

Approximate Learning of Limit-Average Automata

Jakub Michaliszyn 

University of Wrocław, Poland
jmi@cs.uni.wroc.pl

Jan Otop 

University of Wrocław, Poland
jotop@cs.uni.wroc.pl

Abstract

Limit-average automata are weighted automata on infinite words that use average to aggregate the weights seen in infinite runs. We study approximate learning problems for limit-average automata in two settings: passive and active. In the passive learning case, we show that limit-average automata are not PAC-learnable as samples must be of exponential-size to provide (with good probability) enough details to learn an automaton. We also show that the problem of finding an automaton that fits a given sample is NP-complete. In the active learning case, we show that limit-average automata can be learned almost-exactly, i.e., we can learn in polynomial time an automaton that is consistent with the target automaton on almost all words. On the other hand, we show that the problem of learning an automaton that approximates the target automaton (with perhaps fewer states) is NP-complete. The abovementioned results are shown for the uniform distribution on words. We briefly discuss learning over different distributions.

2012 ACM Subject Classification Theory of computation → Automata over infinite objects; Theory of computation → Quantitative automata

Keywords and phrases weighted automata, learning, expected value

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2019.17

Related Version A full version of the paper is available at <https://arxiv.org/abs/1906.11104>.

Funding The National Science Centre (NCN), Poland under grant 2017/27/B/ST6/00299.

1 Introduction

Quantitative verification has been proposed to verify non-Boolean system properties such as performance, energy consumption, fault-tolerance, etc. There are two main challenges in applying quantitative verification in practice. First, formalization of quantitative properties is difficult as specifications are given directly as weighted automata, which extend finite automata with weights on transitions [18, 15]. Quantitative logics, which ease specification, have either limited expression power or their model-checking problem is undecidable [11, 13]. Second, there is little research on abstraction in the quantitative setting [14], which would allow us to reduce the size of quantitative models, presented as weighted automata.

We approach both problems using learning of weighted automata. We can apply the learning framework to facilitate writing quantitative specifications and make it more accessible to non-experts. For abstraction purposes, we study approximate learning, having in mind that approximation of a system can significantly reduce its size.

We focus on weighted automata over infinite words which compute the long-run average of weights [15]. Such automata express interesting quantitative properties [15, 23] and admit polynomial-time algorithms for basic decision questions [15]. Some of the interesting properties can be only expressed by non-deterministic automata [15], but every non-deterministic weighted automaton is approximated by some deterministic weighted automaton on almost all words [29]. This means that, by allowing some small margin of error, we can focus on



© Jakub Michaliszyn and Jan Otop;
licensed under Creative Commons License CC-BY

30th International Conference on Concurrency Theory (CONCUR 2019).

Editors: Wan Fokkink and Rob van Glabbeek; Article No. 17; pp. 17:1–17:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

deterministic automata while being able to model important system properties, such as minimal response time, minimal number of errors, the edit distance problem [23], and the *specification repair* framework from [10].

We therefore focus on approximate learning under probabilistic semantics, which corresponds to the average-case analysis in quantitative verification [16]. We treat words as random events and functions defined by weighted automata as random variables. In this setting, an automaton \mathcal{A}' ϵ -approximates \mathcal{A} if the expected difference between \mathcal{A} and \mathcal{A}' (over all words) is bounded by ϵ . We consider two learning approaches: passive and active.

In passive learning, we think of an automaton as a black-box. This might be a program, a specification, a working correct system to be replaced or abstracted, or even only some examples of how the automaton should work. Our goal is to construct a weighted automaton based on this black-box model by observing only inputs and outputs of the model.

In active learning, we assume presence of an interactive *teacher* that reveals the values of given examples and verifies whether a provided automaton is as intended; if not, the teacher responses with a witness, which is a word showing the difference between the constructed automaton and the intended automaton. This does not necessary mean that the teacher is familiar with the weighted automata formalism – to provide a witness, they may simply run the constructed automaton in parallel with the black-box and, if at any point their behaviors differ, provide the appropriate input to the learning algorithm.

Our contributions. We start with a discussion on the setting for learning problems. The first step is to find a suitable representation for samples; not all infinite words have finite representations. A natural idea is to use ultimately periodic words in samples; we study such samples. However, the probability distribution over ultimately periodic words is very different from the uniform distribution over infinite words. Therefore, we also consider samples consisting of a finite word u labeled with the expected value over all extensions of u . The probability distribution over such samples is closer to the uniform distribution over infinite words.

Then we study the passive learning problem. We show that for unique characterization of an automaton, we need a sample of exponential size. We study the complexity of the problem of finding an automaton that fits the whole sample. The problem, without additional restrictions, has a trivial and overfitting solution. To mitigate this, we impose bounds on automata size, and show that then the problem is NP-complete.

For active learning, we show that the problem of learning an almost-exact automaton can be solved in polynomial time, and finding an automaton of bounded size that approximates the target one cannot be done in polynomial time if $P \neq NP$. We conclude with a discussion on different probability distributions.

Related work. The *probably approximately correct* (PAC) learning, introduced in [35], is a general passive learning framework applied to various objects (DNF/CNF formulas, decision trees, automata, etc.) [26]. PAC learning of deterministic finite automata (DFA) has been extensively studied despite negative indicators. First, the *sample fitting* problem for DFA, where the task is to construct a minimal-size DFA consistent with a given sample, has been shown NP-complete [21]. Even approximate sample fitting, where we ask for a DFA at most polynomially greater than a minimal-size DFA, remains NP-complete [34]. Second, it has been shown that existence of a polynomial-time PAC learning algorithm for DFA would break certain cryptographic systems (such as RSA) and hence it is unlikely [25]. Despite these negative results, it has been empirically shown that DFA can be efficiently

learned [27]. In particular, if we assume *structural completeness* of a sample, then it determines a minimal DFA [32]. Pitt posed a question whether DFA are PAC-learnable under the uniform distribution [33], which remains open [2].

Angluin showed that DFA can be learned in polynomial time, if the learning algorithm can ask membership and equivalence queries [1]. This approach proved to be very fruitful and versatile. Angluin's algorithm has been adapted to learn NFA [12], automata over infinite words [3, 20], nominal automata over infinite alphabets [31], weighted automata over words [9] and weighted automata over trees [22, 28].

Recently, there has been a renewed interest in learning weighted automata [8, 28, 6, 5, 7]. These results apply to weighted automata over fields [18], which work over finite words. We, however, consider limit-average automata, which work over infinite words and cannot be defined using a field or even a semiring. Furthermore, we consider weighted (limit-average) automata under probabilistic semantics [16, 29], i.e., we consider functions represented by automata as random variables.

2 The setting

Given a finite alphabet Σ of letters, a *word* w is a finite or infinite sequence of letters. We denote the set of all finite words over Σ by Σ^* , and the set of all infinite words over Σ by Σ^ω . For a word w , we define $w[i]$ as the i -th letter of w , and we define $w[i, j]$ as the subword $w[i]w[i+1]\dots w[j]$ of w . We use the same notation for vectors and sequences; we assume that sequences start with 0 index.

A *deterministic finite automaton* (DFA) is a tuple $(\Sigma, Q, q_0, F, \delta)$ consisting of the alphabet Σ , a finite set of states Q , the initial state $q_0 \in Q$, a set of final states F , and a transition function $\delta: Q \times \Sigma \rightarrow Q$.

A (deterministic) LIMAVG-automaton extends a DFA with a function $C: \delta \rightarrow \mathbb{Q}$ that defines rational *weights* of transitions. The size of such an automaton \mathcal{A} , denoted by $|\mathcal{A}|$, is the sum of the number of states of \mathcal{A} and the lengths of binary encodings of all the weights.

A *run* π of a LIMAVG-automaton \mathcal{A} on a word w is a sequence of states $\pi[0]\pi[1]\dots$ such that $\pi[0]$ is the initial state and for every $i > 0$ we have $\delta(\pi[i-1], w[i]) = \pi[i]$. We do not consider ω -accepting conditions and assume that all infinite runs are accepting. Every run π of \mathcal{A} on an infinite word w defines a sequence of weights $C(\pi)$ of successive transitions of \mathcal{A} , i.e., $C(\pi)[i] = C(\pi[i-1], w[i], \pi[i])$. The *value* of the run π is then defined as $\text{LIMAVG}(\pi) = \limsup_{k \rightarrow \infty} \text{AVG}(C(\pi)[0, k])$, where for finite runs π we have $\text{AVG}(C(\pi)) = \text{SUM}(C(\pi))/|C(\pi)|$. The value of a word w assigned by the automaton \mathcal{A} , denoted by $\mathcal{L}_{\mathcal{A}}(w)$, is the value of the run of \mathcal{A} on w .

We consider three classes of probability measures on words over the alphabet Σ .

- \mathcal{U}^n , for $n \in \mathbb{N}$, is the uniform probability distribution on Σ^n assigning each word the probability $|\Sigma|^{-n}$.
- $\mathcal{G}(\lambda)$, for a termination probability $\lambda \in \mathbb{Q}^+ \cap (0, 1)$, is such that for $u \in \Sigma^*$ we have $\mathcal{G}(\lambda)(u) = |\Sigma|^{-|u|} \cdot (1 - \lambda)^{|u|} \cdot \lambda$. Observe that the probability of words of the same length is equal, and the probability of generating a word of length k is $(1 - \lambda)^k \cdot \lambda$. We can consider this as process generating finite words, which stops after each step with probability λ .
- \mathcal{U}^∞ is the uniform probability measure on Σ^ω . Formally, we define \mathcal{U}^∞ on basic sets $u\Sigma^\omega = \{uw \mid w \in \Sigma^\omega\}$ as follows: $\mathcal{U}^\infty(u\Sigma^\omega) = |\Sigma|^{-|u|}$. Then, \mathcal{U}^∞ is the unique extension on all Borel sets in Σ^ω considered with the product topology [19, 4].

Automata as random variables. A weighted automaton \mathcal{A} defines the measurable function $\mathcal{L}_{\mathcal{A}}(w): \Sigma^{\omega} \mapsto \mathbb{R}$ that assigns values to words. We interpret such functions as random variables w.r.t. the probabilistic measure \mathcal{U}^{∞} . Hence, for a given automaton \mathcal{A} , we consider the following quantities:

1. $\mathbb{E}(\mathcal{A})$ – the expected value of the random variable $\mathcal{L}_{\mathcal{A}}$ defined by \mathcal{A} w.r.t. the uniform distribution \mathcal{U}^{∞} on Σ^{ω} , and
2. $\mathbb{E}(\mathcal{A} \mid u\Sigma^{\omega})$ – the conditional expected value, defined for \mathcal{U}^{∞} as the expected value of \mathcal{L}^u such that $\mathcal{L}^u(w) = \mathcal{L}_{\mathcal{A}}(uw)$.

We consider automata as generators of random variables; automata defining the same random variable are equivalent. Formally, LIMAVG-automata \mathcal{A}_1 and \mathcal{A}_2 are *almost equivalent* if and only if for almost all words w we have $\mathcal{L}_{\mathcal{A}_1}(w) = \mathcal{L}_{\mathcal{A}_2}(w)$. Note that almost all words means all except for words from some Y of probability 0.

LIMAVG-automata considered over probability distributions are equivalent to Markov chains with the long-run average objectives presented in [4, Section 10.5.2].

► **Theorem 1** ([4]). *Let \mathcal{A} be a LIMAVG-automaton. (1) If \mathcal{A} is strongly connected, then almost all words have the same value equal to $\mathbb{E}(\mathcal{A})$. (2) For almost all words w , the run on w eventually reaches some bottom strongly connected component (SCC) of \mathcal{A} .*

Theorem 1 has important consequences. First, we can contract each bottom SCC to a single state with self-loops of the same weight x being the expected value of that SCC. We refer to x as the value of that SCC. Such an operation does not affect almost equivalence. Second, while reasoning about LIMAVG-automata, we can neglect all weights except for the values of bottom SCCs. We will omit weights other than the values of bottom SCCs.

Samples. A sample is a set of *labeled examples* of some (hidden) function $\mathcal{L}: \Sigma^{\omega} \rightarrow \mathbb{R}$. In the classical automata-learning approach words are finite, and hence they can be presented as examples along with the information whether the word belongs to the hidden language. Infinite-words, however, cannot be given directly to a learning algorithm. To mitigate this problem we consider two approaches: one is to restrict examples to ultimately periodic words, which have finite presentation, and the other is to consider finite words and ask for conditional expected values. We discuss both approaches below.

To distinguish samples with different types of labeled examples, we call them U -samples, E -samples and (E, n) -samples.

1. *Ultimately periodic words.* Consider an example being an ultimately periodic word uv^{ω} . It can be presented as a pair of finite words (u, v) and we consider labeled examples (u, v, x) , where $u, v \in \Sigma^*$ and $x = \mathcal{L}(uv^{\omega})$. A set of labeled examples (u, v, x) is called an U -sample. To draw a random U -sample, we consider finite words u, v to be selected independently at random according to distributions $\mathcal{G}(\lambda_1)$ and $\mathcal{G}(\lambda_2)$ for some λ_1, λ_2 . For such a set of pairs of words, we label them according to the function \mathcal{L} .
2. *Conditional expected values.* We consider examples, which are finite words $u \in \Sigma^*$. A labeled example is a pair (u, x) , where $x = \mathbb{E}(\mathcal{L} \mid u\Sigma^{\omega})$ is the conditional expected value of \mathcal{L} under the condition that random words start with u . For such labeled examples we consider E -samples consisting of labeled words of various length, and (E, n) -samples consisting of words of length n . We assume that finite words for random E -samples are drawn according to a distribution $\mathcal{G}(\lambda)$ for some λ , and finite words for random (E, n) -samples are drawn according to the uniform distribution \mathcal{U}^n .

We only consider minimal consistent samples, i.e., samples that do not contain examples whose value can be computed from other examples in the sample. For instance, $\{(a, a, 0), (a, aa, 1)\}$ is an inconsistent U -sample, and $\{(aa, 0), (ab, \frac{1}{2}), (a, 1)\}$ is an inconsistent E -sample over $\{a, b\}$.

► **Remark 2 (Incompatible distributions).** Note that the distribution on ultimately periodic words differ from the uniform distribution on infinite words. The set of ultimately periodic words is countable and hence it has probability 0 (according to the distribution \mathcal{U}^∞). Moreover, almost all infinite words contain all finite words as infixes, whereas this is not the case for ultimately periodic words under any probability distribution.

► **Remark 3 (Feasibility of conditional expectation).** Consider a LIMAVG-automaton \mathcal{A} computing $\mathcal{L}: \Sigma^\omega \rightarrow \mathbb{R}$. For a finite prefix u , we can compute $\mathbb{E}(\mathcal{L} \mid u\Sigma^\omega)$ in polynomial time in $|\mathcal{A}|$ [4]. If we consider \mathcal{A} to be a black-box, which can be controlled, then $\mathbb{E}(\mathcal{L} \mid u\Sigma^\omega)$ can be approximated in the following way. We pick random words v_1, \dots, v_k of length k , compute partial averages in \mathcal{A} of uv_1, \dots, uv_k and then take the average of these values. The probability that this process returns a value ϵ -close to $\mathbb{E}(\mathcal{L} \mid u\Sigma^\omega)$ converges to 1 at exponential rate with k .

3 Passive learning

Passive learning corresponds to a scenario with an uncontrolled working black-box system. The learner can only observe system's output, and its goal is to create an approximate model of the system. This task comprise of two problems. The first problem, *characterization*, is to assess whether the observations cover most, if not all, behaviors of the system. The second one, called *sample fitting*, is to create a reasonable automaton consistent with the observations. In this section we discuss both problems.

3.1 Characterization

A sample can cover only small part of the system. It is sometimes argued [27, 32], however, that if a sample is large enough, then it is likely to cover most, if not all, important behaviors. We show that for some LIMAVG-automata, randomly drawn samples of size less than exponential are unlikely to demonstrate any probable values.

Let $\|S\|$ denote the sum of the lengths of all the examples in a sample S . A sample *distinguishes* two automata if it is consistent with exactly one of them. We show the following.

► **Theorem 4.** *For any n there are two automata $\mathcal{A}_n, \bar{\mathcal{A}}_n$ of size $n+4$ such that for almost all words w we have $|\mathcal{L}_{\mathcal{A}_n}(w) - \mathcal{L}_{\bar{\mathcal{A}}_n}(w)| = 2$, but a random U -sample, E -sample or (E, k) -sample (for any k) S distinguishes \mathcal{A}_n and $\bar{\mathcal{A}}_n$ with the probability at most $\frac{\|S\|}{2^n}$.*

Proof. Consider the alphabet $\{a, b\}$ and $n \in \mathbb{N}$. We construct a LIMAVG-automaton \mathcal{A}_n with $n+4$ states. We use $n+2$ states to find the first occurrence of the infix $a^n b$ in the standard manner. When such an infix is found, the automaton moves to a state q_a if the following letter is a and to a state q_b otherwise. In q_a it loops with the weight 1 and in q_b it loops with the weight -1 . All other weights are 0. So \mathcal{A}_n returns 1 if the first occurrence of $a^n b$ is followed by a , -1 if it is followed by b , and 0 if there is no $a^n b$. The automaton $\bar{\mathcal{A}}_n$ has the same structure as \mathcal{A}_n , but the weights -1 and 1 are swapped.

We first observe that \mathcal{A}_n and $\bar{\mathcal{A}}_n$ differ over almost all infinite words. Indeed, an infinite word with probability 1 contains the infix $a^n b$, and so on almost all words one of the automata returns -1 and the other one 1. Consider a sample S (it can be an E -sample, (E, k) -sample

or U -sample). This sample distinguishes automata \mathcal{A}_n and $\overline{\mathcal{A}}_n$ only if it contains an example with the infix a^nb (in case of U -samples, this means that this infix occurs in w of some example (u, v)) - all other examples for both automata are the same. The probability that S contains a^nb as an infix of one of its examples is bounded by $\frac{\|S\|}{2^{n+1}}$. Indeed, the number of positions in all words is $\|S\|$ and the probability that a^nb occurs on some specific position is at most $\frac{1}{2^{n+1}}$ for all types of samples and any $k > 0$ (if $k < n + 1$, the probability is 0). ◀

Therefore to be able to distinguish just two automata with probability $1 - \epsilon$, we need a sample such that $\frac{\|S\|}{2^{n+1}} > 1 - \epsilon$, which for fixed $\epsilon < 1$ is of exponential size. We show that the exponential upper bound is sufficient.

U -samples. If automata $\mathcal{A}_1, \mathcal{A}_2$ of size n recognize different languages, then there is a word uw^ω such that $\mathcal{L}_{\mathcal{A}_1}(uw^\omega) \neq \mathcal{L}_{\mathcal{A}_2}(uw^\omega)$ and the length of u and w is bounded by n^2 . Assume for simplicity that $\lambda = \lambda_1 = \lambda_2$. If the sample size is at least $|\Sigma|^{2n^2} \cdot \ln \frac{|\Sigma|^{2n^2}}{\epsilon} \cdot (1 - \lambda)^{-2n^2} \cdot \lambda^{-2}$, then with probability $1 - \epsilon$ a random sample contains all such words, and so distinguishes all the automata of size n .

E -samples. E -samples do not distinguish almost-equivalent automata, hence we cannot learn automata exactly. However, exponential samples are enough to learn automata up to almost equivalence. To see that consider two automata \mathcal{A}_1 and \mathcal{A}_2 of size n that are not almost equivalent. Due to Theorem 1 there is a word $u \in \Sigma^*$ such that u reaches bottom SCCs in both \mathcal{A}_1 and \mathcal{A}_2 , and these bottom SCCs have different expected values. Using standard pumping argument, we can reduce the size of u to at most n^2 . So if the sample size is at least $|\Sigma|^{n^2} \cdot \ln \frac{|\Sigma|^{n^2}}{\epsilon} \cdot (1 - \lambda)^{-n^2} \cdot \lambda^{-1}$, then with probability $1 - \epsilon$ it contains all words u of size n^2 , and therefore distinguishes all automata of size n .

For (E, n) -samples, the reasoning is the similar, assuming n is quadratic in the size of the automaton: the sufficient sample size is $|\Sigma|^n \cdot \ln \frac{|\Sigma|^n}{\epsilon}$.

3.2 Consequences for PAC learning


We discuss the consequences of our results to the *probably approximately correct* (PAC) model of learning [35]. In the PAC framework, the learning algorithm should work independently of the probability distribution on samples. However, variants of the PAC framework have been considered where the distribution on samples is uniform [24]. In particular, PAC learning of DFA under the uniform distribution over words is a long-standing open problem [33, 2].

We restrict the classical PAC model and assume that observations are drawn according to the distributions $\mathcal{U}^n, \mathcal{G}(\lambda)$ (as discussed in Section 2) and the quality of the learned automaton is assessed using the uniform distribution over infinite words \mathcal{U}^∞ .

► **Problem 5** (PAC learning under fixed distributions). *Given $\epsilon, \delta \in \mathbb{Q}^+$, $n \in \mathbb{N}$ and an oracle returning random labeled examples consistent with some automaton \mathcal{A}^T of size n , construct an automaton \mathcal{A} such that with probability $1 - \epsilon$ we have $\mathbb{E}(|\mathcal{L}_{\mathcal{A}} - \mathcal{L}_{\mathcal{A}^T}|) < 1 - \delta$.*

As a consequence of Theorem 4, there is no PAC-learning algorithm for LIMAVG-automata with U -samples (resp., E -samples or (E, k) -samples) that uses samples of polynomial size; in particular, there is no such algorithm working in polynomial time.

► **Theorem 6.** *The class of LIMAVG-automata is not PAC-learnable with U -samples, E -samples or (E, k) -samples.*



■ **Figure 1** The canonical automaton \mathcal{A}_φ from the proof of Theorem 8.

3.3 Sample fitting

Once we have a sample, the problem of finding an automaton fitting the sample can be solved in polynomial time in a trivial way: we create an automaton that is a tree such that every word of a given sample leads to a different leaf in this tree, and then we add loops with appropriate values in the leaves (similarly to a prefix tree acceptor [17] for finite automata). This solution leads to an automaton that *overfits* the samples, as it works well only for the sample and it is unlikely to work well with words not included in the sample. Besides, the automaton is linear in the size of the sample, not in the size of the black-box system. For a fixed automaton we can construct arbitrarily large U-samples (or E-samples) consistent with it and hence the gap between the size of such an automaton and the black-box system is arbitrarily large. To exclude such solutions, we restrict the size of the automaton to be constructed. We study the following problem.

► **Problem 7** (Sample fitting). *Given a sample S and $n \in \mathbb{N}$, construct a LIMAVG-automaton with at most n states, which is consistent with S .*

The decision version of this problem only asks whether such an automaton exists. We show that this problem is NP-complete, regardless of the sample representation. For hardness, we reduce the NP-complete problem SAT_0 [17], which is the SAT problem restricted to CNF formulas such that each clause contains only positive literals or only negative literals.

► **Theorem 8.** *The sample fitting problem is NP-complete for U-samples.*


Proof. The membership in NP follows from the following observation: if n is greater than the total length of the samples, then return yes as the tree-like solution works. Otherwise, non-deterministically pick an automaton of the size n and check whether it fits the sample.

The NP-hardness proof is inspired by the construction from [17, Theorem 6.2.1]. For a given instance of the SAT_0 problem $\varphi = \bigwedge_{i=1}^n C_i$ over variables x_1, \dots, x_n (not all variables need to occur in φ), we construct a U-sample S_φ such that there is an automaton with n states fitting S_φ if and only if φ is satisfiable.

We fix the alphabet $\{a_i, c_i, d_i \mid i = 1, \dots, n\} \cup \{b, t\}$. The sample S_φ consists of:

- S1** (c_i, a_j, x) for each $i, j \in \{1, \dots, n\}$, where x is 1 if $i = j$, and 0 otherwise.
- S2** (c_i, d_j, x) for each $i, j \in \{1, \dots, n\}$, where x is 1 if x_i is in C_j , and 0 otherwise.
- S3** $(c_i b, d_i, 1)$ for each $i \in \{1, \dots, n\}$.
- S4** $(c_i b, t, x)$ for each $i \in \{1, \dots, n\}$, where x is 1 if the clause C_i contains only positive literals and 0 if it contains only negative literals.

Assume that φ is satisfiable and let $\sigma : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ be a satisfying valuation. Then, we construct an automaton \mathcal{A}_φ consistent with the sample S_φ starting from the structure presented in Figure 1. Then, we add the following transitions:



■ **Figure 2** A picture of the canonical automaton from the proof of Theorem 9. All the weights are 0 except from transitions from the state q_T .

- for each i , a loop in q_i on the letter t with the value $\sigma(x_i)$,
- for each i, j , a loop in q_i on the letter d_j with the value 1 if x_i is in C_j and 0 otherwise,
- for each clause C_i , if C_i is satisfied because of a variable x_j , then we add a transition from q_i to q_j over b (if there are multiple possible variables, we choose any).

The remaining transitions can be set in arbitrary way. The obtained automaton \mathcal{A}_φ is consistent with the sample S_φ .

Now assume that there is an automaton \mathcal{A} of n states, which is consistent with the sample. We show that the valuation σ such that $\sigma(x_i) = \mathcal{L}_\mathcal{A}(c_i t^\omega)$ satisfies φ . Let q_i be the state where the automaton \mathcal{A} is after reading the word c_i . By **S1**, all the states q_1, \dots, q_n are pairwise different. Since there are only n states, q_1, \dots, q_n are all the states of \mathcal{A} . Now consider any clause C_i . Let q_j be the state of \mathcal{A} after reading $c_i b$. Notice that by **S3**, the value of d_i^ω in q_j is 1, and by **S2**, this means that x_j is in C_i . If C_i contains only positive literals, then the value of t^ω in q_j is 1 by **S4**, which means that $\sigma(x_j) = 1$ and that C_i is satisfied. The other case is symmetric. ◀

► **Theorem 9.** *The sample fitting problem is NP-complete for E -samples.*

Proof. The proof is similar to the proof of Theorem 8. For the NP-hardness, the sample now is obtained from the sample in the proof of Theorem 8 by replacing every triple (u, v, x) by the pair (uv, x) . However, now we ask for an automaton of size $n + 2$. If there is a valuation that satisfies a given set of clauses, then one can construct an automaton based on the one presented in Figure 2.

On the other hand, if there is an automaton fitting the sample, then it has to have a state where the expected value of any word is 0, a state where the expected value of any word is 1, and n different states reachable by each c_1, \dots, c_n . The rest of the proof is virtually the same as in Theorem 8, except that now we define σ such that $\sigma(x_i)$ is the expected value of words with the prefix $c_i t$. ◀

The above proofs also work with some natural relaxations of the sample fitting problem. For example, if we only require the automaton to fit the samples up to some $\epsilon < \frac{1}{2}$, then the proofs still hold since we use only weights 0 and 1. Another relaxation for the E -samples case is to allow the automaton to give wrong values for some samples as long as the summarized probability of the examples with wrong value is less than some ϵ . However, since all the words are of the length at most three, the probability of $u \Sigma^\omega$ for each example u is greater than $\frac{1}{(3n+2)^4}$ (recall that $|\Sigma| = 3n + 2$), which means that for any $\epsilon < \frac{1}{(3n+2)^4}$ for every example some its extension must fit and hence the whole sample must fit.

4 Active learning

In the active case, the learning algorithm can ask queries to an oracle, which is called *the teacher*, which has a (hidden) function $\mathcal{L}: \Sigma^\omega \rightarrow \mathbb{R}$ and answers two types of queries:

- *expectation queries*: given a finite word u , the teacher returns $\mathbb{E}(\mathcal{L} \mid u\Sigma^\omega)$,
- *ϵ -consistency queries*: given an automaton \mathcal{A} , if $\mathcal{L}_{\mathcal{A}}$ ϵ -approximates \mathcal{L} , the teacher answers YES, otherwise it returns a word u such that $|\mathbb{E}(\mathcal{L} \mid u\Sigma^\omega) - \mathbb{E}(\mathcal{L}_{\mathcal{A}} \mid u\Sigma^\omega)| > \epsilon$.

Observe that if $\mathcal{L}_{\mathcal{A}}$ does not ϵ -approximate \mathcal{L} , i.e., $\mathbb{E}(|\mathcal{L} - \mathcal{L}_{\mathcal{A}}|) > \epsilon$, a counterexample for the ϵ -consistency query exists. Indeed, if $\mathcal{L}_1, \mathcal{L}_2$ are defined by LIMAVG-automata and $\mathbb{E}(|\mathcal{L}_1 - \mathcal{L}_2|) > \epsilon$, then there is a word $u\Sigma^\omega$ such that $\mathbb{E}(|\mathcal{L}_1 - \mathcal{L}_2| \mid u\Sigma^\omega) > \epsilon$ and \mathcal{L}_1 (resp., \mathcal{L}_2) returns $\mathbb{E}(\mathcal{L}_1 \mid u\Sigma^\omega)$ (resp., $\mathbb{E}(\mathcal{L}_2 \mid u\Sigma^\omega)$) on almost all words from $u\Sigma^\omega$. Therefore, $|\mathbb{E}(\mathcal{L}_1 \mid u\Sigma^\omega) - \mathbb{E}(\mathcal{L}_2 \mid u\Sigma^\omega)| = \mathbb{E}(|\mathcal{L}_1 - \mathcal{L}_2| \mid u\Sigma^\omega) > \epsilon$.

In the active learning case we consider two problems: approximate learning and rigid approximate learning. We first define approximate learning:

► **Problem 10** (Approximate learning). *Given $\epsilon \in \mathbb{Q}^+ \cup \{0\}$ and a teacher with a (hidden) function \mathcal{L} , construct a LIMAVG-automaton \mathcal{A} such that $\mathcal{L}_{\mathcal{A}}$ ϵ -approximates \mathcal{L} and \mathcal{A} has the minimal number of state among such automata.*

4.1 Approximate learning

We define a decision problem, called approximate minimization, which can be solved in polynomial-time having a polynomial-time approximate learning algorithm.

► **Problem 11** (Approximate minimization). *Given a LIMAVG-automaton \mathcal{A} , $n \in \mathbb{N}$ and $\epsilon \in \mathbb{Q}^+$, the approximate minimization problem asks whether there exists a LIMAVG-automaton \mathcal{A}' with at most n states such that $\mathbb{E}(|\mathcal{L}_{\mathcal{A}} - \mathcal{L}_{\mathcal{A}'}|) \leq \epsilon$.*

An efficient learning algorithm can be used to efficiently compute approximate minimization of a given LIMAVG-automaton \mathcal{A} ; we can run it and compute answers to queries of the learning algorithm in polynomial time in $|\mathcal{A}|$ [4]. We show that the approximate minimization problem is NP-complete, which means that approximate learning cannot be done in polynomial time if $P \neq NP$.

► **Theorem 12.** *The approximate minimization problem is NP-complete.*

Proof sketch. The problem is contained in NP as we can non-deterministically pick an automaton with n states and check whether it ϵ -approximates \mathcal{A} .

For a vector $\vec{v} \in \mathbb{R}^m$ we define $\|\vec{v}\|_1 = \sum_{i=1}^m |\vec{v}[i]|$. For NP-hardness, consider the following problem: *Binary k -Median Problem (BKMP)*: given numbers $n, m, k, t \in \mathbb{N}$ and a set of Boolean vectors $\mathcal{C} = \{\vec{v}_1, \dots, \vec{v}_n\} \subseteq \{0, 1\}^m$, decide whether there are vectors $\vec{u}_1, \dots, \vec{u}_k \in \{0, 1\}^m$ and a partition $\mathcal{D}_1, \dots, \mathcal{D}_k$ of \mathcal{C} such that $\sum_{j=1}^k \sum_{\vec{v} \in \mathcal{D}_j} \|\vec{v} - \vec{u}_j\|_1 \leq t$?

BKMP has been shown NP-complete in [30]. To ease the reduction, we consider a variant of BKMP, called Modified BKMP, where we assume that $\vec{0}, \vec{1}$ belong to the instance and the solution and some additional constraints. The Modified BKMP is also NP-complete.

For $\mathcal{C} = \{\vec{v}_1, \dots, \vec{v}_n\} \subseteq \{0, 1\}^m$ we define $\mathcal{L}_{\mathcal{C}}$ over the alphabet $\Sigma = \{a_1, \dots, a_m\}$ as follows: for all $a_i, a_j \in \Sigma$, $w \in \Sigma^\omega$

- if $i \leq n$, we set $\mathcal{L}_{\mathcal{C}}(a_i a_j w) = \vec{v}_i[j]$.
- if $i \in \{n+1, m\}$, we set $\mathcal{L}_{\mathcal{C}}(a_i a_j w) = \vec{v}_n[j]$.

Intuitively, we select the vector with the first letter and the vector's component with the second letter. This language can be defined with a tree-like LIMAVG-automaton \mathcal{A}_C .

We show that, if an instance (\mathcal{C}, k, t) of Modified BKMP has a solution $\vec{u}_1, \dots, \vec{u}_k$ and $\mathcal{D}_1, \dots, \mathcal{D}_k$, then there exists an automaton \mathcal{A} with $k + 1$ states that $\frac{t}{nm}$ -approximates \mathcal{L}_C . Let $\vec{v}_1 = \vec{u}_1 = \vec{0}$ and $\vec{v}_2 = \vec{u}_2 = \vec{1}$. The automaton \mathcal{A} consists of the initial state q_0 and the successors of the initial state, q_1, \dots, q_k , which correspond to vectors $\vec{u}_1, \dots, \vec{u}_k$, i.e., q_1 is a bottom SCC of the value 0, q_2 is a bottom SCC of the value 1, and for $i > 2$ the successors of q_i encode \vec{u}_i via $\delta(q_i, a_j) = q_1$ if $\vec{u}_i[j] = 0$ and $\delta(q_i, a_j) = q_2$ otherwise. The successors of q_0 are defined based on the partition $\mathcal{D}_1, \dots, \mathcal{D}_k$, i.e., if \vec{v}_i belongs to C_j , then $\delta(q_0, a_i) = q_j$. Observe that \mathcal{A} ϵ -approximates \mathcal{A}_C .

Conversely, consider \mathcal{A}' with $k + 1$ states that $\frac{t}{nm}$ -approximates \mathcal{L}_C . We define vectors $\vec{p}_1, \dots, \vec{p}_m \in \mathbb{R}^m$ such that $\vec{p}_i[j] = \mathbb{E}(\mathcal{A}' \mid a_i a_j \Sigma^\omega)$. The structure of Modified BKMP, implies that the initial state of \mathcal{A}' has no self loops and hence it has at most k different successor states. Therefore, there are at most k different vectors among $\vec{p}_1, \dots, \vec{p}_m$. Finally, we observe that since we consider $\|\cdot\|_1$, w.l.o.g. we can assume that $\vec{p}_1, \dots, \vec{p}_m \in \{0, 1\}^m$. Therefore, these vectors give us a solution to the instance (\mathcal{C}, k, t) of Modified BKMP. \blacktriangleleft

4.2 Rigid approximate learning

One of the drawbacks of the standard approximation is that the counterexamples may be dubious, if not useless. We illustrate this with an example.

► **Example 13** (Dubious counterexamples). Consider a minimal DFA \mathcal{B} with 2^n states whose language consists of words of length m for some $n < m$, and a word $v \in \Sigma^n$. We define a function $\mathcal{L}_{\mathcal{B}, v}: \{a, b\}^\omega \rightarrow \mathbb{R}$ such that for all u, w

- $\mathcal{L}_{\mathcal{B}, v}(auw) = 0$ if $|u| = n$ and \mathcal{B} accepts u ,
- $\mathcal{L}_{\mathcal{B}, v}(auw) = 0.3$ if $|u| = n$ and \mathcal{B} rejects u ,
- $\mathcal{L}_{\mathcal{B}, v}(bw)$ is 0.4 if the first occurrence of v in w is followed by a , and 1 otherwise.


Fix $\epsilon = 0.1$. Observe that $\mathcal{L}_{\mathcal{B}, v}$ can be 0.1-approximated with an automaton \mathcal{A} , which is faithful to $\mathcal{L}_{\mathcal{B}, v}$ on $b\Sigma^\omega$ and returns 0.15 for all other words. \mathcal{A} has $n + O(1)$ states.

Assume that the teacher gives only counterexamples starting with a , and hence ϵ -consistency queries do not give any information about the values of words starting with b . The teacher can do it as long as the algorithm does not know the whole \mathcal{B} , which takes $\Omega(|\mathcal{B}|)$ queries to learn. Yet even if the algorithm learns the whole \mathcal{B} and returns the 0.7 for the words starting with b , the expected difference is $0.5 \cdot 0.3 = 0.15$. It follows that to learn the approximation, the algorithm needs to learn something about v .

Suppose that the algorithm did not learn the whole \mathcal{B} . Then, to learn something non-trivial about words starting with b , it has to ask an expectation query containing v . Since the learning algorithm is deterministic, it asks the same expectation queries for different words v . Therefore, for every learning algorithm there are words v that can be learned only after asking queries of the total length $2^{\Omega(|v|)}$.

It follows that any learning algorithm has to ask queries of total length $\Omega(|\mathcal{B}|)$ or $2^{\Omega(|v|)}$, which totals to $2^{\Omega(n)}$.

In Example 13 we assumed an antagonistic teacher, which misleads the algorithm on purpose. But even with a stochastic teacher, it is not known whether “fixing” a given random counterexample is a step towards better approximation. To resolve this issue, we consider a stronger notion of approximation, called *rigid approximation*, where we require all conditional expected values to be ϵ -close, i.e., for all words $u \in \Sigma^*$ we have $\mathbb{E}(|\mathcal{L}_1 - \mathcal{L}_2| \mid u\Sigma^\omega) \leq \epsilon$. In



■ **Figure 3** The automaton \mathcal{A}_1 approximated by two minimal non-equivalent automata and the automaton \mathcal{A}_{exp} approximated by exponentially many non-equivalent automata.

this framework counterexamples are certain, i.e., if for some $u \in \Sigma^*$, the expectation over $u\Sigma^\omega$ is more than ϵ off the intended value, it has to be modified. Formally, we define the problem as follows:

► **Problem 14** (Rigid approximate learning). *Given $\epsilon \in \mathbb{Q}^+$ and a teacher with a (hidden) function \mathcal{L} , construct an automaton \mathcal{A} such that $\mathcal{L}_{\mathcal{A}}$ is a rigid ϵ -approximation of \mathcal{L} and \mathcal{A} has the minimal number of state among such automata.*

Even though the counterexamples are certain in this framework, we observe that this does not eliminate ambiguity. For instance, there can be multiple automata with the minimal number of states.

► **Example 15** (Non-unique minimalization). Consider the automaton \mathcal{A}_1 depicted in Figure 3 and $\epsilon = \frac{1}{4}$. Any automaton \mathcal{A} , which is a rigid ϵ -approximation of $\mathcal{L}_{\mathcal{A}_1}$, has at least two bottom SCCs and hence it requires at least 3 states. Therefore the automata $\mathcal{A}_2, \mathcal{A}_3$ depicted in Figure 3, which ϵ -approximate $\mathcal{L}_{\mathcal{A}_1}$, have the minimal number of states. This shows that there can be multiple correct answers to in the rigid approximate learning problem.

Based on this example, we construct the automaton \mathcal{A}_{exp} parametrized by $n \in \mathbb{N}$ depicted in Figure 3, which has $O(n)$ states and there are exponentially many (in n) minimal non-equivalent automata, which are rigid $\frac{1}{8n}$ -approximations of $\mathcal{L}_{\mathcal{A}_{exp}}$.


As in the approximate learning case, efficient rigid approximate learning enables us to solve efficiently the following *rigid approximate minimization* problem.

► **Problem 16** (Rigid approximate minimization). *Given a LIMAVG-automaton \mathcal{A} , $n \in \mathbb{N}$ and $\epsilon \in \mathbb{Q}^+$, construct a LIMAVG-automaton \mathcal{A}' with at most n states such that for all words $u \in \Sigma^*$ we have $\mathbb{E}(|\mathcal{L}_{\mathcal{A}} - \mathcal{L}_{\mathcal{A}'}| \mid u\Sigma^\omega) \leq \epsilon$.*

First, we consider a naive approach to solve the rigid approximate minimization problem based on state merging. We start with an input automaton \mathcal{A} , merge its states to maintain the property that the automaton with merged states \mathcal{A}' rigidly ϵ -approximates $\mathcal{L}_{\mathcal{A}}$. We terminate if the automaton is *minimal*, i.e., merging any two states of \mathcal{A}' violates the property. However, it may happen that q_1 can be merged either with q_2 or q_3 , but not with both. We show in the following example that the choice of merged states can have profound impact on the size of a minimal automaton.

► **Example 17** (Minimal automata of different size). Assume $n \in \mathbb{N}$ and the function \mathcal{L} that returns 0 on words from $a\Sigma^n a\Sigma^\omega$, 1 on $b\Sigma^n b\Sigma^\omega$, and 0.5 otherwise. Let \mathcal{A} be a minimal automaton defining such a language, as depicted in Figure 4.

17:12 Approximate Learning of Limit-Average Automata



■ **Figure 4** Minimal automata of different size.

Let $\epsilon = \frac{1}{4}$. To minimize \mathcal{A} , we can merge it to the automaton \mathcal{A}_S with 3 states or to \mathcal{A}_L with $n + 3$ states. Observe that \mathcal{A}_S and \mathcal{A}_L are rigid ϵ -approximations of $\mathcal{L}_{\mathcal{A}}$. For \mathcal{A}_S , it is because for every word au we have $\mathbb{E}(\mathcal{A}_B \mid au\Sigma^\omega) \in \{0, 0.5\}$ and hence it is ϵ -close to 0.25, and similarly for ba . For \mathcal{A}_L , it is because for every u of length at least $n + 2$ the expected value of \mathcal{A}_L is 0.75 or 0.25 depending on the $(n + 2)$ th letter of u whereas for \mathcal{A} it is in $\{0.5, 1\}$ or in $\{0, 0.5\}$, resp. For shorter u , we simply observe that the difference of expected values w.r.t. to a word does not exceed the maximal difference in its suffixes.

The automaton \mathcal{A}_L is minimal, as there are no states that can be merged. Therefore, the difference and even the ratio between the sizes of both automata are unbounded.

We show that the rigid approximate minimization problem is NP-complete, which implies that there is no polynomial-time rigid approximate learning algorithm (unless $P = NP$).

► **Theorem 18.** *The rigid approximate minimization problem is NP-complete.*

Proof sketch. The rigid approximate minimization problem is in NP. We show that it is NP-hard. We define a problem, which is an intermediate step in our reduction.

Given $n, k \in \mathbb{N}$ and vectors $\mathcal{C} = \{\vec{v}_1, \dots, \vec{v}_n\} \subseteq \{0, \frac{1}{2}, 1\}^n$, the $\frac{1}{4}$ -vector cover problem asks whether there exist $\vec{u}_1, \dots, \vec{u}_k \in \mathbb{R}^n$ such that for every $\vec{v} \in \mathcal{C}$ there is \vec{u}_j with $\|\vec{v} - \vec{u}_j\|_\infty \leq \frac{1}{4}$?

The $\frac{1}{4}$ -vector cover problem is NP-complete; the NP-hardness is via reduction from the dominating set problem. We show that the $\frac{1}{4}$ -vector cover problem reduces to the rigid approximate minimization problem

Let $\mathcal{C} = \{\vec{v}_1, \dots, \vec{v}_n\} \subseteq \{0, \frac{1}{2}, 1\}^n$. We define $\mathcal{L}_{\mathcal{C}}$ over the alphabet $\Sigma = \{a_1, \dots, a_m\}$ such that for all $a_i, a_j \in \Sigma$, $w \in \Sigma^\omega$ we have $\mathcal{L}_{\mathcal{C}}(a_i a_j w) = \vec{v}_i[j]$. Such $\mathcal{L}_{\mathcal{C}}$ can be defined with a tree-like LIMAVG-automaton $\mathcal{A}_{\mathcal{C}}$. We show that an instance n, k, \mathcal{C} of the $\frac{1}{4}$ -vector cover problem has a solution if and only if $\mathcal{A}_{\mathcal{C}}$ has a rigid $\frac{1}{4}$ -approximation with $k + 3$ states.

Assume that \mathcal{A}' is an automaton with $k + 3$ states such that for all words $u \in \Sigma^*$ we have $\mathbb{E}(|\mathcal{L}_{\mathcal{A}_{\mathcal{C}}} - \mathcal{L}_{\mathcal{A}'}| \mid u\Sigma^\omega) \leq \frac{1}{4}$. Let p_0 be initial in \mathcal{A}' ; we show that it has at most k different successors because at least 3 states cannot be successors of p_0 : p_0 itself, and p_1, p_2 being states in two different bottom SSCs (that need to exist). Otherwise, if u leads to p_i ($i = 0, 1, 2$) we have $\mathbb{E}(|\mathcal{L}_{\mathcal{C}}(w) - \mathcal{L}_{\mathcal{A}'}(uw)| \mid w \in \Sigma^\omega) > \frac{1}{2}$, which contradicts \mathcal{A}' being $\frac{1}{4}$ -approximation. Thus, p_0 has at most k different successors; We define vectors $\vec{u}_1, \dots, \vec{u}_k$ from the successors of p_0 . Formally, for $i \in \{1, \dots, n\}$ we define a vector \vec{y}_i as $\vec{y}_i[j] = \mathbb{E}(\mathcal{L}_{\mathcal{A}'} \mid a_i a_j \Sigma^\omega)$. There are at most k different vectors among $\vec{y}_1, \dots, \vec{y}_n$ and we take these distinct vectors as $\vec{u}_1, \dots, \vec{u}_k$. The condition $\mathbb{E}(|\mathcal{L}_{\mathcal{A}_{\mathcal{C}}} - \mathcal{L}_{\mathcal{A}'}| \mid u\Sigma^\omega) \leq \frac{1}{4}$ implies that $\vec{u}_1, \dots, \vec{u}_k$ form a solution to n, k, \mathcal{C} .

Conversely, assume that the instance n, k, \mathcal{C} has a solution. Observe that w.l.o.g. we can assume that the solution vectors $\vec{u}_1, \dots, \vec{u}_k$ belong to $\{0.25, 0.75\}^n$. Based on this solution we define \mathcal{A}' with $k + 3$ states $q_0, q_1, \dots, q_k, s_1, s_2$ such that q_0 is initial, q_1, \dots, q_k are successors

of q_0 , and s_1, s_2 are single-state bottom SCC of the values 0.25 and 0.75. The states q_1, \dots, q_k correspond to vectors $\vec{u}_1, \dots, \vec{u}_k$, i.e., the successors of q_i encode \vec{u}_i via $\delta(q_i, a_j) = s_1$ if $\vec{u}_i[j] = 0.25$ and $\delta(q_i, a_j) = s_2$ otherwise. The successors of q_0 are defined based on the matching from the vector cover, i.e., if $\delta(q_0, a_i) = q_j$, then \vec{u}_j is $\frac{1}{4}$ -close to \vec{v}_i . Observe that \mathcal{A}' is a rigid $\frac{1}{4}$ -approximation of $\mathcal{A}_{\mathcal{L}}$. ◀

5 Almost-exact learning

The almost-exact learning the minimal automaton is defined as follows.

► **Problem 19** (Almost-exact learning). *Given a teacher with a (hidden) function \mathcal{L} , construct a LIMAVG-automaton \mathcal{A} such that $\mathcal{L}_{\mathcal{A}}$ is a 0-approximation of \mathcal{L} .*

Notice that for functions \mathcal{L} and $\mathcal{L}_{\mathcal{A}}$, the following conditions are equivalent:

- $\mathcal{L}_{\mathcal{A}}$ is a 0-approximation of \mathcal{L} . ■ $\mathbb{P}(\{w : \mathcal{L}_{\mathcal{A}}(w) \neq \mathcal{L}(w)\}) = 0$.
- $\mathcal{L}_{\mathcal{A}}$ is a rigid 0-approximation of \mathcal{L} . ■ $\mathbb{P}(\{w : \mathcal{L}_{\mathcal{A}}(uw) \neq \mathcal{L}(uw)\}) = 0$ for each u .

We show that there is a polynomial-time algorithm for almost-exact learning.

► **Theorem 20.** *The almost-exact learning problem for LIMAVG-automata can be solved in polynomial time in the size of the minimal automaton that is almost equivalent with the target function \mathcal{L} and the maximal length of counterexamples returned by the teacher.*

Proof. We define a relation $\equiv_{\mathcal{L}}^0$ on finite words $u, v \in \Sigma^*$ as follows:

$$u \equiv_{\mathcal{L}}^0 v \text{ if and only if } \mathbb{P}(\{w : \mathcal{L}(uw) \neq \mathcal{L}(vw)\}) = 0$$

Clearly, $\equiv_{\mathcal{L}}^0$ is an equivalence relation. We show that $\equiv_{\mathcal{L}}^0$ is a right congruence, i.e., if $u \equiv_{\mathcal{L}}^0 v$ and $a \in \Sigma$, then $ua \equiv_{\mathcal{L}}^0 va$. Indeed, consider $X_1 = \{w : \mathcal{L}(uaw) \neq \mathcal{L}(vaw)\}$ and $X_2 = \{w : \mathcal{L}(uw) \neq \mathcal{L}(vw)\}$. Note that $u \equiv_{\mathcal{L}}^0 v$ implies $\mathbb{P}(X_2) = 0$. For all w , if $w \in X_1$, then $aw \in X_2$. Thus, under the uniform distribution we have $\mathbb{P}(X_1) \leq \frac{1}{|\Sigma|} \mathbb{P}(X_2) = 0$ and hence $\mathbb{P}(X_1) = 0$ and $ua \equiv_{\mathcal{L}}^0 va$.

We now show a counterpart of the Myhill–Nerode theorem: there is a LIMAVG-automaton of n states defining almost-exactly \mathcal{L} if and only if the index of $\equiv_{\mathcal{L}}^0$ is n .

If \mathcal{A} defines almost-exactly \mathcal{L} , then the index of $\equiv_{\mathcal{L}}^0$ is bounded by the number of states of \mathcal{A} . Indeed, if for words u, v the automaton \mathcal{A} starting from the initial state ends up in the same state, then for all words w we have $\mathcal{L}(uw) = \mathcal{L}(vw)$ and hence $u \equiv_{\mathcal{L}}^0 v$. Conversely, assume that $\equiv_{\mathcal{L}}^0$ has a finite index. Then, we construct a LIMAVG-automaton $\mathcal{A}_{\equiv_{\mathcal{L}}^0}$ corresponding to $\equiv_{\mathcal{L}}^0$. The states of $\mathcal{A}_{\equiv_{\mathcal{L}}^0}$ are the equivalence classes of $\equiv_{\mathcal{L}}^0$ and $\mathcal{A}_{\equiv_{\mathcal{L}}^0}$ has a transition from $[u]_{\equiv_{\mathcal{L}}^0}$ to $[v]_{\equiv_{\mathcal{L}}^0}$ over a if and only if $[ua]_{\equiv_{\mathcal{L}}^0} = [va]_{\equiv_{\mathcal{L}}^0}$. Observe that due to Theorem 1, if $[u]_{\equiv_{\mathcal{L}}^0}$ and $[v]_{\equiv_{\mathcal{L}}^0}$ are in a bottom SCC, then $[u]_{\equiv_{\mathcal{L}}^0} = [v]_{\equiv_{\mathcal{L}}^0}$. Then, for every bottom SCCs $[u]_{\equiv_{\mathcal{L}}^0}$ we assign the value of all outgoing transitions (which are self-loops) to $\mathbb{E}(\mathcal{L} \mid u\Sigma^\omega)$. For the remaining transitions we set the weights to 0 (due to Theorem 1 these weights are irrelevant as they do not change any of the expected values). Observe that $\mathcal{A}_{\equiv_{\mathcal{L}}^0}$ computes \mathcal{L} .

A classical result of [1] states that DFA can be learned in polynomial time using membership and equivalence queries. We adapt this result here. The learning algorithm for LIMAVG-automata maintains a pair $(\mathcal{Q}, \mathcal{T})$, where \mathcal{Q} , the set of access words, contains different representatives the right congruence relation, and \mathcal{T} , the set of test words, contains words that approximate the right congruence relation. \mathcal{T} defines the relation $\equiv_{\mathcal{T}}$ such that $u_1 \equiv_{\mathcal{T}} u_2$ if and only if for all $v \in \mathcal{T}$ we have $\mathbb{P}(\{w \mid \mathcal{L}(u_1vw) \neq \mathcal{L}(u_2vw)\}) = 0$.

17:14 Approximate Learning of Limit-Average Automata

The algorithm maintains two properties: *separability*: all different words $u_1, u_2 \in \mathcal{Q}$ belong to different equivalence classes of $\equiv_{\mathcal{T}}$, and *closedness*: for all $u_1 \in \mathcal{Q}$ and $a \in \Sigma$ there is $u_2 \in \mathcal{Q}$ with $u_1 a \equiv_{\mathcal{T}} u_2$. Each separable and closed pair $(\mathcal{Q}, \mathcal{T})$ defines a LIMAVG-automaton $\mathcal{A}_{\mathcal{Q}, \mathcal{T}}$, which can be tested for 0-consistency against the teacher's function. If the teacher provides a counterexample u , then it is used to extend \mathcal{Q} and \mathcal{T} . To do so, we split u into $v_1 a, v_2$ such that $v_1 a$ is the minimal prefix of u such that $\mathbb{E}(\mathcal{L} \mid v_1 \Sigma^\omega) \neq \mathbb{E}(\mathcal{A}_{\mathcal{Q}, \mathcal{T}} \mid u \Sigma^\omega)$. Let v'_1 be a word from \mathcal{Q} that is $\equiv_{\mathcal{T}}$ -equivalent to v_1 . We take $\mathcal{Q}' = \mathcal{Q} \cup \{v'_1 a\}$, $\mathcal{T}' = \mathcal{T} \cup \{v_2\}$ and close $(\mathcal{Q}', \mathcal{T}')$ using expectation queries and test the equivalence again. We repeat this until we get a LIMAVG-automaton defining almost equivalent to \mathcal{L} .

Correctness of this algorithm follows from correctness of the algorithm from [1]. ◀

Non-uniform distributions. So far we only discussed the uniform distribution of words. Here we briefly discuss whether Theorem 20 can be generalized to arbitrary distributions, represented by Markov chains. We assume that the Markov chain is given to a learning algorithm in the input.

A (finite-state discrete-time) *Markov chain* is a tuple $\langle \Sigma, S, s_0, E \rangle$, where Σ is the alphabet of letters, S is a finite set of states, s_0 is an initial state, $E: S \times \Sigma \times S \mapsto [0, 1]$ is an edge probability function, which for every $s \in S$ satisfies $\sum_{a \in \Sigma, s' \in S} E(s, a, s') = 1$.

The probability of a finite word u w.r.t. a Markov chain \mathcal{M} , denoted by $\mathbb{P}_{\mathcal{M}}(u)$, is the sum of probabilities of paths from s_0 labeled by u , where the probability of a path is the product of probabilities of its edges. For basic open sets $u \cdot \Sigma^\omega = \{uw \mid w \in \Sigma^\omega\}$, we have $\mathbb{P}_{\mathcal{M}}(u \cdot \Sigma^\omega) = \mathbb{P}_{\mathcal{M}}(u)$, and then the probability measure over infinite words defined by \mathcal{M} is the unique extension of the above measure [19]. We will denote the unique probability measure defined by \mathcal{M} as $\mathbb{P}_{\mathcal{M}}$. The uniform distribution can be expressed with a (single-state) Markov chain and hence all the lower bounds from Section 3 and Section 4 apply here.

A Markov chain \mathcal{M} is *non-vanishing* if for all words $u \in \Sigma^*$ we have $\mathbb{P}_{\mathcal{M}}(u \Sigma^*) > 0$. The almost-exact learning over distributions given by non-vanishing Markov chains and over \mathcal{M} and over the uniform distribution coincide. Thus, the former can be solved using Theorem 20. Indeed, if X is measurable and \mathcal{M} is non-vanishing, then $\mathbb{P}_{\mathcal{M}}(X) > 0$ if and only if $\mathbb{P}(X) > 0$.

The algorithm for the uniform distribution does not extend to vanishing Markov chains because the relation $\equiv_{\mathcal{L}}^0$ is not a right congruence. This cannot be simply fixed as we show that learning cannot be done in polynomial time, assuming $P \neq NP$. We define the almost-exact minimization problem as an instance of the approximate minimization problem with $\epsilon = 0$. Having a polynomial-time algorithm for almost-exact learning, we can solve the almost-exact minimization problem in polynomial time.

► **Theorem 21.** *The almost-exact minimization problem for LIMAVG-automata under distributions given by Markov chains is NP-complete.*

Proof. The problem is in NP: we can non-deterministically pick an automaton \mathcal{A}' and check in polynomial time whether \mathcal{A} and \mathcal{A}' are almost equivalent w.r.t. a given Markov chain.

We reduce the sample fitting problem for E -samples, which is NP-complete (Theorem 9), to the almost-exact minimization problem. Consider an E -sample S based on words u_1, \dots, u_k . Let \mathcal{M}_S be a Markov chain which assigns probability $\frac{1}{k}$ to $u_i \Sigma^\omega$, for each i , and 0 to words not starting with any of u_i . On each $u_i \Sigma^\omega$, \mathcal{M}_S defines the uniform distribution. Let \mathcal{A}_S be a tree-like LIMAVG-automaton consistent with S . Both, \mathcal{M}_S and \mathcal{A}_S are of polynomial size in S . Observe that every automaton \mathcal{A} , which is consistent with S (according to the uniform distribution over infinite words), is almost equivalent to \mathcal{A}_S (over $\mathbb{P}_{\mathcal{M}}$) and vice versa. Therefore, there is an automaton of n states almost equivalent to \mathcal{A}_S (under the distribution $\mathbb{P}_{\mathcal{M}}$) if and only if the sample fitting problem with S and n has a solution. ◀

References

- 1 Dana Angluin. Learning regular sets from queries and counterexamples. *Information and computation*, 75(2):87–106, 1987.
- 2 Dana Angluin and Dongqu Chen. Learning a Random DFA from Uniform Strings and State Information. In *ALT 2015*, pages 119–133, Cham, 2015. Springer International Publishing.
- 3 Dana Angluin and Dana Fisman. Learning regular omega languages. *Theor. Comput. Sci.*, 650:57–72, 2016.
- 4 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 5 Borja Balle and Mehryar Mohri. Learning Weighted Automata. In *CAI 2015*, pages 1–21, 2015.
- 6 Borja Balle and Mehryar Mohri. On the Rademacher Complexity of Weighted Automata. In *ALT 2015*, pages 179–193, 2015.
- 7 Borja Balle and Mehryar Mohri. Generalization bounds for learning weighted automata. *Theor. Comput. Sci.*, 716:89–106, 2018.
- 8 Borja Balle, Prakash Panangaden, and Doina Precup. A Canonical Form for Weighted Automata and Applications to Approximate Minimization. In *LICS 2015*, pages 701–712, 2015.
- 9 Amos Beimel, Francesco Bergadano, Nader Bshouty, Eyal Kushilevitz, and Stefano Varricchio. Learning Functions Represented as Multiplicity Automata. *Journal of the ACM*, 47, October 1999.
- 10 Michael Benedikt, Gabriele Puppis, and Cristian Riveros. Regular Repair of Specifications. In *LICS 2011*, pages 335–344, 2011.
- 11 Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, and Orna Kupferman. Temporal Specifications with Accumulative Values. *ACM Trans. Comput. Log.*, 15(4):27:1–27:25, 2014.
- 12 Benedikt Bollig, Peter Habermehl, Carsten Kern, and Martin Leucker. Angluin-Style Learning of NFA. In *IJCAI 2009*, pages 1004–1009, 2009.
- 13 Patricia Bouyer, Nicolas Markey, and Raj Mohan Matteplackel. Averaging in LTL. In *CONCUR 2014*, pages 266–280, 2014.
- 14 Pavol Cerný, Thomas A. Henzinger, and Arjun Radhakrishna. Quantitative abstraction refinement. In *POPL 2013*, pages 115–128, 2013.
- 15 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM TOCL*, 11(4):23, 2010.
- 16 Krishnendu Chatterjee, Thomas A. Henzinger, and Jan Otop. Quantitative Automata under Probabilistic Semantics. In *LICS 2016*, pages 76–85, 2016.
- 17 Colin de la Higuera. *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press, New York, NY, USA, 2010.
- 18 Manfred Droste, Werner Kuich, and Heiko Vogler. *Handbook of Weighted Automata*. Springer, 1st edition, 2009.
- 19 W. Feller. *An introduction to probability theory and its applications*. Wiley, 1971.
- 20 Dana Fisman. Inferring regular languages and ω -languages. *J. Log. Algebr. Meth. Program.*, 98:27–49, 2018.
- 21 E Mark Gold. Complexity of automaton identification from given data. *Information and control*, 37(3):302–320, 1978.
- 22 Amaury Habrard and José Oncina. Learning Multiplicity Tree Automata. In *ICGI 2006*, pages 268–280, 2006.
- 23 Thomas A. Henzinger and Jan Otop. Model measuring for discrete and hybrid systems. *Nonlinear Analysis: Hybrid Systems*, 23:166–190, 2017.
- 24 Jeffrey C Jackson. An efficient membership-query algorithm for learning DNF with respect to the uniform distribution. *Journal of Computer and System Sciences*, 55(3):414–440, 1997.
- 25 Michael Kearns and Leslie Valiant. Cryptographic limitations on learning Boolean formulae and finite automata. *Journal of the ACM (JACM)*, 41(1):67–95, 1994.

17:16 Approximate Learning of Limit-Average Automata

- 26 Michael J Kearns, Umesh Virkumar Vazirani, and Umesh Vazirani. *An introduction to computational learning theory*. MIT Press, 1994.
- 27 Kevin J. Lang. Random DFA's Can Be Approximately Learned from Sparse Uniform Examples. In *COLT 1992*, pages 45–52, 1992.
- 28 Ines Marusic and James Worrell. Complexity of equivalence and learning for multiplicity tree automata. *Journal of Machine Learning Research*, 16:2465–2500, 2015.
- 29 Jakub Michaliszyn and Jan Otop. Non-deterministic Weighted Automata on Random Words. In *CONCUR 2018*, volume 118 of *LIPICs*, pages 10:1–10:16, 2018.
- 30 Pauli Miettinen, Taneli Mielikäinen, Aristides Gionis, Gautam Das, and Heikki Mannila. The discrete basis problem. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pages 335–346. Springer, 2006.
- 31 Joshua Moerman, Matteo Sammartino, Alexandra Silva, Bartek Klin, and Michal Szynwelski. Learning nominal automata. In *POPL 2017*, pages 613–625, 2017.
- 32 José Oncina and Pedro Garcia. Inferring regular languages in polynomial updated time. In *Pattern recognition and image analysis: selected papers from the IVth Spanish Symposium*, pages 49–61. World Scientific, 1992.
- 33 Leonard Pitt. Inductive inference, DFAs, and computational complexity. In *International Workshop on Analogical and Inductive Inference*, pages 18–44. Springer, 1989.
- 34 Leonard Pitt and Manfred K Warmuth. The minimum consistent DFA problem cannot be approximated within any polynomial. *Journal of the ACM (JACM)*, 40(1):95–142, 1993.
- 35 Leslie G Valiant. A theory of the learnable. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 436–445. ACM, 1984.