- [GM96] S. A. Goldman and H. D. Mathias. Teaching a smarter learner. *J. Comput. Syst. Sci.*, 52(2):255–267, 1996.
- [Gol67] E.M. Gold. Language identification in the limit. *Information and Control*, 10(5):447–474, 1967.
- [Gol78] E.M. Gold. Complexity of automaton identification from given data. *Information and Control*, 37(3):302–320, 1978.
- [GTW02] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *LNCS*, 2002.
- [HSJC12] F. Howar, B. Steffen, B. Jonsson, and S. Cassel. Inferring canonical register automata. In *Verification, Model Checking, and Abstract Interpretation - 13th Inter. Conference, VMCAI*, pages 251–266, 2012.
- [HV05] P. Habermehl and T. Vojnar. Regular model checking using inference of regular languages. *Electr. Notes Theor. Comput. Sci.*, 138(3):21–36, 2005.
- [KMV06] V. Kumar, P. Madhusudan, and M. Viswanathan. Minimization, learning, and conformance testing of boolean programs. In *CONCUR 2006 - Concurrency Theory, 17th Inter. Conference, CONCUR*, pages 203–217, 2006. URL: https://doi.org/10.1007/11817949_14.
- [KVW00] O. Kupferman, M.Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *J. ACM*, 47(2):312–360, 2000.
- [LCZL21] Y. Li, Y-F. Chen, L. Zhang, and D. Liu. A novel learning algorithm for Büchi automata based on family of dfas and classification trees. *Information and Computation*, 281:104678, 2021.
- [MNRS04] T. Margaria, O. Niese, H. Raffelt, and B. Steffen. Efficient test-based model generation for legacy reactive systems. In *HLDVT*, pages 95–100, 2004.
- [MO20] J. Michaliszyn and J. Otop. Learning deterministic automata on infinite words. In *24th European Conference on Artificial Intelligence (ECAI)*, pages 2370–2377, 2020.
- [MP95] O. Maler and A. Pnueli. On the learnability of infinitary regular sets. *Inf. Comput.*, 118(2):316–326, 1995.

- [MS18] R. Manevich and S. Shoham. Inferring program extensions from traces. In *ICGI*, volume 93 of *Proceedings of Machine Learning Research*, pages 139–154, 2018.
- [NA06] W. Nam and R. Alur. Learning-based symbolic assume-guarantee reasoning with automatic decomposition. In *ATVA*, volume 4218 of *LNCS*, pages 170–185, 2006.
- [NFZ21] D. Nitay, D. Fisman, and M. Ziv-Ukelson. Learning of structurally unambiguous probabilistic grammars. In *Thirty-Fifth Conference on Artificial Intelligence, AAAI*, pages 9170–9178, 2021.
- [OG92] J. Oncina and P. Garcia. Identifying regular languages in polynomial time. In *Advances in structural and syntactic pattern recognition, vol. 5 of series in machine perception and artificial intelligence*, pages 99–108, 1992.
- [Pit89] L. Pitt. Inductive inference, dfas, and computational complexity. In K. P. Jantke, editor, *Analogical and Inductive Inference*, pages 18–44, 1989.
- [PVY99] D. Peled, M.Y. Vardi, and M. Yannakakis. Black box checking. In *FORTE*, pages 225–240, 1999.
- [Sak09] J. Sakarovitch. *Elements of Automata Theory*. Cambridge University Press, 2009.
- [Sch10] S. Schewe. Beyond hyper-minimisation—minimising DBAs and DPAs is np-complete. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*, pages 400–411, 2010.
- [Sch11] S. Schewe. Minimisation of deterministic parity and Buchi automata and relative minimisation of deterministic finite automata. *arXiv, 1007.1333, cs.FL*, 2011. [arXiv:1007.1333](https://arxiv.org/abs/1007.1333).
- [SHV16] M. Schuts, J. Hooman, and F.W. Vaandrager. Refactoring of legacy software using model learning and equivalence checking: An industrial experience report. In *Integrated Formal Methods - 12th Inter. Conference, IFM*, pages 311–325, 2016.
- [Tar72] R.E. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.
- [Vaa17] F. Vaandrager. Model learning. *Commun. ACM*, 60(2):86–95, 2017.
- [Var95] M.Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Banff Higher Order Workshop*, volume 1043 of *LNCS*, pages 238–266, 1995.
- [Var08] M.Y. Vardi. Automata-theoretic model checking revisited. In *Hardware and Software: Verification and Testing, 4th Inter. Haifa Verification Conference, HVC*, page 2, 2008.
- [VSVA05] A. Vardhan, K. Sen, M. Viswanathan, and G. Agha. Using language inference to verify omega-regular properties. In *Tools and Algorithms for the Construction and Analysis of Systems, 11th Inter. Conference, TACAS*, pages 45–60, 2005.
- [Wag75] K.W. Wagner. A hierarchy of regular sequence sets. In *4th Symp. on Mathematical Foundations of Computer (MFCS)*, pages 445–449, 1975.
- [ZZ08] T. Zeugmann and S. Zilles. Learning recursive functions: A survey. *Theoretical Computer Science*, 397(1):4–56, 2008. Forty Years of Inductive Inference: Dedicated to the 60th Birthday of Rolf Wiehagen. URL: <https://www.sciencedirect.com/science/article/pii/S0304397508001400>, doi:10.1016/j.tcs.2008.02.021.