

Active Learning of Symbolic Automata Over Rational Numbers

Sebastián Hagedorn¹, Martín Muñoz^{1, 2, 3}, Cristian Riveros^{1, 2}, Rodrigo Toro Icarte^{1, 4}

¹Department of Computer Science, Pontificia Universidad Católica de Chile, Santiago, Chile

²Instituto Milenio Fundamentos de los Datos (IMFD), Santiago, Chile

³Univ. Artois, CNRS, UMR 8188, Centre de Recherche en Informatique de Lens (CRIL), F-62300 Lens, France

⁴Centro Nacional de Inteligencia Artificial (CENIA), Santiago, Chile

Abstract

Automata learning has many applications in artificial intelligence and software engineering. Central to these applications is the L^* algorithm, introduced by Angluin (1987). The L^* algorithm learns deterministic finite-state automata (DFAs) in polynomial time when provided with a minimally adequate teacher. Unfortunately, the L^* algorithm can only learn DFAs over finite alphabets, which limits its applicability. In this paper, we extend L^* to learn symbolic automata whose transitions use predicates over rational numbers, i.e., over infinite and dense alphabets. Our result makes the L^* algorithm applicable to new settings like (real) RGX, and time series. Furthermore, our proposed algorithm for learning each predicate is optimal in the sense that it asks a number of queries to the teacher that is at most linear with respect to the size of their representation.

1 Introduction

Automata learning has powered numerous advances in artificial intelligence and software engineering. Its applications range from interpretable sequence classification (e.g., Shvo et al. 2021; Katzouris and Paliouras 2022; Roy et al. 2023) to reinforcement learning (e.g., Hasanbeig et al. 2021; Corazza, Gavran, and Neider 2022; Toro Icarte et al. 2023). Key to these advances is the L^* algorithm (Angluin 1987).

Learning the smallest automaton consistent with a fixed set of examples is a well-known NP-hard problem (Gold 1967; Angluin 1980). However, Angluin (1987) showed that minimal automata can be learned in polynomial time under an *active learning* framework. In this setting, the learner interacts with a *minimally adequate teacher (MAT)*, which can answer two types of queries: *membership queries*, which return the correct classification of a given string, and *equivalence queries*, which verify whether a proposed automaton correctly represents the target language. If the hypothesis is incorrect, the teacher provides a counterexample. Under these assumptions, the L^* algorithm can efficiently learn a minimal DFA in polynomial time, assuming constant-time responses to both query types.

The L^* algorithm has had a significant and lasting influence on the field of automata learning, shaping both theoretical developments and practical applications (Vaandrager

et al. 2022). However, it is limited to learning *deterministic finite automata (DFA)*, which can only capture temporal behaviors over finite alphabets. This limitation restricts their use in domains involving infinite or dense input spaces, such as complex event recognition (Giatrakos et al. 2020) and time series learning (Hamilton 2020).

To address these limitations, van Noord and Gerdemann (2001) and Veanes et al. (2012) introduced the concept of symbolic automata, which generalize classical automata by allowing transitions to be labeled with predicates over rich alphabet theories, while preserving many of the desirable properties of traditional automata. The learnability and practical applications of symbolic automata have been extensively explored (Drews and D’Antoni 2017; Mens and Maler 2015; Pitt and Warmuth 1993). However, existing approaches for learning symbolic automata assume that the MAT provides lexicographically minimal counterexamples. Without this assumption, it remains an open question whether such automata can be learned efficiently.

In this paper, we extend the L^* algorithm to enable the learning of symbolic automata over the rational numbers. Our approach can learn any symbolic automaton whose predicates operate over a single rational-valued variable, without imposing any constraints on the form of counterexamples provided by the teacher. Furthermore, our method is query-efficient, requiring only a polynomial number of queries relative to the size of the target automaton. To achieve this, we first introduce a technique for learning *finite piecewise functions* over the rational numbers using a MAT. This is accomplished through a search procedure over a number-theoretic structure known as the Stern–Brocot tree. Interestingly, our learning algorithm is efficient in the sense that it needs a linear number of queries concerning the function size and does not depend on the worst-case size counterexample. We then integrate this learning strategy into the L^* framework to support the learning of symbolic automata over rational numbers. Finally, we prove that our method can learn any such automaton in polynomial time under standard assumptions. Given space restrictions, an extended version of this article can be found in (Hagedorn et al. 2025).

2 Preliminaries

We begin by recalling the learning setting that we will use throughout the paper. We present the general definition here

to later instantiate it for finite piecewise functions in Section 3 and for symbolic automata in Section 4.

Learning domain. Let \mathcal{D} be an infinite domain (e.g., sequences) and Σ a finite domain (e.g., $\{0, 1\}$). A *concept* over \mathcal{D} is a function $\gamma : \mathcal{D} \mapsto \Sigma$. A representation τ over \mathcal{D} is a finite object (i.e., given by some syntax and semantics) that encodes a function $\tau : \mathcal{D} \mapsto \Sigma$. We denote by $\text{size}(\tau)$ the size of τ , namely, the number of bits to store τ . We say that τ is a representation of a concept γ , denoted by $\tau = \gamma$, if $\tau(d) = \gamma(d)$ for every $d \in \mathcal{D}$. We usually use \mathcal{C} and \mathcal{R} to denote a *class of concepts* and a *class of representations*, respectively. We assume that every concept $\gamma \in \mathcal{C}$ has some representation in \mathcal{R} , namely, for every $\gamma \in \mathcal{C}$ there exists $\tau \in \mathcal{R}$ such that $\tau = \gamma$. Finally, we denote by $\text{size}(\gamma) = \min_{\tau: \tau = \gamma} \text{size}(\tau)$ the size of the minimum representation in \mathcal{R} of $\gamma \in \mathcal{C}$.

The MAT learning setting. Let \mathcal{C} and \mathcal{R} be a class of concepts and representations, respectively, over \mathcal{D} . In general, given some concept $\gamma \in \mathcal{C}$, the learning task consists of finding some $\tau \in \mathcal{R}$ such that $\tau = \gamma$. For this task, we follow the active learning setting introduced by (Angluin 1987) called *Minimally Adequate Teacher* (MAT): the Learner has to infer the behavior of an unknown concept $\gamma \in \mathcal{C}$ by having access to two functions, called *membership* and *equivalence oracles*, provided by the Teacher who knows γ beforehand. Specifically, a *membership oracle* MQ of γ answers the following query: given an input $d \in \mathcal{D}$, return the output $\gamma(d)$. An *equivalence oracle* EQ of γ receives a representation $\tau \in \mathcal{R}$ and returns either the pair **(true, \perp)**, meaning that $\tau = \gamma$, or the pair **(false, d^*)** where $d^* \in \mathcal{D}$, meaning that $\tau \neq \gamma$ and d^* is a *counterexample* (i.e., $\tau(d^*) \neq \gamma(d^*)$). For any pair of functions $f, g : \mathbb{N} \rightarrow \mathbb{N}$, we say that an algorithm Λ *learns \mathcal{R} using a MAT with $f(n)$ membership queries and $g(n)$ equivalence queries* if, given a membership oracle MQ and an equivalence oracle EQ of some unknown concept $\gamma \in \mathcal{C}$, Λ finds a representation $\tau \in \mathcal{R}$ such that $\tau = \gamma$ after making at most $O(f(n))$ membership queries and $O(g(n))$ equivalence queries with $\text{size}(\gamma) = n$. As it is common in the literature, for measuring the efficiency of Λ , we are interested in bounding the number of calls that Λ makes to the oracles with respect to $\text{size}(\gamma)$ and, also, to the length of the longest counterexample provided by EQ.

3 Learning Finite Piecewise Functions

In this section, we address the problem of active learning finite piecewise functions over the rational numbers. We start by introducing some basic notation regarding rationals and intervals that we will use throughout the paper to then define the learning problem. We end this section by presenting the main technical result of the paper.

Rationals and intervals Let $\mathbb{N} = \{0, 1, \dots\}$ be the set of natural numbers, \mathbb{Z} the set of integers, \mathbb{Q} the set of rational numbers, and $\mathbb{Q}_\infty = \mathbb{Q} \cup \{\infty, -\infty\}$ where $-\infty < q < \infty$ for every $q \in \mathbb{Q}$. We will use a, b, c, \dots to denote elements in \mathbb{N} or \mathbb{Z} , and p, q, r, \dots for elements in \mathbb{Q} . Further, we represent any $q \in \mathbb{Q}_\infty$ as a *fraction* $q = \frac{a}{b}$ where $a \in \mathbb{Z}$ and $b \in \mathbb{N}$. We note that this includes the fractions $\frac{a}{0} = \infty$ and

$\frac{-a}{0} = -\infty$ for any $a \in \mathbb{N}$. We say that $\frac{a}{b}$ is *irreducible* if there is no natural number $c > 1$ such that c divides both a and b . In the sequel, we will always assume that $\frac{a}{b}$ is irreducible, where $\infty = \frac{1}{0}$ and $-\infty = \frac{-1}{0}$. For $q = \frac{a}{b} \in \mathbb{Q}$ with $a \neq 0$, we define $\text{size}(q) = \log(|a|) + \log(b)$ where $\log(\cdot)$ is in base 2, namely, the number of bits to represent q as a fraction; and when $a = 0$, $\text{size}(q) = 1$.

A subset $I \subseteq \mathbb{Q}$ is an *interval* if I is non-empty, and for every $p, q \in I$ and $r \in \mathbb{Q}$, if $p < r < q$, then $r \in I$. As usual, we represent intervals as pairs $\langle p, q \rangle$ where $\langle \in \{[, (, \{, \}$ and $\rangle \in \{], \}, \})$. In addition, when $p = -\infty$ then $\langle = ($, and when $q = \infty$, then $\rangle =)$ (i.e., $-\infty$ or ∞ are never included). Further, we always write these p and q as irreducible fractions in \mathbb{Q}_∞ . For example:

$$\begin{aligned} \left(\frac{-1}{2}, \frac{3}{1}\right] &= \{p \in \mathbb{Q} \mid -\frac{1}{2} < p \leq 3\} \text{ and} \\ \left[\frac{7}{1}, \frac{1}{0}\right) &= \{p \in \mathbb{Q} \mid 7 \leq p < \infty\} \end{aligned}$$

are intervals over \mathbb{Q} under this notation. For $I = \langle p, q \rangle$ we define $\text{size}(I) = \text{size}(p) + \text{size}(q)$.

For two intervals I_1, I_2 , we write $I_1 < I_2$ iff $p < q$ for every $p \in I_1$ and $q \in I_2$. We say that a sequence of intervals $I_1 I_2 \dots I_k$ is an *interval partition* of \mathbb{Q} iff $I_1 < I_2 < \dots < I_k$ and $\mathbb{Q} = \bigcup_{j=1}^k I_j$. For instance,

$$\left(\frac{-1}{0}, \frac{-2}{3}\right) \left[\frac{-2}{3}, \frac{1}{2}\right] \left(\frac{1}{2}, \frac{3}{2}\right) \left(\frac{3}{2}, \frac{1}{0}\right)$$

is an interval partition of \mathbb{Q} .

Finite piecewise functions. Let Σ be a finite set of labels. We will use A, B, \dots for denoting labels in Σ . The class \mathcal{C} of *concepts* are all functions over the rationals $\gamma : \mathbb{Q} \mapsto \Sigma$. Further, we restrict ourselves to the class of *finite piecewise functions*, namely, functions γ that can be represented by a finite number of intervals. Formally, we say that an interval I is *monochromatic* with respect to γ if $\gamma(p) = \gamma(q)$ for every pair $p, q \in I$. If I is monochromatic with respect to γ , we define $\gamma(I) = \gamma(q)$ where q is any element in I . As an example, consider a function $\gamma_1 : \mathbb{Q} \rightarrow \{A, B\}$:

$$\gamma_1(q) = \begin{cases} A, & \text{if } \frac{-2}{3} \leq q \leq \frac{1}{2} \text{ or } \frac{3}{2} < q, \\ B, & \text{otherwise.} \end{cases}$$

One can check that interval $\left(\frac{1}{2}, \frac{3}{2}\right]$ is monochromatic with respect to γ_1 , and that $\gamma_1\left(\left(\frac{1}{2}, \frac{3}{2}\right]\right) = B$.

We say that γ is a *finite piecewise function* if there exists an interval partition $I_1 I_2 \dots I_k$ of \mathbb{Q} such that each I_j is monochromatic with respect to γ . If this is the case, we can represent γ as the sequence of pairs:

$$\tau = (I_1, \gamma(I_1)) (I_2, \gamma(I_2)) \dots (I_k, \gamma(I_k)). \quad (\dagger)$$

We will write $\tau = \gamma$ to denote that τ represents γ and write $\tau(q)$ to denote the label $\gamma(q)$ for every $q \in \mathbb{Q}$. We define the class \mathcal{R} of all *representations* τ like (\dagger) of finite piecewise functions. We define the size of τ as $\text{size}(\tau) = \sum_j \text{size}(I_j)$.

Note that a representation of a finite piecewise function γ is not unique. For example, $\left(\left(\frac{-1}{0}, \frac{1}{1}\right), A\right) \left(\left[\frac{1}{1}, \frac{1}{0}\right), A\right)$ and $\left(\left(\frac{-1}{0}, \frac{1}{0}\right), A\right)$ represent the same concept γ where $\gamma(q) = A$ for every $q \in \mathbb{Q}$. For this reason, we say that a representation like (\dagger) is the *canonical representation* of γ if, in addition,

$\gamma(I_j) \neq \gamma(I_{j+1})$ for every $j < k$. One can easily check that every concept γ has a unique canonical representation. Finally, for any concept γ with canonical representation τ we define $\text{size}(\gamma) = \text{size}(\tau)$. Continuing the example, γ_1 can be seen to be a finite piecewise function. Its canonical representation is shown in Figure 2.

Main technical result. Herein lies the problem of learning an unknown finite piecewise function $\gamma \in \mathcal{C}$ that can be finitely represented by some $\tau \in \mathcal{R}$. In this work, we present an MAT learning algorithm that does this efficiently, focusing on finding each labeled interval $(I, \gamma(I))$ in time $O(\text{size}(I))$. Specifically, we get the following technical result on learning finite piecewise functions.

Theorem 1. *For the class \mathcal{C} of finite piecewise functions, there exists a learning algorithm that learns an unknown concept $\gamma \in \mathcal{C}$ using MAT with a linear number of membership and equivalence queries over $\text{size}(\gamma)$.*

In other words, no matter the concept γ , the number of calls to the oracles is proportional to the size of the canonical representation of γ , where rational numbers are represented in a binary encoding. Arguably, this MAT learning algorithm is asymptotically optimal in the sense that one cannot take less than $\text{size}(\gamma)$. Furthermore, the learning guarantee does not depend on the size of the longest counterexample as other MAT algorithms do (see Theorem 2 in the next section).

We prove Theorem 1 in Sections 5 and 6. Next, we show how to apply this result on learning symbolic automata.

4 Learning Symbolic Automata

Symbolic finite state automata (SFA) are an extension of classical finite state automata in which the transitions are replaced by predicates instead of the symbols from a finite alphabet. They were first introduced in (van Noord and Gerdemann 2001) with natural language processing motivations and have been studied in recent decades because they enable applications over real-life scenarios that usually consider large or even infinite alphabets (Veanes 2013; Argyros and D’Antoni 2018; Drews and D’Antoni 2017).

This paper addresses the learning of SFAs with one-dimensional predicates on \mathbb{Q} , built out of inequalities. In this section, we formally define SFAs with inequality formulas. Then, we present some applications of SFAs in practice. After this, we show how to apply Theorem 1 to MAT learning of SFAs and discuss how it is compared with similar works.

Inequality formulas. Let x be a fixed variable. An *inequality formula* has the following syntax:

$$\varphi := x < q \mid x > q \mid x \leq q \mid x \geq q \mid \varphi \vee \psi \mid \varphi \wedge \psi$$

with $q \in \mathbb{Q}$. We say that a value $p \in \mathbb{Q}$ *satisfies* φ if one can replace every appearance of x in φ by p and obtain a binary formula that evaluates to true. We denote this by $p \models \varphi$. An inequality formula φ naturally defines a function $\llbracket \varphi \rrbracket : \mathbb{Q} \mapsto \{0, 1\}$ (also called a *predicate*) such that $\llbracket \varphi \rrbracket(p) = 1$ if, and only if, $p \models \varphi$. Two inequality formulas φ and ψ are *equivalent* if $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$. As an example, consider:

$$\varphi = \left(\frac{-2}{3} \leq x \wedge x \leq \frac{1}{3} \right) \vee \left(\frac{3}{2} < x \right)$$

which defines $\llbracket \varphi \rrbracket = \left[\frac{-2}{3}, \frac{1}{3} \right] \cup \left(\frac{3}{2}, -\infty \right)$ (similar to γ_1).

We write **true** or **false** as a shorthand for two fix formulas that are always 1 or 0, respectively. Further, for $\sim_1, \sim_2 \in \{<, \leq\}$, we write $q_1 \sim_1 x \sim_2 q_2$ as a shorthand for $q_1 \sim_1 x \wedge x \sim_2 q_2$. One can easily check that for every interval I , there exists a formula φ_I that defines it (e.g., if $I = \left[\frac{-2}{3}, \frac{1}{3} \right]$ then $\varphi_I := \frac{-2}{3} \leq x \leq \frac{1}{3}$). We call the rational numbers that appear in φ the *coefficients* of φ . We define the *size* of φ , denoted by $\text{size}(\varphi)$, as the sum of the sizes of the coefficients and symbols appearing in φ . Finally, we denote the set of all inequality formulas as \mathcal{F} .

Symbolic Automata. We introduce a subset of Symbolic Finite state Automata (SFA) that carry predicates that are represented by inequality formulas. The difference between SFA and classical finite automata is that finite automata have transitions of the form $s_1 \xrightarrow{a} s_2$ with s_1, s_2 states and a a symbol in some finite alphabet; whereas SFA have transitions of the form $s_1 \xrightarrow{\varphi} s_2$ with s_1, s_2 states and φ an inequality formula. Consequently, there is (possibly) an infinite number of values that can satisfy φ . Hence these type of transitions enable finite state automata to handle an infinite alphabet such as the set of rational numbers.

Formally, a (non-deterministic) *Symbolic Finite state Automata with inequality predicates* (SFA) is a tuple:

$$\mathcal{A} = (S, \Delta, s_0, F) \quad (*)$$

where S is a finite set of states, $\Delta \subseteq S \times \mathcal{F} \times S$ is a finite relation, $s_0 \in S$ is the initial state, and $F \subseteq S$ is the set of final states. We say that Δ is the *transition relation* of \mathcal{A} and $(s_1, \varphi, s_2) \in \Delta$ is a *transition* of \mathcal{A} that we usually write as $s_1 \xrightarrow{\varphi} s_2$. A *run* of \mathcal{A} over a string $w = q_1 q_2 \dots q_n \in \mathbb{Q}^*$ is a sequence of transitions:

$$\rho := s_0 \xrightarrow{\varphi_1} s_1 \xrightarrow{\varphi_2} \dots \xrightarrow{\varphi_n} s_n$$

such that $(s_{i-1}, \varphi_i, s_i) \in \Delta$ and $q_i \models \varphi_i$ for every $i \in \{1, \dots, n\}$. A *run* is accepting if $s_n \in F$. The language accepted by \mathcal{A} is defined as $\mathcal{L}(\mathcal{A}) = \{w \in \mathbb{Q}^* \mid \mathcal{A} \text{ accepts } w\}$. We define $\text{size}(\mathcal{A}) = |S| + |\Delta| + \sum_{(s_1, \varphi, s_2) \in \Delta} \text{size}(\varphi)$ as the *size* of \mathcal{A} .

An SFA \mathcal{A} of the form $(*)$ is *deterministic* if, for all pairs of transitions $(s, \varphi, s_1), (s, \psi, s_2) \in \Delta$, if $\llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket \neq \emptyset$, then $\varphi = \psi$ and $s_1 = s_2$. In (D’Antoni and Veanes 2014), it was shown that for every SFA there exists an equivalent deterministic SFA of at most exponential size. As it is standard in the MAT learning setting of finite automata (Angluin 1987; Drews and D’Antoni 2017), in the sequel we assume that all SFAs are deterministic.

In the following, we present two examples of how to use SFAs in practice that motivate their learning problem.

RGX. In practice, regular expressions (RGX) process text documents where symbols are encoded in *Unicode*, where each symbol is an integer in the range from 0 to 1.114.111 written in hexadecimal as U+0000 to U+10FFFF. For instance, the digits $\{0, 1, \dots, 9\}$ run from 48 (U+0030) to 57 (U+0039), the latin alphabet in uppercase $\{A, \dots, Z\}$ runs from 65 (U+0041) to 90 (U+005A), and the latin alphabet in lowercase $\{a, \dots, z\}$ runs from 97 (U+0061) to




Figure 1: Two SFA with inequality formulas. (1) is the SFA of a RGX and (2) is the SFA of a time series pattern.

122 (U+007A). Therefore, a text document in practice can be seen as a sequence of integers $q_1 \dots q_n \in \mathbb{N}^*$ (note that although SFA and our results used \mathbb{Q} , they also applied to \mathbb{N} given that $\mathbb{N} \subseteq \mathbb{Q}$). For a simple example, suppose that we want to check an uppercase letter directly followed by a lowercase letter in the Latin alphabet. This pattern can be defined by a RGX¹ (in PCRE or Perl syntax) as:

$$[A-Z][a-z]$$

where $[A-Z]$ and $[a-z]$ are called *char classes* denoting any symbol from A to Z (i.e. interval $[65, 90]$) and from a to z (i.e. interval $[97, 122]$). In Figure 1 (1), we show how to represent this RGX with an SFA by using inequalities. Note that this pattern can be defined with standard finite state automata; however, we will need a large number of transitions between states to encode the inequality formulas.

Time series. Consider a scenario where a sensor measures the temperature of the environment. More precisely, consider that the sensor emits a signal in the form of a sequence $q_1 q_2 \dots q_n \in \mathbb{Q}^*$ where each $q_i \in \mathbb{Q}$ is a temperature measured in degrees Celsius and where the order represents the time (i.e., q_1 was measured before q_2 , and so on).

Consider that a user is interested in detecting the following pattern of the sensor signals:

$$\frac{13}{2} < t_1 \leq \frac{23}{3} ; t_2 > \frac{13}{1}$$

which means that it wants to see a temperature t_1 in the range $(\frac{13}{2}, \frac{23}{3}]$ followed by a temperature t_2 over 13. We use $;$ to denote that two events occur sequentially, but not necessarily consecutively. One can check that the SFA illustrated in Figure 1 (2) detects the previously defined pattern.

The two previous examples are simplified cases to showcase the use of SFA with inequality formulas in practice. This setting can easily be extended to more complex scenarios, for example, where data symbols can be a tuple of the form (a, q) where a is a label and q is a rational number. For the sake of presentation, we prefer to present SFA and their examples only over rational numbers, but one can easily extend them to other scenarios.

¹Notice that, in practice, RGX are unanchored, namely, they start the evaluation in any part of the document.

Learning SFAs using MAT. The following result² presented in (Argyros and D’Antoni 2018) is crucial for deriving our result regarding MAT learning SFAs.

Theorem 2 (Theorem 1 in (Argyros and D’Antoni 2018)). *Let Λ be an algorithm that learns \mathcal{F} using MAT with $f(n)$ membership queries and $g(n)$ equivalence queries. Then SFA can be learned using MAT with $n^2 \cdot f(n) + n^2 \cdot g(n) \cdot \log(m)$ membership queries and $n^2 \cdot g(n)$ equivalence queries where m is the length of the longest counterexample.*

In other words, a good MAT learning algorithm for the class of inequality formulas \mathcal{F} implies a good MAT learning algorithm for SFA with inequality predicates. One can easily note that inequality formulas are a special case of finite piecewise functions. Indeed, we already noticed that a formula $\varphi \in \mathcal{F}$ defines a function $\llbracket \varphi \rrbracket : \mathbb{Q} \mapsto \Sigma$ where $\Sigma = \{0, 1\}$. Furthermore, $\llbracket \varphi \rrbracket$ is a finite piecewise function and, if $\llbracket \varphi \rrbracket$ has a canonical representation of the form (\dagger) :

$$\varphi^* := \bigvee_{j: \llbracket \varphi \rrbracket(I_j)=1} \varphi_{I_j}$$

where φ_I is the formula representing the interval I . Then, $\text{size}(\varphi^*) \in O(\text{size}(\llbracket \varphi \rrbracket))$, that is, the minimum size of a representation in \mathcal{F} for φ is proportional to the size of the canonical representation of $\llbracket \varphi \rrbracket$. Then, by applying Theorem 1 for learning finite piecewise functions with Theorem 2 for learning SFAs we get the following corollary.

Corollary 1. *An SFA of size n can be learned using MAT with $n^3 \cdot \log(m)$ membership queries and n^3 equivalence queries where m is the length of the longest counterexample.*


It is important to note that previous works like (Drews and D’Antoni 2017; Argyros and D’Antoni 2018; Mens and Maler 2015) settled the basis on MAT learning SFAs (e.g., Theorem 2); however, their results rely on simple predicates (e.g., equality predicates) or strong assumptions over the teacher (e.g., the equivalence query always returns the minimal counterexample). Instead, to the best of our knowledge, Corollary 1 is the first result that shows *fully* MAT learning of SFA over non-trivial predicates. Furthermore, it is efficient in the sense that Theorem 1 is the best that one could achieve for combining it with Theorem 2.

5 The Stern-Brocot Tree and Break Links

In this section, we present the Stern-Brocot tree of rational numbers and some of its properties. Then we introduce the concept of a break link and its connection with convergent fractions, which will be fundamental in supporting the efficient learning algorithm of finite piecewise functions.

The Stern-Brocot tree. The Stern-Brocot tree (T_{SB}) is a construction to represent the set of positive rational numbers presented by (Stern 1858; Brocot 1861). In this paper we consider an extended version of this tree that contains the

²(Argyros and D’Antoni 2018) presents Theorem 2 with more details regarding the exact complexity of its MAT learning algorithm. Here, we oversimplified the upper bound on the size of the representation, which better fits our purposes.



$$\tau_1 = \left(\left(\frac{-1}{0}, \frac{-2}{3} \right), B \right) \left(\left(\frac{-2}{3}, \frac{1}{2} \right), A \right) \left(\left(\frac{1}{2}, \frac{3}{2} \right), B \right) \left(\left(\frac{3}{2}, \frac{1}{0} \right), A \right)$$

Figure 2: Top: The finite piecewise function γ_1 (see Section 3) shown in a Stern–Brocot tree up to depth 4. Nodes q for which $\gamma(q) = A$ (resp. B) are represented by gray (resp. white) blocks. Bottom: The canonical representation of γ_1 .

whole set of rational numbers \mathbb{Q} . The theory for the Stern–Brocot tree is well-studied, so we present here the main definitions that are important for our algorithmic result (see Graham, Knuth, and Patashnik (1994) for further details).

The *Stern–Brocot tree* is a directed rooted tree $T_{\text{SB}} = (V(T_{\text{SB}}), E(T_{\text{SB}}))$ such that T_{SB} is an infinite complete binary tree where $V(T_{\text{SB}}) = \mathbb{Q}$ are the set of vertices (written as irreducible fractions), $E(T_{\text{SB}}) = \{(u, \text{left}(u)), (u, \text{right}(u)) \mid u \in V(T_{\text{SB}})\}$ is the set of edges, and $\frac{0}{1}$ is the root. Figure 2 illustrates the structure of the Stern–Brocot tree up to depth 4. Every vertex u in $V(T_{\text{SB}})$ has exactly two children, the left child $\text{left}(u)$ and the right child $\text{right}(u)$, which are placed from left to right, following the left-to-right order of their children.

The construction of T_{SB} is as follows: we define the *boundary relation* $B \subseteq \mathbb{Q}_{\infty} \times \mathbb{Q} \times \mathbb{Q}_{\infty}$ as the smallest set that satisfies that $(\frac{-1}{0}, \frac{0}{1}, \frac{1}{0}) \in B$, and that if $(\frac{c_1}{d_1}, \frac{a}{b}, \frac{c_2}{d_2}) \in B$, then $(\frac{c_1}{d_1}, \frac{c_1+a}{d_1+b}, \frac{a}{b}) \in B$ and $(\frac{a}{b}, \frac{a+c_2}{b+d_2}, \frac{c_2}{d_2}) \in B$. It can be shown that every irreducible fraction, except $\frac{-1}{0}$ and $\frac{1}{0}$, appears as the middle component of some triple in B exactly once. With the relation B , we define the functions $\text{left}, \text{right} : \mathbb{Q} \rightarrow \mathbb{Q}$ of T_{SB} by:

$$\text{left}\left(\frac{a}{b}\right) = \frac{a+c_1}{b+d_1} \quad \text{and} \quad \text{right}\left(\frac{a}{b}\right) = \frac{a+c_2}{b+d_2},$$

for every $\frac{a}{b} \in \mathbb{Q}$ such that $(\frac{c_1}{d_1}, \frac{a}{b}, \frac{c_2}{d_2}) \in B$. One can check that this relation holds in Figure 2.

Since T_{SB} forms a binary tree and every $\frac{a}{b} \in \mathbb{Q}$ is uniquely represented as a vertex in $V(T_{\text{SB}})$, then there is a unique path from the root $\frac{0}{1}$ to $\frac{a}{b}$. We can encode this path as a string in $\{L, R\}^*$ of the left-child (L) and right-child (R) moves from $\frac{0}{1}$ to $\frac{a}{b}$. For instance, one can check in Figure 2 that LRL represents $\frac{-2}{3}$ and RLL represents $\frac{1}{3}$. In particular, negative numbers starts with L , positive numbers with R , and $\frac{0}{1}$ is represented by the empty string. Furthermore, we can succinctly encode a string in $\{L, R\}^*$ by its *run-length encoding* of the contiguous sequences of R and L symbols. For example, the string $RLLLLRLLRRRL$ (which represents $\frac{23}{108}$) can be grouped as $R^1L^4R^1L^2R^3L^1$. In general,

for every $q \in \mathbb{Q}$ we can define its *SB-encoding*:

$$\text{enc}_{\text{SB}}(q) = \sim[a_1, \dots, a_n]$$

where $\sim \in \{+, -\}$, $n \geq 1$, and $a_i \geq 1$ such that:

- if $\sim = +$ and n is even, then $S = R^{a_1}L^{a_2} \dots R^{a_{n-1}}L^{a_n}$;
- if $\sim = +$ and n is odd, then $S = R^{a_1}L^{a_2} \dots L^{a_{n-1}}R^{a_n}$;
- if $\sim = -$ and n is even, then $S = L^{a_1}R^{a_2} \dots L^{a_{n-1}}R^{a_n}$;
- if $\sim = -$ and n is odd, then $S = L^{a_1}R^{a_2} \dots R^{a_{n-1}}L^{a_n}$;

and S is the path from $\frac{0}{1}$ to q in T_{SB} . In particular, we define the empty array as the SB-encoding of $\frac{0}{1}$. For example, from Figure 2 one can check that $\text{enc}_{\text{SB}}(\frac{-2}{3}) = -[1, 1, 1]$ and $\text{enc}_{\text{SB}}(\frac{1}{3}) = +[1, 2]$. It is worth noticing that the SB-encoding is directly related to the *continued fraction* representation of rational numbers (see Graham, Knuth, and Patashnik (1994))—yet we will present the notion solely as the path of left- and right-child moves in T_{SB} , as this suits our purposes better.

An important property of the SB-encoding of rational numbers is that its size is not much bigger than the irreducible representation of the number.

Proposition 1. *Given any rational number $q = \frac{a}{b}$ and its corresponding SB-encoding $\sim[a_1, \dots, a_n]$, it holds that $\sum_{i=1}^n \log(1 + a_i) \in O(\text{size}(q))$.*

Notice that every array $\sim[a_1, \dots, a_n]$ uniquely represents a rational number in T_{SB} . For every $q \in \mathbb{Q}$ with $\text{enc}_{\text{SB}}(q) = \sim[a_1, \dots, a_n]$ we call the rational number p_m with $\text{enc}_{\text{SB}}(p_m) = \sim[a_1, \dots, a_m]$ (for $1 \leq m \leq n$) the *m-th convergent fraction* of q . In addition, we say that p_m is a strict convergent fraction of q if $m < n$. For instance, $\frac{-1}{1}$ and $\frac{-1}{2}$ are strict convergent fractions of $\frac{-2}{3}$ and $\frac{1}{1}$ is the only strict convergent fraction of $\frac{1}{3}$. One can think of the convergent fractions of a rational number q as the nodes at the turning points in the path to q from the root in T_{SB} . Convergent fractions will be key to our main characterization.

Break links. For learning an unknown finite piecewise function γ with canonical representation τ of the form (\dagger) , the objective is to identify every interval I_j and the corresponding value $\gamma(I_j)$. In the next section, we present a learning algorithm that, using the Stern–Brocot tree, constructs τ by finding the convergent fractions of each endpoint of the intervals in τ . In this subsection, we introduce a notion inside the Stern–Brocot tree, called *break links*, that we will use as a proxy to find all the values that appear in τ .

Formally, given an interval $I = \langle p, q \rangle$, we denote by $\text{bounds}(I) = \{p, q\}$, namely, the lower and upper bounds of I . For a finite piecewise function γ with canonical representation τ like (\dagger) , we define $\text{bounds}(\gamma) = \{\text{bounds}(I_j) \mid j \leq k\} \setminus \{-\infty, \infty\}$. That is, $\text{bounds}(\gamma)$ contains the bounding values inside γ except $-\infty$ and ∞ . Therefore, the main goal in our learning algorithm will be to find all the values in $\text{bounds}(\gamma)$ for γ .

In the following, we denote the *unordered interval* of two rational numbers p, q as: $[p, q] = [\min(p, q), \max(p, q)]$. Let γ be a finite piecewise function and (p, q) be an edge in T_{SB} . We say that (p, q) is a *break link* of γ if p is $\frac{0}{1}$ or $[p, q]$ is not monochromatic with respect to γ , meaning that there

exist $x, y \in [p, q]$, such that $\gamma(x) \neq \gamma(y)$. In Figure 2, we show the Stern-Brocot tree up to depth 4 along with the finite piecewise function γ_1 introduced in Section 3 where nodes q for which $\gamma(q) = A$ (resp. B) are represented by gray (resp. white) blocks. One can check that $(\frac{0}{1}, \frac{1}{1})$ and $(\frac{1}{2}, \frac{2}{3})$ are break links of γ_1 but $(\frac{-1}{1}, \frac{-2}{1})$ and $(\frac{1}{2}, \frac{1}{3})$ are not.

Next, we present the relation between break links and the convergent fractions of the bounds in a piecewise function γ .

Proposition 2. *Given a finite piecewise function γ , (1) for every convergent fraction q of a bound in $\text{bounds}(\gamma)$, there exists a break link (p_1, p_2) of γ such that $q \in \{p_1, p_2\}$; and (2) for every break link (p_1, p_2) of γ at least one among p_1 and p_2 is a convergent fraction of some bound in $\text{bounds}(\gamma)$.*

The previous proposition is crucial for our learning algorithm to find the endpoints in γ . Given that every value in \mathbb{Q} is a convergent fraction of itself, (1) implies that, if we find all break links of γ in T_{SB} , then we will have all the elements in $\text{bounds}(\gamma)$. Furthermore, by (2) and Propositions 1 the number of break links of γ will not be bigger than $\text{size}(\gamma)$. In the next section, we will show how to use this connection to learn any finite piecewise function γ .

6 The Learning Algorithm

The goal of this section is to present the learning algorithm that, given a membership oracle MQ and an equivalence oracle EQ of an unknown finite piecewise function γ , computes its representation τ of the form (\ddagger) . The algorithm works by incrementally adding break links of γ until every break link has been discovered by the algorithm. Then, from the break links we construct the final hypothesis. We will start the section by describing the data structure and an algorithm to build a representation from the collected data. Then, we present the learning algorithm and prove its correctness.

Building a consistent representation. As already mentioned, the main strategy of the algorithm is to find all break links (p, q) of γ . Instead of keeping pairs (p, q) , we will maintain the points p and q in a list D of the form:

$$D = (q_1, A_1), (q_2, A_2), \dots, (q_m, A_m) \quad (\ddagger)$$

where $q_i \in \mathbb{Q}$ and $\gamma(q_i) = A_i$ for every $i \leq m$, and $q_1 < q_2 < \dots < q_m$. Furthermore, every q_i is part of a break link of γ , namely, $q_i \in \{p, q\}$ for some break link (p, q) of γ . We call D the *data structure* of our learning algorithm.

Our first task for the learning algorithm is to show how to construct a representation τ from a given data structure D . Specifically, we say that a representation τ is *consistent* with a data structure D like (\ddagger) if $\tau(q_i) = A_i$ for every $i \leq m$. For a finite piecewise function γ with canonical representation τ like (\ddagger) , we say that $q \in \text{bounds}(\gamma)$ is a *left-bound* (right-bound) of γ if $q \in I_j$ and $q = \inf(I_j)$ (resp. $q = \sup(I_j)$) for some $j \leq k$. In other words, I_j is of the form $[q, p]$ (resp. $\langle p, q]$) for some $p \in \mathbb{Q}_\infty$. Note that every $q \in \text{bounds}(\gamma)$ is either a left- or a right-bound of γ or both (e.g., $I_j = [q, q]$).

So, how do we build a consistent representation τ from D ? We use a strategy to determine which of the values in D are endpoints in $\text{bounds}(\gamma)$ and, additionally, which are left- or right-bounds (or both). Intuitively, if q_{i-1} is a descendant of

Algorithm 1: Construct a representation τ from D

Input: A data structure D like (\ddagger) .

Output: A representation τ consistent with D .

```

1: Procedure ConstructRepresentation( $D$ )
2:    $\tau \leftarrow \emptyset$ 
3:    $\langle \leftarrow (, \text{leftBnd} \leftarrow -\infty, \text{currLabel} \leftarrow A_1$ 
4:   for  $i \in \{1, \dots, m\}$  do
5:     if  $q_{i-1}$  is descendant of  $q_i$  and  $A_{i-1} \neq A_i$  then
6:        $\tau \leftarrow \tau \cup (\langle \text{leftBnd}, q_i), \text{currLabel}$ 
7:        $\langle \leftarrow [, \text{leftBnd} \leftarrow q_i, \text{currLabel} \leftarrow A_i$ 
8:     if  $q_i$  is an ancestor of  $q_{i+1}$  and  $A_i \neq A_{i+1}$  then
9:        $\tau \leftarrow \tau \cup (\langle \text{leftBnd}, q_i), \text{currLabel}$ 
10:       $\langle \leftarrow (, \text{leftBnd} \leftarrow q_i, \text{currLabel} \leftarrow A_{i+1}$ 
11:    $\tau \leftarrow \tau \cup (\langle \text{leftBnd}, \infty), A_m)$ 
12:   return  $\tau$ 

```

q_i in T_{SB} and $\gamma(q_{i-1}) \neq \gamma(q_i)$ for some $i > 1$, then this is evidence that $q_i \in \text{bounds}(\gamma)$ and, furthermore, q_i is a left-bound of γ . Similarly, if q_i is an ancestor of q_{i+1} and $\gamma(q_i) \neq \gamma(q_{i+1})$ for some $i < m$, then q_i is a right-bound of γ . This intuition, in fact, is guaranteed to completely characterize $\text{bounds}(\gamma)$ when D contains all break links of γ .

Proposition 3. *Assume that D like (\ddagger) contains all break links of γ . Then (1) q_i is a left-bound of γ iff q_{i-1} is a descendant of q_i and $A_{i-1} \neq A_i$, and (2) q_i is a right-bound of γ iff q_i is an ancestor of q_{i+1} and $A_i \neq A_{i+1}$.*

In Algorithm 1, we implement this strategy for constructing a representation from D where lines 5-7 encode (1) and lines 8-10 encode (2) plus the border cases (i.e., $-\infty$ and ∞) which we resolve in lines 3 and 11. Notice that before D has all break links, $\text{ConstructRepresentation}(D)$ will not necessarily produce the correct representation for γ . Nevertheless, as the following result shows, it always produces a consistent representation of D .

Theorem 3. *Algorithm 1 outputs a representation τ that is consistent with D . Moreover, if D contains all break links of γ , then $\tau = \gamma$.*

By the previous result, we know that the representation τ is always consistent with D and, moreover, if we aim to find all break links, then eventually we will have that $\tau = \gamma$.

Finding all break links. The learning algorithm is presented in Algorithm 2. It receives as input a membership oracle MQ and an equivalence oracle EQ, both compatible with an unknown finite piecewise function γ in \mathcal{C} . Its goal is to construct a representation τ for γ of the form (\ddagger) after a finite number of iterations. This is done by maintaining a data structure D that stores all the break links of γ that have been found up to the current iteration. Since γ is a finite piecewise function, the interval partition I_1, I_2, \dots, I_k of \mathbb{Q} must be finite and, furthermore, every endpoint q in $\text{bounds}(\gamma)$ has a finite encoding $\text{enc}_{\text{SB}}(q)$. Therefore, by Proposition 2 and Theorem 3, the number of break links of γ must be finite, and, after collecting all of them in D , we can construct a representation τ of γ .

Algorithm 2 works by first adding a pair $(0, \gamma(0))$ into D (line 1), and then, at each iteration, adding the components

Algorithm 2: The Learning Algorithm

Input: A Membership Oracle MQ and an Equivalence Oracle EQ compatible with some finite p.w. function γ .

Output: A representation τ equivalent to γ .

```
1:  $D \leftarrow \{(\frac{0}{1}, \text{MQ}(\frac{0}{1}))\}$ 
2: while true do
3:    $\tau \leftarrow \text{ConstructRepresentation}(D)$ 
4:    $(\text{ans}, q^*) \leftarrow \text{EQ}(\tau)$ 
5:   if ans = true then
6:     return  $\tau$ 
7:    $r \leftarrow \text{FindClosestAncestor}(q^*, D)$ 
8:    $(p, q) \leftarrow \text{FindBreakLink}(r, q^*, D)$ 
9:    $D \leftarrow D \cup \{(p, \text{MQ}(p)), (q, \text{MQ}(q))\}$ 
```

Algorithm 3: Find a break link of γ from a node in T_{SB}

Input: An origin node r , counterexample q^* and array D .

Output: A break link (p, q) of γ not present in D .

```
1: Procedure FindBreakLink( $r, q^*, D$ )
2:   if  $q^* < r$  then
3:      $(p, q) \leftarrow \text{SearchLeft}(r, q^*, D)$ 
4:   else
5:      $(p, q) \leftarrow \text{SearchRight}(r, q^*, D)$ 
6:   return  $(p, q)$ 
```

p and q of a new break link (p, q) of γ into D , along with the labels $\gamma(p)$ and $\gamma(q)$ (line 9). Each iteration does the following: First, it constructs a representation τ using Algorithm 1 and the array D (line 3). Then, it asks the equivalence oracle EQ whether $\tau = \gamma$; if it is true, the algorithm ends (lines 4–6). If not, this means that we are missing some break link in D and, then, in lines 7–8 the algorithm searches the new break link for γ in T_{SB} . Finally, in line 9, it adds the information of the new break link into D , maintaining the array ordered, and avoiding adding any existing tuple again. By Theorem 3, we know that if the equivalence oracle returns a counterexample q^* , then D is missing some break link of γ . Thus, if we prove that Algorithm 2 always adds a new break link in lines 7–9, then its correctness will follow.

The question that remains to answer is where to find a new break link given the counterexample q^* . Towards this goal, Algorithm 2 first computes the closest ancestor of q^* in D (line 7), namely, the element $r \in \{q_1, \dots, q_m\}$ that is the closest ancestor to q^* in T_{SB} assuming that D has form (\ddagger) . One can show that one can compute method FindClosestAncestor(q^*, D) in time $O(|D|)$ and without calling MQ or EQ. To this end, we show that, if D is missing a break link, then one can be found in the left branch or right branch that stems from r , depending on whether $q^* < r$ or $q^* > r$, respectively. The case $q^* = r$ would never happen since D is assumed to be consistent with γ .

Proposition 4. *Assume that $\tau \neq \gamma$ in line 5 of Algorithm 2 and r is the closest ancestor of q^* in D . Then there is a missing break link in the left branch of r when $q^* < r$, or in the right branch when $q^* > r$.*

The method FindBreakLink(r, q^*, D) in Algorithm 3 ex-

ploits Proposition 4 to find a break link. Specifically, it works as follows: given an origin node r , a counterexample q^* and array D , it searches over the left branch or right branch that stem from r using two subroutines searchLeft and searchRight, respectively (see (Hagedorn et al. 2025) for the implementation of both methods). Each subroutine uses galloping search (Bentley and Yao 1976) and membership queries to efficiently find a break link over the branches. If the distance from r to an undiscovered break link of γ is k , then the number of membership queries it takes to find this break link is in $O(\log(k))$.

Theorem 4. *Assume that $\tau \neq \gamma$ in line 5 of Algorithm 2 and r is the closest ancestor of q^* that appears in D . The method FindBreakLink(r, q^*, D) returns an undiscovered break link (p, q) of γ not present in D making at most $O(\log(k))$ membership queries to MQ where k is the distance in T_{SB} between r and p .*

Finally, let us analyze the number of membership and equivalence queries needed to learn an unknown finite piecewise function γ using Algorithm 2. The number of membership queries is the total of queries to MQ made in all calls to FindBreakLink. We see that each $q \in \text{bounds}(\gamma)$ with SB-encoding $\sim[a_1, \dots, a_n]$, by Proposition 2 and Theorem 4, contributes $O(\log(1 + a_i))$ calls to find the break link at a_i . Thus, the total number of membership queries is in $O(\sum_{i=1}^n \log(1 + a_i))$ for q and, by Proposition 1, we find that the number of membership queries is in $O(\text{size}(\gamma))$.

The number of equivalence queries used in Algorithm 2 is bounded by the number of iterations, which is equal to the number of break links for γ . As we reasoned above, this number is in $O(\text{size}(\gamma))$ by Proposition 1.

Theorem 5. *Algorithm 2 returns the canonical representation τ of a finite piecewise function γ using no more than $O(\text{size}(\gamma))$ membership queries and no more than $O(\text{size}(\gamma))$ equivalence queries.*

7 Conclusions and Future Work

We presented a MAT learning algorithm for finite piecewise functions that directly applies to the learning of symbolic automata with interval formulas. Furthermore, this algorithm is efficient in the sense that it takes at most a linear number of queries to the Teacher concerning the size of the concept.

The techniques developed in this paper could have the potential to work in other problems regarding MAT learning. One open problem is to learn a finite piecewise function over \mathbb{Q}^k (i.e., k -dimensions) where a finite number of square regions specify the function. This problem also has applications in symbolic automata, where, for example, time series are sequences of data tuples (e.g., the sensor measures temperature and humidity together). Further, the techniques proposed in this paper could have application in learning timed automata (An et al. 2020), where clock conditions can be seen as a special case of finite piecewise functions.

Acknowledgements

The work of Muñoz and Riveros was supported by ANID – Millennium Science Initiative Program – Code ICN17_002. Riveros was also supported by ANID Fondecyt project 1230935. The work of Toro Icarte was supported by the National Center for Artificial Intelligence CENIA FB210017 (Basal ANID) and Fondecyt project 11230762.

References

- An, J.; Chen, M.; Zhan, B.; Zhan, N.; and Zhang, M. 2020. Learning One-Clock Timed Automata. In Biere, A.; and Parker, D., eds., *Tools and Algorithms for the Construction and Analysis of Systems*, 444–462. Cham: Springer International Publishing. ISBN 978-3-030-45190-5.
- Angluin, D. 1980. Inductive inference of formal languages from positive data. *Information and Control*, 45(2): 117–135.
- Angluin, D. 1987. Learning regular sets from queries and counterexamples. *Information and Computation*, 75(2): 87–106.
- Argyros, G.; and D’Antoni, L. 2018. The Learnability of Symbolic Automata. In Chockler, H.; and Weissenbacher, G., eds., *Computer Aided Verification*, 427–445. Cham: Springer International Publishing. ISBN 978-3-319-96145-3.
- Bentley, J. L.; and Yao, A. C. 1976. An Almost Optimal Algorithm for Unbounded Searching. *Inf. Process. Lett.*, 5(3): 82–87.
- Brocot, A. 1861. Calcul des rouages par approximation, nouvelle méthode, *Revue chronométrique. Journal des horlogers, scientifique et pratique*, 186–194.
- Corazza, J.; Gavran, I.; and Neider, D. 2022. Reinforcement learning with stochastic reward machines. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6429–6436.
- D’Antoni, L.; and Veanes, M. 2014. Minimization of Symbolic Automata. volume 49, 541–553.
- Drews, S.; and D’Antoni, L. 2017. Learning Symbolic Automata. In Legay, A.; and Margaria, T., eds., *Tools and Algorithms for the Construction and Analysis of Systems*, 173–189. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-662-54577-5.
- Gitrakos, N.; Alevizos, E.; Artikis, A.; Deligiannakis, A.; and Garofalakis, M. 2020. Complex event recognition in the big data era: a survey. *The VLDB Journal*, 29(1): 313–352.
- Gold, E. M. 1967. Language identification in the limit. *Information and Control*, 10(5): 447–474.
- Graham, R. L.; Knuth, D. E.; and Patashnik, O. 1994. *Concrete Mathematics: A Foundation for Computer Science*. Reading, MA: Addison-Wesley, second edition. ISBN 0201558025 9780201558029 0201580438 9780201580433 0201142368 9780201142365.
- Hagedorn, S.; Muñoz, M.; Riveros, C.; and Toro Icarte, R. 2025. Active learning of symbolic automata over rational numbers. *CoRR*.
- Hamilton, J. D. 2020. *Time series analysis*. Princeton university press.
- Hasanbeig, M.; Jeppu, N. Y.; Abate, A.; Melham, T.; and Kroening, D. 2021. DeepSynth: Automata synthesis for automatic task segmentation in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 7647–7656.
- Katzouris, N.; and Paliouras, G. 2022. Learning Automata-Based Complex Event Patterns in Answer Set Programming. In *International Conference on Inductive Logic Programming*, 52–68. Springer.
- Mens, I.-E.; and Maler, O. 2015. Learning Regular Languages over Large Ordered Alphabets. *Logical Methods in Computer Science*, Volume 11, Issue 3.
- Pitt, L.; and Warmuth, M. K. 1993. The minimum consistent DFA problem cannot be approximated within any polynomial. *J. ACM*, 40(1): 95–142.
- Roy, R.; Gaglione, J.-R.; Baharisangari, N.; Neider, D.; Xu, Z.; and Topcu, U. 2023. Learning interpretable temporal properties from positive examples only. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 6507–6515.
- Shvo, M.; Li, A. C.; Icarte, R. T.; and McIlraith, S. A. 2021. Interpretable sequence classification via discrete optimization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 9647–9656.
- Stern, M. 1858. Ueber eine zahlentheoretische Funktion. *Journal für die reine und angewandte Mathematik*, 55: 193–220.
- Toro Icarte, R.; Klassen, T. Q.; Valenzano, R.; Castro, M. P.; Waldie, E.; and McIlraith, S. A. 2023. Learning reward machines: A study in partially observable reinforcement learning. *Artificial Intelligence*, 323: 103989.
- Vaandrager, F.; Garhewal, B.; Rot, J.; and Wißmann, T. 2022. A new approach for active automata learning based on apartness. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 223–243. Springer.
- van Noord, G.; and Gerdemann, D. 2001. Finite State Transducers with Predicates and Identities. *Grammars*, 4.
- Veanes, M. 2013. Applications of Symbolic Finite Automata. In Konstantinidis, S., ed., *Implementation and Application of Automata*, 16–23. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-39274-0.
- Veanes, M.; Hooimeijer, P.; Livshits, B.; Molnar, D.; and Bjorner, N. 2012. Symbolic finite state transducers: Algorithms and applications. In *Proceedings of the 39th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, 137–150.