

Learning DFA from Simple Examples*

RAJESH PAREKH

Allstate Research and Planning Center, 321 Middlefield Road, Menlo Park CA 94025, USA

rpare@allstate.com

VASANT HONAVAR**

Department of Computer Science, Iowa State University, Ames IA 50011, USA

honavar@cs.iastate.edu

Abstract. Efficient learning of DFA is a challenging research problem in *grammatical inference*. It is known that both exact and approximate (in the PAC sense) identifiability of DFA is hard. Pitt, in his seminal paper posed the following open research problem: “Are DFA PAC-identifiable if examples are drawn from the uniform distribution, or some other known simple distribution?” [25]. We demonstrate that the class of simple DFA (i.e., DFA whose canonical representations have logarithmic Kolmogorov complexity) is efficiently PAC learnable under the Solomonoff Levin universal distribution. We prove that if the examples are sampled at random according to the universal distribution by a teacher that is knowledgeable about the target concept, the entire class of DFA is efficiently PAC learnable under the universal distribution. Thus, we show that DFA are efficiently learnable under the PACS model [6]. Further, we prove that any concept that is learnable under Gold’s model for learning from characteristic samples, Goldman and Mathias’ polynomial teachability model, and the model for learning from example based queries is also learnable under the PACS model.

1. Introduction

The problem of learning the minimum state DFA that is consistent with a given sample has been actively studied for over two decades. DFAs are recognizers for *regular* languages which constitute the simplest class in the Chomsky hierarchy of formal languages [5, 15]. An understanding of the issues and problems encountered in learning regular languages (or equivalently, identification of the corresponding DFA) are therefore likely to provide insights into the problem of learning more general classes of languages.

Exact learning of the target DFA from an arbitrary presentation of labeled examples is a hard problem [11]. Gold showed that the problem of identifying the minimum state DFA consistent with a presentation S comprising of a finite non-empty set of positive examples S^+ and possibly a finite non-empty set of negative examples S^- is *NP-hard*. Under the standard *complexity theoretic* assumption $P \neq NP$, Pitt and Warmuth showed that no polynomial time algorithm can be guaranteed to produce a DFA with at most $N^{(1-\epsilon)\log\log(N)}$ states from a set of labeled examples corresponding to a DFA with N states [26].

Efficient learning algorithms for identification of DFA assume that additional information is provided to the learner. Trakhtenbrot and Barzdin described a polynomial time algorithm for constructing the smallest DFA consistent with a *complete labeled sample* i.e., a sample that includes all strings up to a particular length and the corresponding label that states whether the string is accepted by the target DFA or not [28]. Angluin showed that given a *live-complete* set of examples (that contains a representative string for each live state of the target DFA) and a knowledgeable teacher to answer *membership queries* it is possible to exactly learn the target DFA [1]. In a later paper, Angluin relaxed the requirement of a live-complete set and has designed a polynomial time inference algorithm using both *membership* and *equivalence* queries [2]. The *regular positive and negative inference* (RPNI) algorithm is a framework for identifying in polynomial time, a DFA consistent with a given sample S [21]. Further, if S is a superset of a characteristic set (see section 2.1) for the target DFA then the DFA output by the RPNI algorithm is guaranteed to be equivalent to the target [21, 9].

Pitt surveyed several approaches for *approximate* identification of DFA [25]. Valiant’s distribution-independent model of learning, also called the *probably approximately correct* (PAC) learning model [29],

* The results in section 5 were presented earlier in [24].

** This research was partially supported by the National Science Foundation grants IRI-9409580 and IRI-9643299 to Vasant Honavar.

is a widely used framework for approximate learning of concept classes. PAC learning models natural learning in that it is fast (learning takes place in polynomial time) and it suffices to learn approximately. When adapted to the problem of learning DFA, the goal of a PAC learning algorithm is to obtain in polynomial time, with high probability, a DFA that is a good approximation of the target DFA. We define PAC learning of DFA more formally in section 2. Angluin’s L^* algorithm [2] that learns DFA in polynomial time using *membership* and *equivalence* queries can be recast under the PAC framework to learn by posing membership queries alone. Even approximate learnability of DFA was proven to be a hard problem. Pitt and Warmuth showed that the problem of polynomially approximate predictability of the class of DFA is hard [27]. They used *prediction preserving reductions* to show that if DFAs are polynomially approximately predictable then so are other known hard to predict concept classes such as *boolean formulas*. Further, Kearns and Valiant showed that an efficient algorithm for learning DFA would entail efficient algorithms for solving problems such as breaking the RSA cryptosystem, factoring Blum integers, and detecting *quadratic residues* [17]. Under the standard *cryptographic assumptions* these problems are known to be hard to solve. Thus, they argued that DFA learning is a hard problem.

The PAC model’s requirement of learnability under all conceivable distributions is often considered too stringent for practical learning scenarios. Pitt’s paper identified the following open research problem: “Are DFA’s PAC-identifiable if examples are drawn from the uniform distribution, or some other known simple distribution?” [25]. Using a variant of Trakhtenbrot and Barzdin’s algorithm, Lang empirically demonstrated that random DFAs are approximately learnable from a sparse uniform sample [18]. However, exact identification of the target DFA was not possible even in the average case with a randomly drawn training sample. Several efforts have been made to study the learnability of concept classes under restricted classes of distributions. Li and Vitányi proposed a model for PAC learning with *simple* examples called the *simple PAC* model wherein the class of distributions is restricted to *simple* distributions (see section 4). Following up on this, Denis *et al* proposed a model of learning where examples are drawn at random according to the universal distribution by a teacher that is knowledgeable about the target concept [6]. This model is known as the PACS learning model. In this paper, we present a method for efficient PAC learning of DFA from simple examples. We will prove that the class of *simple* DFA (see section 4) is learnable under the simple PAC model and the entire class of DFA is learnable under the PACS model. Further, we demonstrate how the model of learning from simple examples naturally extends the model of *learning concepts from representative examples* [11], the *polynomial teachability* model [12], and the model of *learning from example based queries* [3] to a probabilistic framework.

This paper is organized as follows: Section 2 briefly introduces some concepts used in the results described in this paper. This includes a discussion of the PAC learning model, Kolmogorov complexity, and the universal distribution. Section 3 reviews the RPNI algorithm for learning DFA. Section 4 discusses the PAC learnability of the class of *simple* DFA under the simple PAC learning model. Section 5 demonstrates the PAC learnability of the entire class of DFA under the PACS learning model. Section 6 analyzes the PACS model in relation with other models for concept learning. Section 7 addresses the issue of collusion that arises because a helpful teacher can potentially encode the target DFA as a labeled training example. Section 8 concludes with a summary of our contributions and a discussion of several interesting directions for future research.

2. Preliminaries

Let Σ be a finite set of symbols called the *alphabet*; Σ^* be the set of strings over Σ ; α, β, γ be strings in Σ^* ; and $|\alpha|$ be the length of the string α . λ is a special string called the *null* string and has length 0. Given a string $\alpha = \beta\gamma$, β is the *prefix* of α and γ is the *suffix* of α . Let $Pref(\alpha)$ denote the set of all prefixes of α . A *language* L is a subset of Σ^* . The set $Pref(L) = \{\alpha \mid \alpha\beta \in L\}$ is the set of *prefixes* of the language and the set $L_\alpha = \{\beta \mid \alpha\beta \in L\}$ is the set of *tails* of α in L . The *standard order* of strings of the alphabet Σ is denoted by $<$. The standard enumeration of strings over $\Sigma = \{a, b\}$ is $\lambda, a, b, aa, ab, ba, bb, aaa, \dots$. The set of *short prefixes* $S_p(L)$ of a language L is defined as $S_p(L) = \{\alpha \in Pref(L) \mid \exists \beta \in \Sigma^* \text{ such that } L_\alpha = L_\beta \text{ and } \beta < \alpha\}$. The *kernel* $N(L)$ of a language

L is defined as $N(L) = \{\lambda\} \cup \{\alpha a \mid \alpha \in S_p(L), a \in \Sigma, \alpha a \in Pref(L)\}$. Given two sets S_1 and S_2 , let $S_1 \setminus S_2$ and $S_1 \oplus S_2$ denote the *set difference* and the *symmetric difference* respectively. Let \ln and \lg denote the log to the bases e and 2 respectively.

2.1. Finite Automata

A *deterministic finite state automaton* (DFA) is a quintuple $A = (Q, \delta, \Sigma, q_0, F)$ where, Q is a finite set of states, Σ is the finite alphabet, $q_0 \in Q$ is the start state, $F \subseteq Q$ is the set of accepting states, and δ is the transition function: $Q \times \Sigma \rightarrow Q$. A state $d_0 \in Q$ such that $\forall a \in \Sigma, \delta(d_0, a) = d_0$ is called a *dead state*. The extension of δ to handle input strings is standard and is denoted by δ^* . The set of all strings accepted by A is its language, $L(A)$. The language accepted by a DFA is called a *regular language*. Fig. 1 shows the state transition diagram for a sample DFA. A *non-deterministic finite automaton* (NFA) is defined just like the DFA except that the transition function δ defines a mapping from $Q \times \Sigma \rightarrow 2^Q$. In general, a finite state automaton (FSA) refers to either a DFA or a NFA.

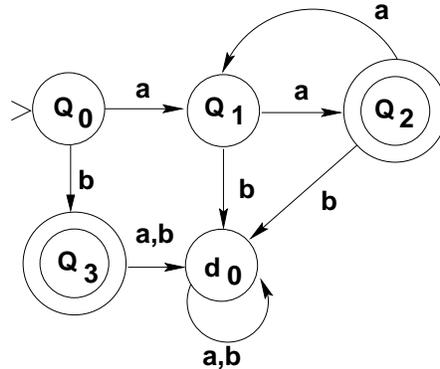


Figure 1. Deterministic Finite State Automaton.

Given any FSA A' , there exists a minimum state DFA (also called the *canonical DFA*, A) such that $L(A) = L(A')$. Without loss of generality, we will assume that the target DFA being learned is a canonical DFA. Let N denote the number of states of A . It can be shown that any canonical DFA has at most one dead state [15]. One can define a standard encoding of DFA as binary strings such that any DFA with N states is encoded as a binary string of length $O(N \lg N)$. A labeled example $(\alpha, c(\alpha))$ for A is such that $\alpha \in \Sigma^*$ and $c(\alpha) = +$ if $\alpha \in L(A)$ (i.e., α is a positive example) or $c(\alpha) = -$ if $\alpha \notin L(A)$ (i.e., α is a negative example). Let S^+ and S^- denote the set of *positive* and *negative* examples of A respectively. A is consistent with a *sample* $S = S^+ \cup S^-$ if it accepts all positive examples and rejects all negative examples.

A set S^+ is said to be *structurally complete* with respect to a DFA A if S^+ covers each transition of A (except the transitions associated with the dead state d_0) and uses every element of the set of final states of A as an accepting state [22, 23, 10]. It can be verified that the set $S^+ = \{b, aa, aaaa\}$ is structurally complete with respect to the DFA in Fig. 1. Given a set S^+ , let $PTA(S^+)$ denote the *prefix tree acceptor* for S^+ . $PTA(S^+)$ is a DFA that contains a path from the start state to an accepting state for each string in S^+ modulo common prefixes. Clearly, $L(PTA(S^+)) = S^+$. Learning algorithms such as the RPNI (see section 3) require the states of the PTA to be numbered in standard order. If we consider the set $Pref(S^+)$ of prefixes of the set S^+ then each state of the PTA corresponds to a unique element in the set $Pref(S^+)$ i.e., for each state q_i of the PTA there exists exactly one string α_i in the set $Pref(S^+)$ such that $\delta^*(q_0, \alpha_i) = q_i$ and vice-versa. The strings of $Pref(S^+)$ are sorted in standard order $<$ and each state q_i is numbered by the position of its corresponding string α_i in the sorted list. The PTA for the set $S^+ = \{b, aa, aaaa\}$ is shown in Fig. 2. Note that its states are numbered in standard order.

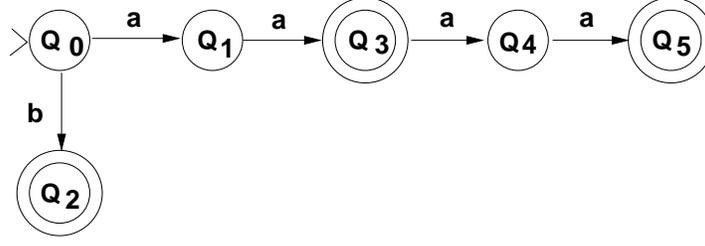


Figure 2. Prefix Tree Automaton.

Given a FSA A and a partition π on the set of states Q of A (ignoring the dead state d_0 and its associated transitions), we define the *quotient automaton* $A_\pi = (Q_\pi, \delta_\pi, \Sigma, B(q_0, \pi), F_\pi)$ obtained by merging the states of A that belong to the same block of the partition π as follows: $Q_\pi = \{B(q, \pi) \mid q \in Q\}$ is the set of states with each state represented uniquely by the block $B(q, \pi)$ of the partition π that contains the state q , $F_\pi = \{B(q, \pi) \mid q \in F\}$ is the set of accepting states, and $\delta_\pi : Q_\pi \times \Sigma \rightarrow 2^{Q_\pi}$ is the transition function such that $\forall B(q_i, \pi), B(q_j, \pi) \in Q_\pi, \forall a \in \Sigma, B(q_j, \pi) = \delta_\pi(B(q_i, \pi), a)$ iff $q_i, q_j \in Q$ and $q_j = \delta(q_i, a)$. Note that a quotient automaton of a DFA might be a NFA and vice-versa. For example, the quotient automaton corresponding to the partition $\pi = \{\{Q_0, Q_1\}, \{Q_2\}, \{Q_3\}\}$ of the set of states of the DFA in Fig. 1 is shown in Fig. 3.

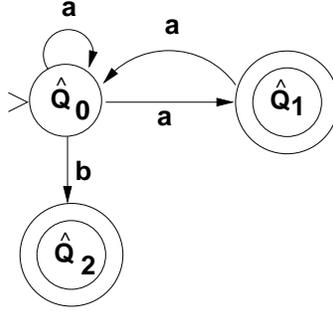


Figure 3. Quotient Automaton.

The set of all quotient automata obtained by systematically merging the states of a DFA A represents a *lattice* of FSA [22]. This lattice is ordered by the *grammar cover* relation \preceq . Given two partitions $\pi_i = \{B_1, B_2, \dots, B_r\}$ and $\pi_j = \{B_1, B_2, \dots, B_k\}$ of the states of A , we say that π_i covers π_j (written $\pi_j \preceq \pi_i$) if $r = k - 1$ and for some $1 \leq l, m \leq k, \pi_i = \{\pi_j \setminus \{B_l, B_m\} \cup \{B_l \cup B_m\}\}$. The *transitive closure* of \preceq is denoted by \ll . We say that $A_{\pi_j} \ll A_{\pi_i}$ iff $L(A_{\pi_j}) \subseteq L(A_{\pi_i})$. Given a canonical DFA A and a set S^+ that is structurally complete with respect to A , the lattice $\Omega(S^+)$ derived from $PTA(S^+)$ is guaranteed to contain A [22, 23, 10].

A sample $S = S^+ \cup S^-$ is said to be *characteristic* with respect to a regular language L (with a canonical acceptor A) if it satisfies the following two conditions [21]:

- $\forall \alpha \in N(L)$, if $\alpha \in L$ then $\alpha \in S^+$ else $\exists \beta \in \Sigma^*$ such that $\alpha\beta \in S^+$.
- $\forall \alpha \in S_p(L), \forall \beta \in N(L)$, if $L_\alpha \neq L_\beta$ then $\exists \gamma \in \Sigma^*$ such that $(\alpha\gamma \in S^+ \text{ and } \beta\gamma \in S^-)$ or $(\beta\gamma \in S^+ \text{ and } \alpha\gamma \in S^-)$.

Intuitively, $S_p(L)$, the set of short prefixes of L is a live complete set with respect to A in that for each live state $q \in Q$, there is a string $\alpha \in S_p(L)$ such that $\delta^*(q_0, \alpha) = q$. The kernel $N(L)$ includes the set of short prefixes as a subset. Thus, $N(L)$ is also a live complete set with respect to A . Further,

$N(L)$ covers every transition between each pair of live states of A . i.e., for all live states $q_i, q_j \in Q$, for all $a \in \Sigma$, if $\delta(q_i, a) = q_j$ then there exists a string $\beta \in N(L)$ such that $\beta = \alpha a$ and $\delta^*(q_0, \alpha) = q_i$. Thus, condition 1 above which identifies a suitably defined suffix $\beta \in \Sigma^*$ for each string $\alpha \in N(L)$ such that the augmented string $\alpha\beta \in L$ implies structural completeness with respect to A . Condition 2 implies that for any two distinct states of A there is a suffix γ that would correctly distinguish them. In other words, for any $q_i, q_j \in Q$ where $q_i \neq q_j$, $\exists \gamma \in \Sigma^*$ such that $\delta^*(q_i, \gamma) \in F$ and $\delta^*(q_j, \gamma) \notin F$ or vice-versa. Given the language L corresponding to the DFA A in Fig. 1, the set of short prefixes is $S_p(L) = \{\lambda, a, b, aa\}$ and the kernel is $N(L) = \{\lambda, a, b, aa, aaa\}$. It can be easily verified that the set $S = S^+ \cup S^-$ where $S^+ = \{b, aa, aaaa\}$ and $S^- = \{\lambda, a, aaa, baa\}$ is a characteristic sample for L .

2.2. PAC Learning of DFA

Let \mathcal{X} denote the *sample space* defined as the set of all strings Σ^* . Let $x \subseteq \mathcal{X}$ denote a *concept*. For our purpose, x is a *regular language*. We identify the concept with the corresponding DFA and denote the class of all DFA as the *concept class* \mathcal{C} . The *representation* \mathcal{R} that assigns a name to each DFA in \mathcal{C} is defined as a function $\mathcal{R} : \mathcal{C} \rightarrow \{0, 1\}^*$. \mathcal{R} is the set of standard encodings of the DFA in \mathcal{C} . Assume that there is an unknown and arbitrary but fixed distribution \mathcal{D} according to which the examples of the target concept are drawn. In the context of learning DFA, \mathcal{D} is restricted to a probability distribution on strings of Σ^* of length at most m .

Definition 1. (due to Pitt [25])

DFA's are PAC-identifiable *iff* there exists a (possibly randomized) algorithm \mathcal{A} such that on input of any parameters ϵ and δ , for any DFA M of size N , for any number m , and for any probability distribution \mathcal{D} on strings of Σ^* of length at most m , if \mathcal{A} obtains labeled examples of M generated according to the distribution \mathcal{D} , then \mathcal{A} produces a DFA M' such that with probability at least $1 - \delta$, the probability (with respect to distribution \mathcal{D}) of the set $\{\alpha \mid \alpha \in L(M) \oplus L(M')\}$ is at most ϵ . The run time of \mathcal{A} (and hence the number of randomly generated examples obtained by \mathcal{A}) is required to be polynomial in N , m , $1/\epsilon$, $1/\delta$, and $|\Sigma|$.

If the learning algorithm \mathcal{A} produces a DFA M' such that with probability at least $1 - \delta$, M' is equivalent to M i.e., the probability (with respect to distribution \mathcal{D}) of the set $\{\alpha \mid \alpha \in L(M) \oplus L(M')\}$ is exactly 0 then \mathcal{A} is said to be a *probably exact* learning algorithm for the class of DFA and the class of DFA is said to be *probably exactly learnable* by the algorithm \mathcal{A} .

2.3. Kolmogorov Complexity

Note that the definition of PAC learning requires that the concept class (in this case the class of DFA) must be learnable under any arbitrary (but fixed) probability distribution. This requirement is often considered too stringent in practical learning scenarios where it is not unreasonable to assume that a learner is first provided with *simple* and *representative* examples of the target concept. Intuitively, when we teach a child the rules of *multiplication* we are more likely to first give simple examples like 3×4 than examples like 1377×428 . A *representative set* of examples is one that would enable the learner to identify the target concept exactly. For example, the characteristic set of a DFA would constitute a suitable representative set. The question now is whether we can formalize what simple examples mean. *Kolmogorov complexity* provides a machine independent notion of *simplicity* of objects. Intuitively, the Kolmogorov complexity of an object (represented by a binary string α) is the length of the shortest binary program that computes α . Objects that have regularity in their structure (i.e., objects that can be easily compressed) have low Kolmogorov complexity. For example, consider the string $s_1 = 010101\dots01 = (01)^{500}$. On a particular machine M , a program to compute this string would be “*Print 01 500 times*”. On the other hand consider a totally random string $s_2 = 110011010\dots00111$. Unlike s_1 , it is not possible to compress the string s_2

which means that a program to compute s_2 on M would be “*Print 1100111010000 . . . 00111*” i.e., the program would have to explicitly specify the string s_2 . The length of the program that computes s_1 is shorter than that of the program that computes s_2 . Thus, we could argue that s_1 has lower Kolmogorov complexity than s_2 with respect to the machine M .

We will consider the *prefix* version of the Kolmogorov complexity that is measured with respect to prefix Turing machines and denoted by K . Consider a prefix Turing machine that implements the partial recursive function $\phi : \{0, 1\}^* \xrightarrow{\text{partial}} \{0, 1\}^*$. For any string $\alpha \in \{0, 1\}^*$, the Kolmogorov complexity of α relative to ϕ is defined as $K_\phi(\alpha) = \min\{|\pi| \mid \phi(\pi) = \alpha\}$ where $\pi \in \{0, 1\}^*$ is a program input to the Turing machine. Prefix Turing machines can be effectively enumerated and there exists a *Universal Turing Machine* (U) capable of simulating every prefix Turing machine. Assume that the universal Turing machine implements the partial function ψ . The *Optimality Theorem* for Kolmogorov Complexity guarantees that for any prefix Turing machine ϕ there exists a constant c_ϕ such that for any string α , $K_\psi(\alpha) \leq K_\phi(\alpha) + c_\phi$. Note that we use the name of the Turing Machine (say M) and the partial function it implements (say ϕ) interchangeably i.e., $K_\phi(\alpha) = K_M(\alpha)$. Further, by the *Invariance Theorem* it can be shown that for any two universal machines ψ_1 and ψ_2 there is a constant $\eta \in \mathcal{N}$ (where \mathcal{N} is the set of natural numbers) such that for all strings α , $|K_{\psi_1}(\alpha) - K_{\psi_2}(\alpha)| \leq \eta$. Thus, we can fix a single universal Turing machine U and denote $K(\alpha) = K_U(\alpha)$. Note that there exists a Turing machine that computes the identity function $\chi : \{0, 1\}^* \rightarrow \{0, 1\}^*$ where $\forall \alpha$, $\chi(\alpha) = \alpha$. Thus, it can be shown that the Kolmogorov complexity of an object is bounded by its length i.e., $K(\alpha) \leq |\alpha| + K(|\alpha|) + \eta$ where η is a constant independent of α .

Suppose that some additional information in the form of a string β is available to the Turing machine ϕ . The conditional Kolmogorov complexity of any object α given β is defined as $K_\phi(\alpha \mid \beta) = \min\{|\pi| \mid \phi(\langle \pi, \beta \rangle) = \alpha\}$ where $\pi \in \{0, 1\}^*$ is a program and $\langle x, y \rangle$ is a standard pairing function¹. Note that the conditional Kolmogorov complexity does not charge for the extra information β that is available to ϕ along with the program π . Fixing a single universal Turing machine U we denote the conditional Kolmogorov complexity of α by $K(\alpha \mid \beta) = K_U(\alpha \mid \beta)$. It can be shown that $K(\alpha \mid \beta) \leq K(\alpha) + \eta$ where η is a constant independent of α .

2.4. Universal Distribution

The set of programs for a string α relative to a Turing machine M is defined as $PROG_M(\alpha) = \{\pi \mid M(\pi) = \alpha\}$. The algorithmic probability of α relative to M is defined as $\mathbf{m}_M(\alpha) = \Pr(PROG_M)$. The algorithmic probability of α with respect to the universal Turing machine U is denoted as $\mathbf{m}_U(\alpha) = \mathbf{m}(\alpha)$. \mathbf{m} is known as the Solomonoff-Levin distribution. It is the universal enumerable probability distribution, in that, it multiplicatively dominates all enumerable probability distributions. Thus, for any enumerable probability distribution P there is a constant $c \in \mathcal{N}$ such that for all strings α , $c \mathbf{m}(\alpha) \geq P(\alpha)$. The *Coding Theorem* due independently to Schnorr, Levin, and Chaitin [20] states that $\exists \eta \in \mathcal{N}$ such that $\forall \alpha \mathbf{m}_M(\alpha) \leq 2^{\eta - K(\alpha)}$. Intuitively this means that if there are several programs for a string α on some machine M then there is a short program for α on the universal Turing machine (i.e., α has a low Kolmogorov complexity). By optimality of \mathbf{m} it can be shown that: $\exists \eta \in \mathcal{N}$, such that $\forall \alpha \in \{0, 1\}^*$, $2^{-K(\alpha)} \leq \mathbf{m}(\alpha) \leq 2^{\eta - K(\alpha)}$. We see that the universal distribution \mathbf{m} assigns higher probability to simple objects (objects with low Kolmogorov complexity). Given a string $r \in \Sigma^*$, the universal distribution based on the knowledge of r , \mathbf{m}_r , is defined as is defined as $\mathbf{m}_r(\alpha) = \lambda_r 2^{-K(\alpha \mid r)}$ where $\lambda_r \sum_{\alpha \in \Sigma^*} 2^{-K(\alpha \mid r)} = 1$ (i.e., $\lambda_r \geq 1$) [6]. Further, \mathbf{m}_r is such that $2^{-K(\alpha \mid r)} \leq \mathbf{m}_r(\alpha) \leq 2^{\eta - K(\alpha \mid r)}$ where η is a constant.

The interested reader is referred to [20] for a thorough treatment of Kolmogorov complexity, universal distribution, and related topics.

3. The RPNI Algorithm

The *regular positive and negative inference* (RPNI) algorithm [21] is a polynomial time algorithm for identification of a DFA consistent with a given set $S = S^+ \cup S^-$. If the sample is a characteristic set for the target DFA then the algorithm is guaranteed to return a canonical representation of the target DFA. Our description of the RPNI algorithm is based on the explanation given in [8].

A labeled sample $S = S^+ \cup S^-$ is provided as input to the algorithm. It constructs a prefix tree automaton $PTA(S^+)$ and numbers its states in the standard order. Then it performs an ordered search in the space of partitions of the set of states of $PTA(S^+)$ under the control of the set of negative examples S^- . The partition, π_0 , corresponding to the automaton $PTA(S^+)$ itself is $\{\{0\}, \{1\}, \dots, \{\overline{N} - 1\}\}$ where \overline{N} is the number of states of the PTA. $M_{\pi_0} = PTA(S^+)$ is consistent with all the training examples and is treated as the initial hypothesis. The current hypothesis is denoted by M_π and the corresponding partition is denoted by π . The algorithm is outlined in Fig. 4. The nested *for* loop refines the partition π by merging the states of $PTA(S^+)$ in order. At each step, a partition $\tilde{\pi}$ is obtained from the partition π by merging the two blocks that contain the states i and j respectively. The function *derive* obtains the quotient automaton $M_{\tilde{\pi}}$, corresponding to the partition $\tilde{\pi}$. $M_{\tilde{\pi}}$ might be a NFA in which case the function *deterministic_merge* determinizes it by recursively merging the states that cause non-determinism. For example, if q_i, q_j , and q_k are states of $M_{\tilde{\pi}}$ such that for some $a \in \Sigma$, $\delta(q_i, a) = \{q_j, q_k\}$ then the states q_j and q_k are merged together. This recursive merging of states can go on for at most $\overline{N} - 1$ steps and the resulting automaton $M_{\hat{\pi}}$ is guaranteed to be a DFA [8]. Note that since $\tilde{\pi} \ll \hat{\pi}$ we know by the grammar covers relation that if $M_{\tilde{\pi}}$ accepts a negative example in S^- then so would $M_{\hat{\pi}}$. The function, *consistent*($M_{\hat{\pi}}, S^-$) returns *True* if $M_{\hat{\pi}}$ is consistent with all examples in S^- and *False* otherwise. If a partition $\hat{\pi}$ is found such that the corresponding DFA $M_{\hat{\pi}}$ is consistent with S^- then $M_{\hat{\pi}}$ replaces M_π as the current hypothesis.

Let $\|S^+\|$ and $\|S^-\|$ denote the sums of the lengths of examples in S^+ and S^- respectively. $PTA(S^+)$ has $O(\|S^+\|)$ states. The nested *for* loop of the algorithm performs $O(\|S^+\|^2)$ state merges. Further, each time two blocks of the partition π are merged, the routine *deterministic_merge* in the worst case would cause $O(\|S^+\|)$ state mergings and the function *consistent* that checks for the consistency of the derived DFA with the negative examples would incur a cost of $O(\|S^-\|)$. Hence the time complexity of the RPNI algorithm is $O((\|S^+\| + \|S^-\|) \cdot \|S^+\|^2)$.

Example

We demonstrate the execution of the RPNI algorithm on the task of learning the DFA in Fig. 1. Note that for convenience we have shown the target DFA in Fig. 5 without the dead state d_0 and its associated transitions. Assume that a sample $S = S^+ \cup S^-$ where $S^+ = \{b, aa, aaaa\}$ and $S^- = \{\lambda, a, aaa, baa\}$. It can be easily verified that S is a characteristic sample for the target DFA. The DFA $M = PTA(S^+)$ is depicted in Fig. 2 where the states are numbered in the standard order. The initial partition is $\pi = \pi_0 = \{\{0\}, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}\}$.

The algorithm attempts to merge the blocks containing states 1 and 0 of the partition π . The quotient FSA $M_{\tilde{\pi}}$ and the DFA $M_{\hat{\pi}}$ obtained after invoking *deterministic_merge* are shown in Fig. 6. The DFA $M_{\hat{\pi}}$ accepts the negative example $\lambda \in S^-$. Thus, the current partition π remains unchanged.

Next the algorithm merges the blocks containing states 2 and 0 of the partition π . The quotient FSA $M_{\tilde{\pi}}$ is depicted in Fig. 7. Since $M_{\tilde{\pi}}$ is a DFA, the procedure *deterministic_merge* returns the same automaton i.e., $M_{\tilde{\pi}} = M_{\hat{\pi}}$. $M_{\hat{\pi}}$ accepts the negative example $\lambda \in S^-$ and hence the partition π remains unchanged.

Table 1 lists the different partitions $\tilde{\pi}$ obtained by fusing the blocks of π_0 , the partitions $\hat{\pi}$ obtained by *deterministic_merge* of $\tilde{\pi}$, and the negative example (belonging to S^-), if any, that is accepted by the quotient FSA $M_{\hat{\pi}}$. The partitions marked * denote the partition π for which M_π is consistent with all

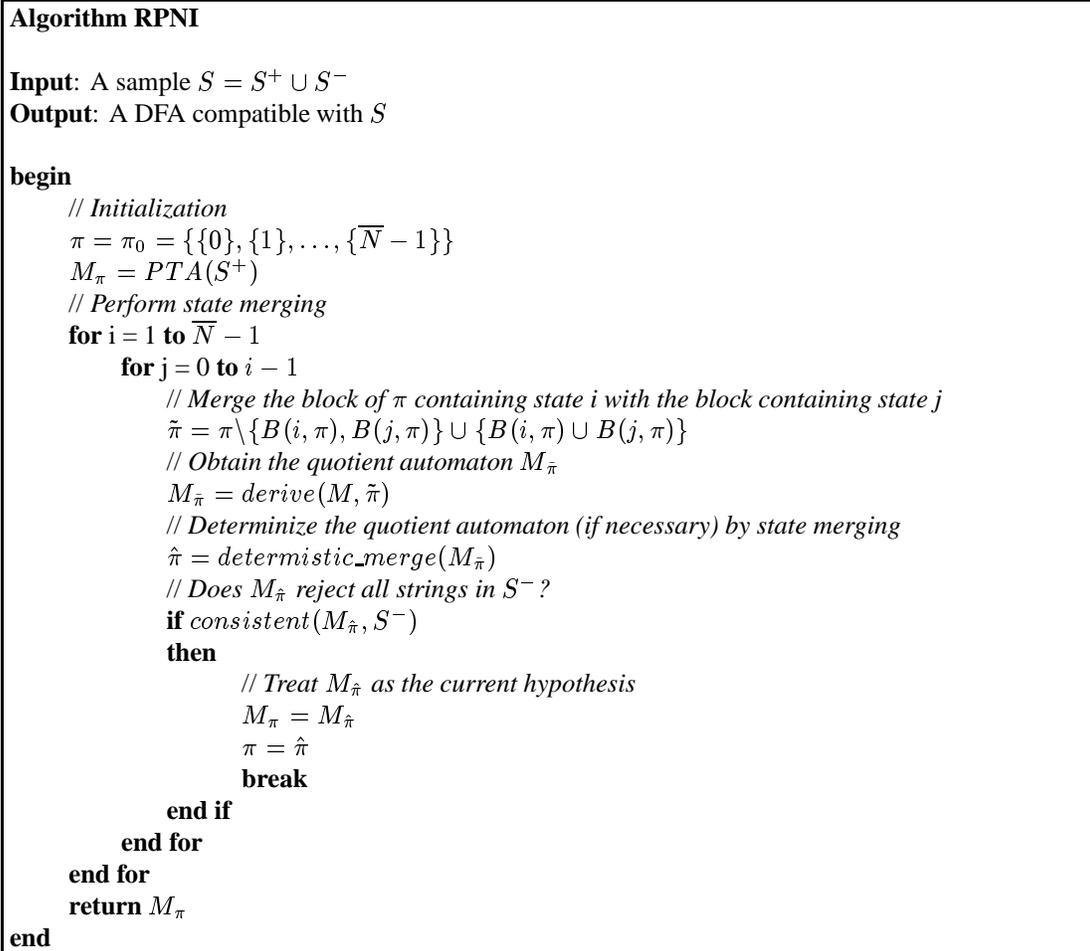


Figure 4. RPNI Algorithm.

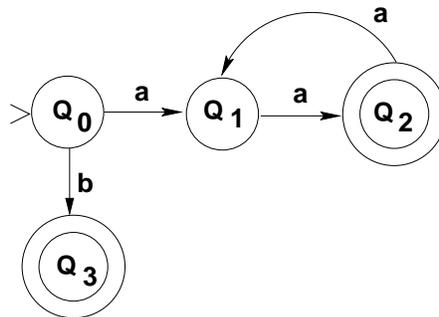


Figure 5. Target DFA A.

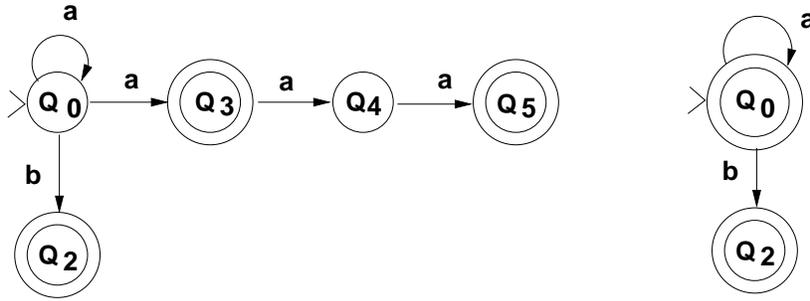


Figure 6. $M_{\tilde{\pi}}$ Obtained by Fusing Blocks Containing the States 1 and 0 of π and the Corresponding $M_{\tilde{\pi}}$

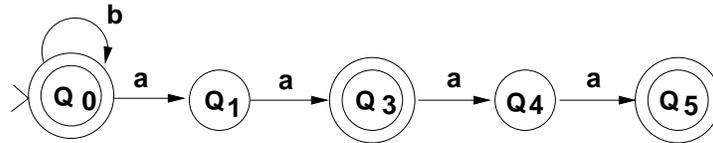


Figure 7. $M_{\tilde{\pi}}$ (same as $M_{\tilde{\pi}}$) Obtained by Fusing Blocks Containing the states 2 and 0 of π .

examples in S^- and hence is the current hypothesis. It is easy to see that the DFA corresponding to the partition $\pi = \{\{0\}, \{1, 4\}, \{2\}, \{3, 5\}\}$ is exactly the target DFA we are trying to learn (Fig. 1).

Table 1. Execution of the RPNI Algorithm.

Partition $\tilde{\pi}$	Partition $\tilde{\pi}$	Negative Example
$\{\{0, 1\}, \{2\}, \{3\}, \{4\}, \{5\}\}$	$\{\{0, 1, 3, 4, 5\}, \{2\}\}$	a
$\{\{0, 2\}, \{1\}, \{3\}, \{4\}, \{5\}\}$	$\{\{0, 2\}, \{1\}, \{3\}, \{4\}, \{5\}\}$	λ
$\{\{0\}, \{1, 2\}, \{3\}, \{4\}, \{5\}\}$	$\{\{0\}, \{1, 2\}, \{3\}, \{4\}, \{5\}\}$	a
$\{\{0, 3\}, \{1\}, \{2\}, \{4\}, \{5\}\}$	$\{\{0, 3\}, \{1, 4\}, \{2\}, \{5\}\}$	λ
$\{\{0\}, \{1, 3\}, \{2\}, \{4\}, \{5\}\}$	$\{\{0\}, \{1, 3, 4, 5\}, \{2\}\}$	a
$\{\{0\}, \{1\}, \{2, 3\}, \{4\}, \{5\}\}$	$\{\{0\}, \{1\}, \{2, 3\}, \{4\}, \{5\}\}$	baa
$\{\{0, 4\}, \{1\}, \{2\}, \{3\}, \{5\}\}$	$\{\{0, 4\}, \{1, 5\}, \{2\}, \{3\}\}$	a
$\{\{0\}, \{1, 4\}, \{2\}, \{3\}, \{5\}\}$	$\{\{0\}, \{1, 4\}, \{2\}, \{3, 5\}\}^*$	—
$\{\{0, 3, 5\}, \{1, 4\}, \{2\}\}$	$\{\{0, 3, 5\}, \{1, 4\}, \{2\}\}$	λ
$\{\{0\}, \{1, 3, 4, 5\}, \{2\}\}$	$\{\{0\}, \{1, 3, 4, 5\}, \{2\}\}$	a
$\{\{0\}, \{1, 4\}, \{2, 3, 5\}\}$	$\{\{0\}, \{1, 4\}, \{2, 3, 5\}\}$	baa
$\{\{0\}, \{1, 4\}, \{2\}, \{3, 5\}\}$	$\{\{0\}, \{1, 4\}, \{2\}, \{3, 5\}\}^*$	—
$\{\{0\}, \{1, 3, 4, 5\}, \{2\}\}$	$\{\{0\}, \{1, 3, 4, 5\}, \{2\}\}$	a

4. Learning Simple DFA under the Simple PAC model

Li and Vitányi have proposed the simple PAC learning model where the class of probability distributions is restricted to *simple* distributions [19]. A distribution is simple if it is multiplicatively dominated by some enumerable distribution. Simple distributions include a variety of distributions including all computable distributions. Further, the *simple distribution independent learning theorem* due to Li and Vitányi says that a concept class is learnable under universal distribution \mathbf{m} iff it is learnable under the entire class of *simple distributions* provided the examples are drawn according to the universal distribution [19]. Thus, the simple PAC learning model is sufficiently general. Concept classes such as *log n-term DNF* and *simple k-reversible DFA* are learnable under the simple PAC model whereas their PAC learnability in the standard sense is unknown [19]. We show that the class of *simple* DFA is polynomially learnable under the simple PAC learning model.

A simple DFA is one with low Kolmogorov complexity. More specifically, a DFA A with N states and a standard encoding (or canonical representation) r is simple if $K(A) = O(\lg N)$. For example, a DFA that accepts all strings of length N is a simple DFA. Note however that this DFA contains a path for every string of length N and hence it has a path of Kolmogorov complexity N . In general, simple DFA might actually have very random paths. We saw in section 2.3 that a natural learning scenario would typically involve learning from a *simple* and *representative* set of examples for the target concept. We adopt Kolmogorov complexity as a measure of simplicity and define simple examples as those with low Kolmogorov complexity i.e., with Kolmogorov complexity $O(\lg N)$. Further, a characteristic set for the DFA A can be treated as its representative set.

We demonstrate that for every simple DFA there exists a characteristic set of simple examples S_c .

LEMMA 1 *For any N state simple DFA (with Kolmogorov complexity $O(\lg N)$) there exists a characteristic set of simple examples S_c such that the length of each string in this set is at most $2N - 1$.*

Proof: Consider the following enumeration of a characteristic set of examples for a DFA $A = (Q, \delta, \Sigma, q_0, F)$ with N states².

1. Fix an enumeration of the shortest paths (in standard order) from the state q_0 to each state in Q except the dead state. This is the set of short prefixes of A . There are at most N such paths and each path is of length at most $N - 1$.
2. Fix an enumeration of paths that includes each path identified above and its extension by each letter of the alphabet Σ . From the paths just enumerated retain only those that do not terminate in the dead state of A . This represents the kernel of A . There are at most $N(|\Sigma| + 1)$ such paths and each path is of length at most N .
3. Let the characteristic set be denoted by $S_c = S_c^+ \cup S_c^-$.
 - (A) For each string α identified in step 2 above, determine the first suffix β in the standard enumeration of strings such that $\alpha\beta \in L(A)$. Since $|\alpha| \leq N$, and β is the shortest suffix in the standard order it is clear that $|\alpha\beta| \leq 2N - 1$. Each such $\alpha\beta$ is a member of S_c^+ .
 - (B) For each pair of strings (α, β) in order where α is a string identified in step 1, β is a string identified in step 2, and α and β lead to different states of A determine the first suffix γ in the standard enumeration of strings such that $\alpha\gamma \in L(A)$ and $\beta\gamma \notin L(A)$ or vice versa. Since $|\alpha| \leq N - 1$, $|\beta| \leq N$, and γ is the shortest distinguishing suffix for the states represented by α and β it is clear that $|\alpha\gamma|, |\beta\gamma| \leq 2N - 1$. The accepted string from among $\alpha\gamma$ and $\beta\gamma$ is a member of S_c^+ and the rejected string is a member of S_c^- .

Trivial upper bounds on the sizes of S_c^+ and S_c^- are $|S_c^+| \leq N^2(|\Sigma| + 1) + N(|\Sigma|)$, $|S_c^-| \leq N^2(|\Sigma| + 1) - N$. Thus, $|S_c| = O(N^2)$. Further, the length of each string in S_c is less than $2N - 1$.

The strings in S_c can be ordered in some way such that individual strings can be identified by an index of length at most $\lg(3|\Sigma|N^2)$ bits. There exists a Turing machine M that implements the above algorithm for constructing the set S_c . M can take as input an encoding of a simple DFA of length $O(\lg N)$ bits and an index of length $\lg(3|\Sigma|N^2)$ bits and output the corresponding string α belonging to S_c . Thus, $\forall \alpha \in S_c$,

$$\begin{aligned} K(\alpha) &\leq k_1 \lg N + \lg(3|\Sigma|N^2) \\ K(\alpha) &\leq k_1 \lg N + k_2 \lg N \\ &= O(\lg N) \end{aligned}$$

This proves the lemma. ■

LEMMA 2 *Suppose a sample S is drawn according to \mathbf{m} . For $0 < \delta \leq 1$, if $|S| = O(N^k \lg(\frac{1}{\delta}))$ then with probability greater than $1 - \delta$, $S_c \subseteq S$ where k is a constant.*

Proof: From lemma 1 we know that $\forall \alpha \in S_c$, $K(\alpha) = O(\lg N)$. Further, $|S_c| = O(N^2)$. By definition, $\mathbf{m}(\alpha) \geq 2^{-K(\alpha)}$. Thus, $\mathbf{m}(\alpha) \geq 2^{-k_1 \lg N}$ or equivalently $\mathbf{m}(\alpha) \geq N^{-k_1}$ where k_1 is a constant.

$$\begin{aligned} \Pr(\alpha \in S_c \text{ is not sampled in one random draw}) &\leq (1 - N^{-k_1}) \\ \Pr(\alpha \in S_c \text{ is not sampled in } |S| \text{ random draws}) &\leq (1 - N^{-k_1})^{|S|} \\ \Pr(\text{some } \alpha \in S_c \text{ is not sampled in } |S| \text{ random draws}) &\leq |S_c|(1 - N^{-k_1})^{|S|} \\ &\leq k_2 N^2 (1 - N^{-k_1})^{|S|} \\ &\quad \text{since } |S_c| = O(N^2) \\ \Pr(S_c \not\subseteq S) &\leq k_2 N^2 (1 - N^{-k_1})^{|S|} \end{aligned}$$

We want this probability to be less than δ .

$$\begin{aligned} k_2 N^2 (1 - N^{-k_1})^{|S|} &\leq \delta \\ k_2 N^2 (e^{-N^{-k_1}})^{|S|} &\leq \delta \text{ since } 1 - x \leq e^{-x} \\ \ln(k_2) + \ln(N^2) - N^{-k_1} |S| &\leq \ln(\delta) \\ |S| &\geq N^{k_1} (\ln(\frac{1}{\delta}) + \ln(k_2) + \ln(N^2)) \\ &= O(N^k \lg(\frac{1}{\delta})) \text{ where } k \text{ replaces } k_1 \end{aligned}$$

Thus, $\Pr(S_c \subseteq S) \geq 1 - \delta$. ■

We now prove that the class of simple DFA is polynomially learnable under \mathbf{m} .

THEOREM 1 *For all N , the class $\mathcal{C}^{\leq N}$ of simple DFA whose canonical representations have at most N states is probably exactly learnable under the simple PAC model.*

Proof: Let A be a simple DFA with at most N states. Let S_c be a characteristic sample of A enumerated as described in lemma 1 above. Recall, that the examples in S_c are simple (i.e., each example has Kolmogorov complexity $O(\lg N)$). Now consider the algorithm \mathcal{A}_1 in Fig. 8 that draws a sample S with the following properties.

1. $S = S^+ \cup S^-$ is a set of positive and negative examples corresponding to the target DFA A .
2. The examples in S are drawn at random according to the distribution \mathbf{m} .
3. $|S| = O(N^k \lg(\frac{1}{\delta}))$.

Lemma 1 showed that for every simple DFA A there exists a characteristic set of simple examples S_c where each example is of length at most $2N - 1$. Lemma 2 showed that if a labeled sample S of size $O(N^k \lg(\frac{1}{\delta}))$ is randomly drawn according to \mathbf{m} then with probability greater than $1 - \delta$, $S_c \subseteq S$. The RPNI algorithm is guaranteed to return a canonical representation of the target DFA A if the set of examples S provided is a superset of a characteristic set S_c . Since the size of S is polynomial in N and $1/\delta$ and the length of each string in S is restricted to $2N - 1$, the RPNI algorithm, and thus the algorithm \mathcal{A}_1 can be implemented to run in time polynomial in N and $1/\delta$. Thus, with probability greater than $1 - \delta$, \mathcal{A}_1 is guaranteed to return a canonical representation of the target DFA A . This proves that the class $\mathcal{C}^{\leq N}$ of simple DFA whose canonical representations have at most N states is exactly learnable with probability greater than $1 - \delta$. ■

Algorithm \mathcal{A}_1

Input: $N, 0 < \delta \leq 1$
Output: A DFA M

begin
 Randomly draw a labeled sample S according to \mathbf{m}
 Retain only those examples in S that have length at most $2N - 1$
 $M = RPNI(S)$
return M
end

Figure 8. A Probably Exact Algorithm for Learning Simple DFA.

5. Learning DFA under the PACS model

In section 4 we proved that the class of simple DFA is learnable under the simple PAC model where the underlying distribution is restricted to the universal distribution \mathbf{m} . Denis *et al* proposed a model of learning where examples are drawn at random according to the universal distribution by a teacher that is knowledgeable about the target concept [6]. Under this model, examples with low conditional Kolmogorov complexity given a representation r of the target concept are called simple examples. Specifically, for a concept with representation r , the set $S_{sim}^r = \{\alpha \mid K(\alpha|r) \leq \mu \lg(|r|)\}$ (where μ is a constant) is the set of simple examples for the concept. Further, $S_{sim,rep}^r$ is used to denote a set of simple and representative examples of r . The PACS model restricts the underlying distribution to \mathbf{m}_r . Formally, the probability of drawing an example α for a target concept with representation r is given as $\mathbf{m}_r(\alpha) = \lambda_r 2^{-K(\alpha|r)}$. Representative examples for the target concept are those that enable the learner to exactly learn the target. As explained earlier, the characteristic set corresponding to a DFA can be treated as a representative set for the DFA. The Occam's Razor theorem proved by Denis *et al* states that if there exists a representative set of simple examples for each concept in a concept class then the concept class is PACS learnable [6].

We now demonstrate that the class of DFA is efficiently learnable under the PACS model. Lemma 3 proves that for any DFA A with standard encoding r there exists a characteristic set of simple examples $S_{sim,rep}^r$.

LEMMA 3 *For any N state DFA with standard encoding r ($|r| = \mathcal{O}(N \lg(N))$), there exists a characteristic set of simple examples (denoted by $S_{sim,rep}^r$) such that each string of this set is of length at most $2N - 1$.*

Proof: Given a DFA $A = (Q, \delta, \Sigma, q_0, F)$, it is possible to enumerate a characteristic set of examples S_c for A as described in lemma 1 such that $|S_c| = \mathcal{O}(N^2)$ and each example of S_c is of length at most $2N - 1$. Individual strings in S_c can be identified by specifying an index of length at most $\lg(3|\Sigma|N^2)$ bits. There exists a Turing machine M that implements the above algorithm for constructing the set S_c . Given the knowledge of the target concept r , M can take as input an index of length $\lg(3|\Sigma|N^2)$ bits and output the corresponding string belonging to S_c . Thus, $\forall \alpha \in S_c$

$$\begin{aligned} K(\alpha|r) &\leq \lg(3|\Sigma|N^2) \\ &\leq \mu \lg(|r|) \text{ where } \mu \text{ is a constant} \end{aligned}$$

We define the set S_c to be the characteristic set of simple examples $S_{sim,rep}^r$ for the DFA A . This proves the lemma. ■

LEMMA 4 (Due to Denis et al [6])

Suppose that a sample S is drawn according to \mathbf{m}_r . For an integer $l \geq |r|$, and $0 < \delta \leq 1$, if $|S| = O(l^\mu \lg(\frac{1}{\delta}))$ then with probability greater than $1 - \delta$, $S_{sim}^r \subseteq S$.

Proof: Claim 4.1: $\forall \alpha \in S_{sim}^r, \mathbf{m}_r(\alpha) \geq l^{-\mu}$

$$\begin{aligned} \mathbf{m}_r(\alpha) &\geq 2^{-K(\alpha|r)} \\ &\geq 2^{-\mu \lg|r|} \\ &\geq |r|^{-\mu} \\ &\geq l^{-\mu} \end{aligned}$$

Claim 4.2: $|S_{sim}^r| \leq 2l^\mu$

$$\begin{aligned} |S_{sim}^r| &\leq |\{\alpha \in \{0,1\}^* \mid K(\alpha|r) \leq \mu \lg(|r|)\}| \\ &\leq |\{\alpha \in \{0,1\}^* \mid K(\alpha|r) \leq \mu \lg(l)\}| \\ &\leq |\{\beta \in \{0,1\}^* \mid |\beta| \leq \mu \lg(l)\}| \\ &\leq 2^{\mu \lg(l)+1} \\ &\leq 2l^\mu \end{aligned}$$

Claim 4.3: If $|S| = O(l^\mu \lg(\frac{1}{\delta}))$ then $\Pr(S_{sim}^r \subseteq S) \geq 1 - \delta$

$$\Pr(\alpha \in S_{sim}^r \text{ is not sampled in one random draw}) \leq (1 - l^{-\mu}) \quad (\text{claim 4.1})$$

$$\Pr(\alpha \in S_{sim}^r \text{ is not sampled in } |S| \text{ random draws}) \leq (1 - l^{-\mu})^{|S|}$$

$$\Pr(\text{some } \alpha \in S_{sim}^r \text{ is not sampled in } |S| \text{ random draws}) \leq 2l^\mu (1 - l^{-\mu})^{|S|} \quad (\text{claim 4.2})$$

$$\Pr(S_{sim}^r \not\subseteq S) \leq 2l^\mu (1 - l^{-\mu})^{|S|}$$

We would like this probability to be less than δ .

$$\begin{aligned} 2l^\mu (1 - l^{-\mu})^{|S|} &\leq \delta \\ 2l^\mu (e^{-l^{-\mu}})^{|S|} &\leq \delta, \quad \text{since } 1 - x \leq e^{-x} \\ \ln(2) + \ln(l^\mu) - |S|l^{-\mu} &\leq \ln(\delta) \\ |S| &\geq l^\mu (\ln(2) + \ln(l^\mu) + \ln(1/\delta)) \\ |S| &\geq O(l^\mu \lg(1/\delta)) \end{aligned}$$

Thus, $\Pr(S_{sim}^r \subseteq S) \geq 1 - \delta$ ■

COROLLARY 1 Suppose that a sample S is drawn according to \mathbf{m}_r . For an integer $l \geq |r|$, and $0 < \delta \leq 1$, if $|S| = O(l^\mu \lg(1/\delta))$ then with probability greater than $1 - \delta$, $S_{sim,rep}^r \subseteq S$.

Proof: Follows immediately from Lemma 3 since $S_{sim,rep}^r \subseteq S_{sim}^r$. ■

We now prove that the class of DFA is polynomially learnable under \mathbf{m}_r .

THEOREM 2 For all N , the class $\mathcal{C}^{\leq N}$ of DFA whose canonical representations have at most N states is probably exactly learnable under the PACS model.

Proof: Let A be a canonical DFA with at most N states and r be its standard encoding. We define the simple representative sample $S_{sim,rep}^r$ to be the characteristic sample of A enumerated as described

in lemma 3. Recall that the length of each example in $S_{sim,rep}^r$ is at most $2N - 1$. Now consider the algorithm \mathcal{A}_2 that draws a sample S with the following properties

1. $S = S^+ \cup S^-$ is a set of positive and negative examples corresponding to the target DFA A
2. The examples in S are drawn at random according to the distribution \mathbf{m}_r
3. $|S| = O(l^\mu \lg(\frac{1}{\delta}))$

Algorithm \mathcal{A}_2

Input: $N, 0 < \delta \leq 1$
Output: A DFA M

begin

Randomly draw a labeled sample S according to \mathbf{m}_r
 Retain only those examples in S that have length at most $2N - 1$
 $M = RPNI(S)$
return(M)

end

Figure 9. A Probably Exact Algorithm for Learning DFA.

Lemma 3 showed that for every DFA A there exists a characteristic set of simple examples $S_{sim,rep}^r$. Corollary 1 showed that if a labeled sample S of size $O(l^\mu \lg(\frac{1}{\delta}))$ is randomly drawn according to \mathbf{m}_r then with probability greater than $1 - \delta$, $S_{sim,rep}^r \subseteq S$. The RPNI algorithm is guaranteed to return a canonical representation of the target DFA A if the set of examples S is a superset of a characteristic set for A . Since the size of S is polynomial in N and $1/\delta$ and the length of each string in S is restricted to $2N - 1$, the RPNI algorithm, and thus the algorithm \mathcal{A}_2 can be implemented to run in time polynomial in N and $1/\delta$. Thus, with probability greater than $1 - \delta$, \mathcal{A}_2 is guaranteed to return a canonical DFA equivalent to the target A . This proves that the class $\mathcal{C}^{\leq N}$ of DFA whose canonical representations have at most N states is exactly learnable with probability greater than $1 - \delta$. ■

Since the number of states of the target DFA (N) might not be known in advance we present a PAC learning algorithm \mathcal{A}_3 that iterates over successively larger guesses of N . At each step the algorithm draws a random sample according to \mathbf{m}_r , applies the RPNI algorithm to construct a DFA, and tests the DFA using a randomly drawn test sample. If the DFA is consistent with the test sample then the algorithm outputs the DFA and halts. Otherwise the algorithm continues with the next guess for N .

THEOREM 3 *The concept class \mathcal{C} of DFA is learnable in polynomial time under the PACS model.*

Proof: Fig. 10 shows a PAC learning algorithm for DFA.

In algorithm \mathcal{A}_3 the polynomial p is defined such that a sample S of size $p(N, \frac{1}{\delta})$ contains the characteristic set of simple examples $S_{sim,rep}^r$ with probability greater than $1 - \delta$. It follows from corollary 1 that $p(N, \frac{1}{\delta}) = O(l^\mu \lg(1/\delta))$ will satisfy this constraint. The polynomial q is defined as $q(i, \frac{1}{\epsilon}, \frac{1}{\delta}) = \frac{1}{\epsilon} [2 \ln(i + 1) + \ln(\frac{1}{\delta})]$.

Consider the execution of the algorithm \mathcal{A}_3 . At any step i where $i \geq N$, the set S will include the characteristic set of simple examples $S_{sim,rep}^r$ with probability greater than $1 - \delta$ (as proved in lemma 4). In this case the RPNI algorithm will return a DFA M that is equivalent to the target A and hence M will be consistent with the test sample T . Thus, with probability at least $1 - \delta$, the algorithm will halt and correctly output the target DFA.

```

Algorithm  $\mathcal{A}_3$ 
Input:  $\epsilon, \delta$ 
Output: A DFA  $M$ 

begin
  1)  $i = 1, EX = \phi, p(0, 1/\delta) = 0$ 
  2) repeat
      Draw  $p(i, 1/\delta) - p(i - 1, 1/\delta)$  examples according to  $\mathbf{m}_r$ 
      Add the examples just drawn to the set  $EX$ 
      Let  $S$  be the subset of examples in  $EX$  of length at most  $2i - 1$ 
       $M = RPNI(S)$ 
      Draw  $q(i, 1/\epsilon, 1/\delta)$  examples according to  $\mathbf{m}_r$  and call this set  $T$ 
      if  $consistent(M, T)$ 
        then Output  $M$  and halt
      else  $i = i * 2$ 
      end if
    until eternity
end

```

Figure 10. A PAC Algorithm for Learning DFA.

Consider the probability that the algorithm halts at some step i and returns a DFA M with an error greater than ϵ .

$$\begin{aligned}
 \Pr(M \text{ and } A \text{ are consistent on some } \alpha) &\leq 1 - \epsilon \\
 \Pr(M \text{ and } A \text{ are consistent on all } \alpha \in T) &\leq (1 - \epsilon)^{|T|} \\
 &\leq (1 - \epsilon)^{\frac{1}{\epsilon}[2 \ln(i+1) + \ln(\frac{1}{\delta})]} \\
 &\leq e^{-[2 \ln(i+1) + \ln(\frac{1}{\delta})]} \text{ since } 1 - x \leq e^{-x} \\
 &\leq \frac{\delta}{(i+1)^2}
 \end{aligned}$$

The probability that the algorithm halts at step i and returns a DFA with error greater than ϵ is less than $\sum_{i=1}^{\infty} \frac{\delta}{(i+1)^2}$ which is in turn strictly less than δ . Thus, we have shown that with probability greater than $1 - \delta$ the algorithm returns a DFA with error at most ϵ . Further, the run time of the algorithm is polynomial in N , $|\Sigma|$, $\frac{1}{\epsilon}$, $\frac{1}{\delta}$, and m (where m is the length of the longest test example seen by the algorithm). Thus, the class of DFA is efficiently PAC learnable under the PACS model. ■

6. Relationship of the PACS Model to other Learning Models

In this section we study the relationship of the PACS model to other prominent learning models such as Gold's model of *polynomial identifiability from characteristic samples* [11], Goldman and Mathias' *polynomial teachability* model [12], and the model of learning from *example based queries* [3]. We explain how the PACS learning model naturally extends these other models to a probabilistic framework. In the discussion that follows we will let \mathcal{X} be the instance space, \mathcal{C} be the concept class and \mathcal{R} be the set of representations of the concepts in \mathcal{C} .

6.1. Polynomial Identifiability from Characteristic Samples

Gold's model for polynomial identifiability of concept classes from characteristic samples is based on the availability of a polynomial sized characteristic sample for any concept in the concept class and an algorithm which when given a superset of a characteristic set is guaranteed to return, in polynomial time, a representation of the target concept.

Definition 2. (due to Colin de la Higuera [14])

\mathcal{C} is polynomially identifiable from characteristic samples iff there exist two polynomials $p_1()$ and $p_2()$ and an algorithm \mathcal{A} such that

1. Given any sample $S = S^+ \cup S^-$ of labeled examples, \mathcal{A} returns in time $p_1(\|S^+\| + \|S^-\|)$ a representation $r \in \mathcal{R}$ of a concept $c \in \mathcal{C}$ such that c is consistent with S .
2. For every concept $c \in \mathcal{C}$ with corresponding representation $r \in \mathcal{R}$ there exists a characteristic sample $S_c = S_c^+ \cup S_c^-$ such that $\|S_c^+\| + \|S_c^-\| = p_2(|r|)$ and if \mathcal{A} is provided with a sample $S = S^+ \cup S^-$ where $S_c^+ \subseteq S^+$ and $S_c^- \subseteq S^-$ then \mathcal{A} returns a representation r' of a concept c' that is equivalent to c .

Using the above definition Gold's result can be restated as follows.

THEOREM 4 (due to Gold [11])

The class of DFA is polynomially identifiable from characteristic samples.

The problem of identifying a minimum state DFA that is consistent with an arbitrary labeled sample $S = S^+ \cup S^-$ is known to be NP-complete [11]. This result does not contradict the one in theorem 4 because a characteristic set is not any arbitrary set of examples but a special set that enables the learning algorithm to correctly infer the target concept in polynomial time (see the RPNI algorithm in section 3).

6.2. Polynomial Teachability of Concept Classes

Goldman and Mathias developed a teaching model for efficient learning of target concepts [12]. Their model takes into account the quantity of information that a good teacher must provide to the learner. An additional player called the *adversary* is introduced in this model to ensure that there is no collusion whereby the teacher gives the learner an encoding of the target concept. A typical teaching session proceeds as follows:

1. The adversary selects a target concept and gives it to the teacher.
2. The teacher computes a set of examples called the *teaching set*.
3. The adversary adds correctly labeled examples to the teaching set with the goal of complicating the learner's task.
4. The learner computes a hypothesis from the augmented teaching set.

Under this model, a concept class for which the computations of both the teacher and the learner takes polynomial time and the learner always learns the target concept is called *polynomially T/L teachable*. Without the restrictive assumption that the teacher's computations be performed in polynomial time, the concept class is said to be *semi-polynomially T/L teachable*. When this model is adapted to the framework of learning DFA the length of the examples seen by the learner must be included as a parameter in the model. In the context of learning DFA the number of examples is infinite (it includes the entire set Σ^*) and further the lengths of these examples grow unboundedly. A scenario in which the teacher constructs

a very small teaching set whose examples are unreasonably long is clearly undesirable and must be avoided. This is explained more formally in the following definition.

Definition 3. (due to Colin de la Higuera [14])

A concept class \mathcal{C} is semi-polynomially T/L teachable iff there exist polynomials $p_1()$, $p_2()$, and $p_3()$, a teacher T , and a learner L , such that for any adversary ADV and any concept c with representation r that is selected by ADV , after the following teaching session the learner returns the representation r' of a concept c' that is equivalent to c .

1. ADV gives r to T .
2. T computes a teaching set S of size at most $p_1(|r|)$ such that each example in the teaching set has length at most $p_2(|r|)$.
3. ADV adds correctly labeled examples to this set, with the goal of complicating the learner's task.
4. The learner uses the augmented set S to compute a hypothesis r' in time $p_3(|S|)$.

Note that from Gold's result (theorem 4) it follows that DFA are semi-polynomially T/L teachable. Further, we demonstrated in lemma 1 that for any DFA there exists a procedure to enumerate a characteristic set corresponding to that DFA. This procedure can be implemented in polynomial time thereby proving a stronger result that DFA are polynomially T/L teachable. It was proved that the model for polynomial identification from characteristic samples and the model for polynomial teachability are equivalent to each other (i.e., by identifying the characteristic set with the teaching sample it was shown that a concept class is polynomially identifiable from characteristic samples iff it is semi-polynomially T/L teachable) [14].

LEMMA 5 *Let $c \in \mathcal{C}$ be a concept with corresponding representation $r \in \mathcal{R}$. If there exists a characteristic sample S_c for c and a polynomial $p_1()$ such that S_c can be computed from r and $|S_c| = p_1(|r|)$ then each example in S_c is simple in the sense that $\forall \alpha \in S_c$, $K(\alpha|r) \leq \mu \lg(|r|)$ where μ is a constant.*

Proof: Fix an ordering of the elements of S_c and define an index to identify the individual elements. Since $|S_c| = p_1(|r|)$ an index that is $\lg(p_1(|r|)) = \mu \lg(|r|)$ bits long is sufficient to uniquely identify each element of S_c . Since S_c can be computed from r we can construct a Turing machine that given r reads as input an index of length $\mu \lg(|r|)$ and outputs the corresponding string of S_c . Thus, $\forall \alpha \in S_c$, $K(\alpha|r) \leq \mu \lg(|r|)$ where μ is a constant independent of α . ■

Let us designate the characteristic set of simple examples S_c identified above to be the set of simple representative examples $S_{sim,rep}^r$ for the concept c represented by r . Lemma 4 and corollary 1 together show that for an integer $l \geq |r|$ and $0 < \delta < 1$ if a sample S of size $|S| = O(l^\mu \lg(\frac{1}{\delta}))$ is drawn at random according to \mathbf{m}_r then with probability greater than $1 - \delta$, $S_{sim,rep}^r \subseteq S$.

THEOREM 5 *Any concept class that is polynomially identifiable from characteristic samples or equivalently semi-polynomially T/L teachable is probably exactly learnable under the PACS model.*

Proof: The proof follows directly from the results of lemma 5, lemma 4, and corollary 1. ■

6.3. Learning from Example Based Queries

A variety of concept classes are known to be learnable in deterministic polynomial time when the learner is allowed access to a teacher (or an oracle) that answers *example based queries* [3]. Example based queries include *equivalence*, *membership*, *subset*, *superset*, *disjointness*, *exhaustive*, *justifying assignment*, and *partial equivalence* queries.

Definition 4. (due to Goldman and Mathias [12])

An example based query is any query of the form

$$\forall (x_1, x_2, \dots, x_k) \in \mathcal{X}^k \text{ does } \phi_r(x_1, x_2, \dots, x_k) = 1?$$

where r is the target concept and k is a constant.

ϕ may use the instances (x_1, \dots, x_k) to compute additional instances on which to perform membership queries. The teacher's response to example based queries is either *yes* or a counter example consisting of $(x_1, x_2, \dots, x_k) \in \mathcal{X}^k$ (along with the correct classification corresponding to each of the x_i 's) for which $\phi_r(x_1, x_2, \dots, x_k) = 0$ and the labeled examples for which membership queries were made in order to evaluate ϕ_r .

THEOREM 6 (due to Goldman and Mathias [12])

Any concept class that is learnable in deterministic polynomial time using example based queries is semi-polynomially T/L teachable.

The above theorem enables us to connect learning from example based queries to PACS learning as follows.

THEOREM 7 *Any concept class that is learnable in deterministic polynomial time using example based queries is also probably exactly learnable under the PACS model*

Proof: Follows directly from theorems 6 and 5. ■

Recently Castro and Guijarro have independently shown that any concept class that is learnable using membership and equivalence queries is also learnable under the PACS model [4]. Further, they have intuitively demonstrated how this result can be extended to all example based queries. Theorem 7 above is an alternate proof of the relationship between query learning and PACS learning. The results presented in this section are more general and enable us to view the PACS learning model in a broader context.

7. Collusion and PACS Learning

Learning models that involve interaction between a knowledgeable teacher and a learner are vulnerable to unnatural *collusion* wherein a the teacher passes information about the representation of the target function as part of the training set [16, 13]. Consider for simplicity that the instance space is $\{0, 1\}^n$ (i.e., the training examples are n bits long). The teacher and learner can *a-priori* agree on some suitable binary encoding of concepts. The teacher can then break the representation of the target concept r in to groups of n bits and use the training set to pass these groups as appropriately labeled examples to the learner. The learner could quickly discover the target concept without even considering the labels of the training examples. The *teaching model* due to Jackson and Tomkins [16] prevents this blatant coding of the target concept by requiring that the learner must still succeed if the teacher is replaced by an adversarial substitute (who does not code the target concept as the teacher above). Further, they argue that in their model the learner can stop only when it is convinced that there is only one concept consistent with the information received from the teacher i.e., the teacher does not tell the learner when to stop. Otherwise learning would be trivialized in that the teacher passes groups of n bits to the learner (as training examples) and when sufficient number of bits have been passed to the learner so as to reconstruct the representation r of the target concept, the teacher tells the learner to stop. Goldman and Mathias' work on *polynomial teachability* [13] shows that an *adversary* whose task it is to embed the training set provided by the teacher (called *teaching set*) into a larger set of correctly labeled examples is sufficient to prevent the type of collusion discussed above. An interesting quirk of the PACS learning model is the fact that the standard encoding of the target concept r is itself a simple example because by definition $K(r|r)$ is low. Thus, r has a high probability under the universal distribution m_r . The PACS learning

scenario wherein examples are drawn at random according to the universal distribution \mathbf{m}_r is similar to the teaching framework in the above teaching models where a knowledgeable teacher draws a helpful set of examples that would enable the learner to learn the target concept quickly. The representation of the target concept r (that is drawn with high probability according to \mathbf{m}_r) could be broken into groups of n bits and passed to the learner as appropriately labeled training examples. Following the argument of Jackson and Tomkins, a learner who is accumulating bits would not know precisely when it has sufficient information to reconstruct the target.

Consider the implication of this issue of collusion on the problem of learning DFA. Note that unlike several concept learning scenarios where the instance space is restricted to fixed length examples in $\{0, 1\}^n$, in DFA learning the individual examples of the teaching set can be of different lengths. In section 5 we presented an algorithm for PAC learning of DFA (under the universal distribution \mathbf{m}_r) that uses the RPNI method. As discussed above, the canonical encoding of the target DFA A (a string r of length $O(N \lg N)$) appears with high probability when the examples are drawn at random according to \mathbf{m}_r . The fact that DFA learning does not require fixed length examples means that r would itself appear with high probability as part of a polynomial sized training set. Of course the learner cannot directly identify which of the example strings represents the target DFA. However, the learner can decode each of the labeled examples. For each example that represents a valid DFA, the learner can test whether the DFA (the decoded example) is consistent with the teaching set and output (say) the first DFA in lexicographic order that is consistent with the training set. With high probability the learner would output the target DFA. This constitutes a PACS algorithm for learning DFA that is computationally more efficient than the RPNI based PACS algorithm presented in this paper. This is another perhaps more subtle form of collusion wherein the learner is provided with an encoding of the target concept but must perform some computation (checking for consistency) in order to suitably identify the target. Goldman and Mathias claim that this is not collusion according to their definition wherein a colluding teacher learner pair is one where the teacher can potentially influence the distribution on the concept class \mathcal{C} when presenting the learner with teaching sets for logically equivalent concepts [13]. Specifically, say the teacher encodes the representation r of the target DFA A as part of the teaching set. The adversary can simply add a correctly labeled example r' which is a representation of a DFA A' that is logically equivalent to A to the teaching set. Both A and A' are consistent with the teaching set and the learner that operates by decoding examples could output either r or r' . While this does not constitute collusion as per their definition the fact is that the strategy of learning by decoding is more attractive as it is computationally more efficient than the RPNI based learning algorithm.

It is clear from the above discussion that the PACS learning framework admits multiple learning algorithms for learning DFA including the seemingly collusive one that learns by decoding labeled examples. A natural question to ask then is whether one can tighten the learning framework suitably so as to avoid any unnatural collusion. Obtaining a satisfactory and general answer to this question would require the development of much more precise definitions of collusion and collusion-free learning algorithms than are currently available. One might argue that sampling of examples according to *helpful* distributions (as done by the PACS framework) by itself constitutes a form of collusion. One might even suggest (perhaps quite unreasonably) that all inductive learning frameworks allow some form of collusion since the training examples have to provide adequate information to identify or approximate the target concept. It is possible that, in general, the definition of collusion-free learning may have to be relative to well-defined restrictions on the specific strategies that can be employed by the learner. A detailed exploration of these issues is beyond the scope of this paper. We restrict ourselves to a few remarks in relation to the DFA learning algorithms discussed in this paper and the collusion strategy that operates by decoding examples and checking for consistency.

In the PACS learning algorithms discussed earlier we have not clearly distinguished between the environment (or the teacher) that provides the labeled examples (drawn according to \mathbf{m}_r) and the learner that uses the training set to produce a hypothesis DFA. It is helpful to keep this distinction in mind for the following discussion. Consider the algorithm \mathcal{A}_2 for probably exact learning of DFA (see Fig. 9). The teacher draws labeled examples according to \mathbf{m}_r . The learning framework allows only labeled examples

of length up to $2N - 1$ to be passed to the learner. Clearly, the encoding of the target DFA which is of length $O(N \lg N)$ cannot appear as a single labeled example in the training set that is restricted to examples of length up to $2N - 1$ only. Note that breaking up the representation of the target DFA into smaller bit strings will not help either because the learner will have to concatenate different strings in different ways in order to determine which concatenation represents a valid DFA. In the worst case this would involve considering all possible permutations of the example strings which is very expensive. Admittedly, it is possible that the teacher and the learner have *a-priori* agreed that the first few strings will contain the necessary representation of the target DFA. Further, since the learner knows N (the number of states of the target) it knows exactly how many bits would be needed to encode the representation of the target. To overcome this scenario it is sufficient to include an adversary (as in the model of Goldman and Mathias). By simply shuffling the set of strings in the teaching set the adversary can cripple a learner that was expecting to receive the target DFA encoded in the first few labeled example strings. The iterative algorithm \mathcal{A}_3 (see Fig. 10) is used when the number of states of the target DFA is not known in advance. In this algorithm, although the training set is restricted to labeled examples of length $2i - 1$ (where i is the current guess of the number of states of the target DFA), the test set T is not thus restricted. The learner could thus potentially receive an encoding of the target concept as a labeled example in the test set. One way to prevent this is to again resort to the adversary. In addition to shuffling the elements of the training set (as discussed above) the adversary will have to take over the task of testing the hypothesis produced by the learner. The learning would proceed as follows: The teacher draws a polynomial (in i and $1/\delta$) sized sample of labeled examples according to \mathbf{m}_r . The adversary takes a subset of the above sample such that each example of the subset is of length at most $2i - 1$, randomly shuffles the subset, and gives it to the learner. The learner uses this training set to produce a hypothesis. The adversary tests the learner's hypothesis for consistency with a polynomial sized test sample drawn according \mathbf{m}_r and informs the learner (without actually revealing the test set) whether the hypothesis was consistent with the test sample or not. The learner decides whether to halt or continue with another iteration of the above steps. This framework will not allow the collusive learning algorithm that operates on decoding examples.

8. Discussion

The problem of exactly learning the target DFA from an arbitrary set of labeled examples and the problem of approximating the target DFA from labeled examples under Valiant's PAC learning framework are both known to be hard problems. Thus, the question as to whether DFA are efficiently learnable under some restricted yet fairly general and practically useful classes of distributions was clearly of interest. In this paper, we have answered this question in the affirmative by providing a framework for efficient PAC learning of DFA from simple examples.

We have demonstrated that the class of simple DFA is polynomially learnable under the universal distribution \mathbf{m} (the simple PAC learning model) and the entire class of DFA is shown to be learnable under the universal distribution \mathbf{m}_r (the PACS learning model). When an upper bound on the number of states of the target DFA is unknown, the algorithm for learning DFA under \mathbf{m}_r can be used iteratively to efficiently PAC learn the concept class of DFAs for any desired error and confidence parameters⁴. These results have an interesting implication on the framework for incremental learning of the target DFA. In the RPNI2 incremental algorithm for learning DFA, the learner maintains a hypothesis that is consistent with all labeled examples seen thus far and modifies it whenever a new inconsistent example is observed [8]. The convergence of this algorithm relies on the fact that sooner or later, the set of labeled examples seen by the learner will include a characteristic set. If in fact the stream of examples provided to the learner is drawn according to a simple distribution, our results show that in an incremental setting the characteristic set would be made available relatively early (during learning) with a sufficiently high probability and hence the algorithm will converge quickly to the desired target. Finally, we have shown the applicability of the PACS learning model in a more general setting by proving that all concept classes that are polynomially identifiable from characteristic samples according to Gold's model, semi-

polynomially T/L teachable according to Goldman and Mathias' model, and learnable in deterministic polynomial time from example based queries are also probably exactly learnable under the PACS model.

The class of simple distributions includes a large variety of probability distributions (including all computable distributions). It has been shown that a concept class is efficiently learnable under the universal distribution if and only if it is efficiently learnable under each simple distribution provided that sampling is done according to the universal distribution [19]. This raises the possibility of using sampling under the universal distribution to learn under all computable distributions. However, the universal distribution is not computable. Whether one can instead get by with a polynomially computable approximation of the universal distribution remains an open question. It is known that the universal distribution for the class of polynomially-time bounded simple distributions is computable in exponential time [19]. This opens up a number of interesting possibilities for learning under simple distributions. In a recent paper Denis and Gilleron have proposed a new model of learning under *helpful distributions* [7]. A helpful distribution is one in which examples belonging to the characteristic set for the concept (if there exists one) are assigned non-zero probability. A systematic characterization of the class of helpful distributions would perhaps give us a more practical framework for learning from simple examples. For instance it might help in adapting the PACS learning model to the incremental learning scenario. A helpful teacher could start out by drawing simple examples based on its knowledge of the target concept but as the learning progresses the teacher could potentially draw examples by combining its knowledge of the target concept with the current hypothesis output by the learner.

A related question of interest has to do with the nature of environments that can be modeled by simple distributions. In particular, if Kolmogorov complexity is an appropriate measure of the intrinsic complexity of objects in nature and if nature (or the teacher) has a propensity for simplicity, then it stands to reason that the examples presented to the learner by the environment are likely to be generated by a simple distribution. Against this background, empirical evaluation of the performance of the proposed algorithms using examples that come from natural domains is clearly of interest.

The issue of collusion that we addressed briefly also opens up several avenues that merit further investigation. Collusion and collusion-free learning need to be defined more precisely. It is of interest to identify whether there are any concept classes that are efficiently learnable by a collusive learning strategy (such as the one that relies on decoding training examples) but are not otherwise efficiently learnable. We have demonstrated one possible modification to the PACS framework for learning DFA to prevent unnatural collusion whereby the target DFA is passed to the learner as a suitably labeled example. It is natural to inquire if the learning framework can be modified similarly for other concept classes such that this type of collusion is prevented.

Some of the negative results in approximate identification of DFA are derived by showing that an efficient algorithm for learning DFA would entail algorithms for solving known hard problems such as *learning boolean formulae* [26] and *breaking the RSA cryptosystem* [17]. It would be interesting to explore the implications of our results on efficient learning of DFA from simple examples on these problems.

Acknowledgments

The authors wish to thank Jack Lutz for introducing them to Kolmogorov complexity, Giora Slutzki for several helpful discussions on automata theory, Pierre Dupont for a careful review of an earlier draft of this paper and several helpful suggestions, and Colin de la Higuera and François Denis for discussions that helped clarify some of the issues related to collusion.

Notes

1. Define $\langle x, y \rangle = bd(x)01y$ where bd is the bit doubling function defined as $bd(0) = 00$, $bd(1) = 11$, and $bd(ax) = aabd(x)$, $a \in \{0, 1\}$.
2. This enumeration strategy applies to any DFA and is not restricted to simple DFA alone.

3. Note that if the sum of the lengths of the examples belonging to a set is k then clearly, the number of examples in that set is at most $k + 1$.
4. Recently it has been shown that if a concept class is learnable under the PACS model using an algorithm that satisfies certain properties then simple concepts of that concept class are learnable under the simple PAC learning model [4]

References

1. D. Angluin. A note on the number of queries needed to identify regular languages. *Information and Control*, 51:76–87, 1981.
2. D. Angluin. Learning regular sets from queries and counterexamples. *Information and Computation*, 75:87–106, 1987.
3. D. Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988.
4. J. Castro and D. Guijarro. Query, pacs and simple-pac learning. Technical Report LSI-98-2-R, Universitat Polyècnica de Catalunya, Spain, 1998.
5. N. Chomsky. Three models for the description of language. *PGIT*, 2(3):113–124, 1956.
6. F. Denis, C. D’Halluin, and R. Gilleron. Pac learning with simple examples. *STACS’96 - Proceedings of the 13th Annual Symposium on the Theoretical Aspects of Computer Science*, pages 231–242, 1996.
7. F. Denis and R. Gilleron. Pac learning under helpful distributions. In *Proceedings of the Eighth International Workshop on Algorithmic Learning Theory (ALT’97), Lecture Notes in Artificial Intelligence 1316*, pages 132–145, Sendai, Japan, 1997.
8. P. Dupont. Incremental regular inference. In L. Miclet and C. Higuera, editors, *Proceedings of the Third ICGI-96, Lecture Notes in Artificial Intelligence 1147*, pages 222–237, Montpellier, France, 1996.
9. P. Dupont. *Utilisation et Apprentissage de Modèles de Langage pour la Reconnaissance de la Parole Continue*. PhD thesis, Ecole Normale Supérieure des Télécommunications, Paris, France, 1996.
10. P. Dupont, L. Miclet, and E. Vidal. What is the search space of the regular inference? In *Proceedings of the Second International Colloquium on Grammatical Inference (ICGI’94)*, pages 25–37, Alicante, Spain, 1994.
11. E. M. Gold. Complexity of automaton identification from given data. *Information and Control*, 37(3):302–320, 1978.
12. S. Goldman and H. Mathias. Teaching a smarter learner. In *Proceedings of the Workshop on Computational Learning Theory (COLT’93)*, pages 67–76. ACM Press, 1993.
13. S. Goldman and H. Mathias. Teaching a smarter learner. *Journal of Computer and System Sciences*, 52:255–267, 1996.
14. Colin de la Higuera. Characteristic sets for polynomial grammatical inference. In L. Miclet and C. Higuera, editors, *Proceedings of the Third ICGI-96, Lecture Notes in Artificial Intelligence 1147*, pages 59–71, Montpellier, France, 1996.
15. J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison Wesley, 1979.
16. J. Jackson and A. Tomkins. A computational model of teaching. In *Proceedings of the Workshop on Computational Learning Theory (COLT’92)*, ACM Press, pages 319–326, 1992.
17. M. Kearns and L. G. Valiant. Cryptographic limitations on learning boolean formulae and finite automata. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, New York, pages 433–444, 1989.
18. K. J. Lang. Random dfa’s can be approximately learned from sparse uniform sample. In *Proceedings of the 5th ACM workshop on Computational Learning Theory*, pages 45–52, 1992.
19. M. Li and P. Vitányi. Learning simple concepts under simple distributions. *SIAM Journal of Computing*, 20:911–935, 1991.
20. M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*, 2^d edition. Springer Verlag, New York, 1997.
21. J. Oncina and P. García. Inferring regular languages in polynomial update time. In N. et al Pérez, editor, *Pattern Recognition and Image Analysis*, World Scientific, pages 49–61, 1992.
22. T. Pao and J. Carr. A solution of the syntactic induction-inference problem for regular languages. *Computer Languages*, 3:53–64, 1978.
23. R. G. Parekh and V. G. Honavar. Efficient learning of regular languages using teacher supplied positive examples and learner generated queries. In *Proceedings of the Fifth UNB Conference on AI*, Fredricton, Canada, pages 195–203, 1993.
24. R. G. Parekh and V. G. Honavar. Learning dfa from simple examples. In *Proceedings of the Eighth International Workshop on Algorithmic Learning Theory (ALT’97), Lecture Notes in Artificial Intelligence 1316*, pages 116–131, Sendai, Japan, 1997. Springer. Also presented at the *Workshop on Grammar Inference, Automata Induction, and Language Acquisition (ICML’97)*, Nashville, TN, July 12, 1997.
25. L. Pitt. Inductive inference, DFAs and computational complexity. In *Analogical and Inductive Inference, Lecture Notes in Artificial Intelligence 397*, Springer-Verlag, pages 18–44, 1989.
26. L. Pitt and M. K. Warmuth. Reductions among prediction problems: on the difficulty of predicting automata. In *Proceedings of the 3rd IEEE Conference on Structure in Complexity Theory*, pages 60–69, 1988.
27. L. Pitt and M. K. Warmuth. The minimum consistency dfa problem cannot be approximated within any polynomial. In *Proceedings of the 21st ACM Symposium on the Theory of Computing*, ACM, pages 421–432, 1989.
28. B. Trakhtenbrot and Ya. Barzdin. *Finite Automata: Behavior and Synthesis*. North Holland Publishing Company, Amsterdam, 1973.
29. L. Valiant. A theory of the learnable. *Communications of the ACM*, 27:1134–1142, 1984.