



Model Learning as a Satisfiability Modulo Theories Problem

Rick Smetsers, Paul Fiterău-Broștean^(✉), and Frits Vaandrager

Institute for Computing and Information Sciences,
Radboud University, Nijmegen, The Netherlands
fiteraup@yahoo.com

Abstract. We explore an approach to model learning that is based on using *satisfiability modulo theories* (SMT) solvers. To that end, we explain how DFAs, Mealy machines and register automata, and observations of their behavior can be encoded as logic formulas. An SMT solver is then tasked with finding an assignment for such a formula, from which we can extract an automaton of minimal size. We provide an implementation of this approach which we use to conduct experiments on a series of benchmarks. These experiments address both the scalability of the approach and its performance relative to existing active learning tools.

1 Introduction

We are interested in algorithms that construct black-box state diagram models of software and hardware systems by observing their behavior and performing experiments. Developing such algorithms is a fundamental research problem that has been widely studied. Roughly speaking, two approaches have been pursued in the literature: *passive learning* techniques, where models are constructed from (sets of) runs of the system, and *active learning* techniques, that accomplish their task by actively doing experiments on the system.

Gold [11] showed that the passive learning problem of finding a minimal DFA that is compatible with a finite set of positive and negative examples, is NP-hard. In spite of these hardness results, many DFA identification algorithms have been developed over time, see [13] for an overview. Some of the most successful approaches translate the DFA identification problem to well-known computationally hard problems, such as SAT [12], vertex coloring [10], or SMT [18], and then use existing solvers for those problems.

Angluin [5] presented an efficient algorithm for active learning a regular language L , which assumes a *minimally adequate teacher* (MAT) that answers two types of queries about L . With a *membership query*, the algorithm asks whether or not a given word w is in L , and with an *equivalence query* it asks whether or

This work is supported by the Netherlands Organization for Scientific Research (NWO) projects 628.001.009 on Learning Extended State Machine for Malware Analysis (LEMMA), and 612.001.216 on Active Learning of Security Protocols (ALSeP).

This paper greatly extends earlier work [24].

not the language L_H of an hypothesized DFA H is equal to L . If L_H and L are different, a word in the symmetric difference of the two languages is returned. Angluin’s algorithm has been successfully adapted for learning models of real-world software and hardware systems [19, 21, 25], as shown in Fig. 1. A membership query (MQ) is implemented by bringing the *system under learning* (SUL) in its initial state and the observing the outputs generated in response to a given input sequence, and an equivalence query (EQ) is approximated using a *conformance testing tool* (CT) [17] via a finite number of *test queries* (TQ).




Fig. 1. Model learning within the MAT framework.

Walkinshaw et al. [26] observed that from each passive learning algorithm one can trivially construct an active learning algorithm that only poses equivalence queries. Starting from the empty set of examples, the passive algorithm constructs a first hypothesis H_1 that is forwarded to the conformance tester. The first counterexample w_1 of the conformance tester is then used to construct a second hypothesis H_2 . Next counterexamples w_1 and w_2 are used to construct hypothesis H_3 , and so on, until no more counterexamples are found.

In this article, we compare the performance of existing active learning algorithms with passive learning algorithms that are ‘activated’ via the trick of Walkinshaw et al. [26]. At first, this may sound like a crazy thing to do: why would one compare an efficient active learning algorithm, polynomial in the size of the unknown state machine, with an algorithm that makes a possibly super-polynomial number of calls [6] to a solver for an NP-hard problem? The main reason is that in practical applications i/o interactions often take a significant amount of time. In [22], for instance, a case study of an interventional X-ray system is described in which a single i/o interaction may take several seconds. Therefore, the main bottleneck in these applications is the total number of membership and test queries, rather than the time required to decide which queries to perform. Also, in practical applications the state machines are often small, with at most a few dozen states (see for instance [3, 4, 22]). Therefore, even though passive learning algorithms do not scale well, there is hope that they can still

handle these applications. Active learning algorithms rely on asking a large number of membership queries to construct hypotheses. Passive learning algorithms pose no membership queries, but instead need a larger number of equivalence queries, which are then approximated using test queries. A priori, it is not clear which approach performs best in terms of the total number of membership and test queries needed to learn a model.

Our experiments compare the original L^* [5] and the state-of-the-art TTT [15] active learning algorithm with an SMT-based passive learning algorithm on a number of practical benchmarks. We encode the question whether there exists a state machine with n states that is consistent with a set of observations into a logic formula, and then use the Z3 SMT solver [9] to decide whether this formula is satisfiable. By iteratively incrementing the number of states we can find a minimal state machine consistent with the observations. As equivalence oracle we use a state-of-the-art conformance testing algorithm based on adaptive distinguishing sequences [16, 23]. In line with our expectations, the passive learning approach is competitive with the active learning algorithms in terms of the number of membership and test queries needed for learning.

An advantage of SMT encodings, when compared for instance with encodings based on SAT or vertex coloring, is the expressivity of the underlying logic. In recent years, much progress has been made in extending active learning algorithms to richer classes of models, such as register automata [2, 8, 14] in which data may be tested and stored in registers. We show that the problem of finding a register automaton that is consistent with a set of observations can be expressed as an SMT problem, and compare the performance of the resulting learning algorithm with that of Tomte [2], a tool for active learning of register automata, on some simple benchmarks. New algorithms for active learning of FSMs, Mealy machines and various types of register automata are often extremely complex, and building tools implementations often takes years [2, 8, 15]. Adapting these tools to slightly different scenarios is typically a nightmare. One such scenario is when the system is missing *reset* functionality. This renders most active learning tools impractical, as these rely on the ability to reset the system. Developing SMT-based learning algorithms for register automata in settings with and without resets only took us a few weeks. This shows that the SMT-approach can be quite effective as a means for prototyping learning algorithms in various settings.

The rest of this paper is structured as follows. Section 2 describes how one can encode the problem of learning a minimal consistent automaton in SMT. The scalability and effectiveness of our approach, and its applicability in practice are assessed in Sect. 3. Conclusions are presented in Sect. 4.

2 Model Learning as an SMT Problem

This section describes how to express the existence of an automaton A of at most n states that is consistent with a set of observations S in a logic formula. *If and only if* there exists an assignment to the variables of this formula that makes it true, then exists an automaton A with at most n states that is consistent with S .

We use an SMT solver to find such an assignment. If the SMT solver concludes that the formula is satisfiable, then its solution provides us with A .

We distinguish three types of *constraints*:

- *axioms* must be satisfied for A to behave as intended by its definition.
- *observation constraints* must be satisfied for A to be consistent with S .
- *size constraints* must be satisfied for A to have n states or less.

Hence, the problem can be solved by iteratively incrementing n until the encoding of the axioms, observation constraints and size constraints is satisfiable.

In the following subsections, we present encodings for deterministic finite automata (Sect. 2.1), register automata (Sect. 2.2), and input-output register automata (Sect. 2.3). Not included here are encodings for Mealy and Moore machines. These are available in the extended version of this work¹.

2.1 An Encoding for Deterministic Finite Automata

A *deterministic finite automaton* (DFA) accepts and rejects *strings*, which are sequences of labels. We define a DFA as follows.

Definition 1. A DFA is a tuple (L, Q, q_0, δ, F) comprising a finite, non-empty set of labels L , a finite non-empty set of states Q , an initial state $q_0 \in Q$, a transition function $\delta : Q \times L \rightarrow Q$ and a set of accepting states $F \subseteq Q$.

Let x be a string. We use x_i to denote i th label of x . We use $x_{[i,j]}$ to denote the substring of x starting at position i and ending at position j (inclusive), i.e. $x = x_{[1,|x|]}$. A DFA accepts a string x if its computation ends in an accepting state, or more formally, if $\delta(q_0, x) \in F$ where δ is extended to strings.

Let S_+ be a set of strings that should be accepted, and let S_- be a disjoint set of strings that should be rejected. Let S be the set that contains all of these strings, along with their labels, i.e. $S = \{(x, \text{true}) : x \in S_+\} \cup \{(x, \text{false}) : x \in S_-\}$. A DFA is *consistent* with S if it accepts all strings in S_+ , and rejects all strings in S_- .

In our DFA encoding, Q is a finite subset of the (non-negative) natural numbers \mathbb{N} with $q_0 = 0$, while F is encoded as a function $\lambda : Q \rightarrow \mathbb{B}$, such that $q \in F \iff \lambda(q) = \text{true}$.

Similarly to Heule and Verwer [12] and Bruynooghe et al. [7], we arrange elements of S in an *observation tree* (OT) with the following definition:

Definition 2. An OT for $S = \{S_+, S_-\}$ is a tuple (L, Q, λ) , where L is a set of labels, $Q = \{x \in L^* : x \text{ is a prefix of a string in } S_+ \cup S_-\}$, $\lambda : S_+ \cup S_- \rightarrow \mathbb{B}$ is a output function for the strings, with $x \in S_+ \iff \lambda(x) = \text{true}$.

An OT allows us to provide efficient encodings for the labeled strings in S . Suppose an OT $T = (L, Q^T, \lambda^T)$ for a given set $S = \{S_+, S_-\}$. To find a DFA

¹ <https://gitlab.science.ru.nl/rick/z3gi/blob/lata/resources/paper.pdf>.

$A = (L, Q^A, q_0, \delta^A, F)$ consistent with S , we define a surjective (i.e. many-to-one) function $map : Q^T \rightarrow Q^A$ encoding correspondence between prefixes in the OT and states in A . We then introduce the following observation constraints:

$$map(\epsilon) = q_0 \quad (1)$$

$$\forall xl \in Q^T : x \in L^*, l \in L \quad \delta^A(map(x), l) = map(xl) \quad (2)$$

$$\forall x \in S_+ \cup S_- \quad \lambda^A(map(x)) = \lambda^T(x) \quad (3)$$

Equation 1 maps the empty string to the initial state of A . Equation 2 encodes the observed prefixes as transitions of A while Eq. 3 encodes the observed outputs, with λ^A encoding F .

Lastly, we limit A to at most n states by the following size constraint:

$$\forall q \in \{0, \dots, n-1\} \quad \forall l \in L \quad \bigvee_{q'=0}^{n-1} \delta(q, l) = q' \quad (4)$$

2.2 An Encoding for Register Automata

A *register automaton* (RA) can be seen as an automaton that is extended with a set of *registers* that can store data parameters. The values in these registers can then be used to express conditions over the transitions of the automaton, or *guards*. If the guard is satisfied the transition is fired, possibly storing the provided data parameter (this is called an *assignment*) and bringing the automaton from the current *location* to the next. As such, an RA can be used to accept or reject sequences of label-value pairs. Unlike finite automata, “states” in a register automaton are called locations because the *state* of the automaton also comprises the values of the registers. Therefore, an infinite number of possible states can be modeled using a small number of locations and registers.

The RAs we define have the following restrictions. Transitions in an RA do not imply (dis)equality of distinct registers (*right invariance*), values cannot be moved from one register to another (*non-swapping*) and registers always store unique values (*unique-valued*). The first two restrictions help simplify the encoding while the third is necessary to avoid the non-determinism caused by two used registers holding the same value. RAs with these restrictions can require more locations to be consistent with a given set of action strings, than those without. However, they are equally expressive. For a formal treatment of these restrictions and their implications, we refer to [2].

We define an RA as follows.

Definition 3. *An RA is a tuple $(L, R, Q, q_0, \delta, \lambda, \tau, \pi)$. Therein, L , Q , q_0 and λ are a set of labels, a set of locations, the start location, and a location output function respectively. R is a finite set of registers. $\delta : Q \times L \times (R \cup \{r_\perp\}) \rightarrow Q$ is a register transition function. $\tau : Q \times R \rightarrow \mathbb{B}$ is a register use predicate, and $\pi : Q \times L \rightarrow (R \cup \{r_\perp\})$ is a register update function.*

We call a label-value pair an *action* and denote it $l(v)$ for input label l and parameter v . We assume w.l.o.g. that parameter values are integers (\mathbb{Z}). A sequence of actions is called an *action string*, and is denoted by σ . A set of observations S for an RA comprises a set of action strings that should be accepted S_+ , and a set of action strings that should be rejected S_- . An RA is consistent with $S = \{S_+, S_-\}$ if it accepts all action strings in S_+ , and rejects all action strings in S_- .

Roughly speaking, an RA can be described as a DFA (Definition 1) enriched with a finite set of registers R and two additional functions. The first function, τ , specifies which registers are in use in a location. In a location q there can be two types of transitions for a label l and parameter value v . If v is equal to some used register r , then the transition $\delta(q, l, r)$ is taken. Else (v is different to all used registers), the *fresh* transition $\delta(q, l, r_\perp)$ is taken.

The second function, π , specifies if and where to store a value v when a fresh transition ($\delta(q, l, r_\perp)$) is taken. If $\pi(q, l) = r_\perp$ then the value v on transition $\delta(q, l, r_\perp)$ is not stored. Else (if $\pi(q, l) = r$ for some register $r \in R$), the value v on transition $\delta(q, l, r_\perp)$ is stored in register r .

For the RA to behave as intended we introduce the following axioms. First, we require that no registers are used in the initial location:

$$\forall r \in R \quad \tau(q_0, r) = \text{false} \quad (5)$$

Second, if a register is used after a transition, it was used before or updated:

$$\begin{aligned} \forall q \in Q \quad \forall l \in L \quad \forall r \in R \quad \forall r' \in (R \cup \{r_\perp\}) \\ \tau(\delta(q, l, r'), r) = \text{true} \implies (\tau(q, r) = \text{true} \vee (r' = r_\perp \wedge \pi(q, l) = r)) \end{aligned} \quad (6)$$

Third, if a register is updated, then it is used afterwards:

$$\forall q \in Q \quad \forall l \in L \quad \forall r \in R \quad \pi(q, l) = r \implies \tau(\delta(q, l, r_\perp), r) = \text{true} \quad (7)$$

Our goal is to learn an RA that is consistent with a set of action strings $S = \{S_+, S_-\}$. For this, we need to define a function that keeps track of the valuation of registers during runs over these action strings. Let $A = (L, R^A, Q^A, q_0, \delta^A, \lambda^A, \tau^A, \pi^A)$ be an RA, and let $T = (L \times \mathbb{Z}, Q^T, \lambda^T)$ be an OT for S . In addition to the *map* function ($\text{map} : Q^T \rightarrow Q^A$), we define a *valuation function* $\text{val} : Q^T \times R^A \rightarrow \mathbb{Z}$ that maps a state of T and a register of A to the value the register contains in that state.

Before constructing constraints for action strings, we first *canonize* them by making them *neat* [2, Definition 7]. An action string is *neat* if each parameter value is equal to either a previous value, or to the largest preceding value plus one. To exemplify, let \mathbf{a} be an input label, and let $\mathbf{a}(3)\mathbf{a}(1)\mathbf{a}(3)\mathbf{a}(45)$ be an action string, then $\mathbf{a}(0)\mathbf{a}(1)\mathbf{a}(0)\mathbf{a}(2)$ is its corresponding neat action string. Aarts et al. [1] show that an RA can be learned by just studying its neat action strings.

Observation constraints for an RA are constructed as follows. First, we map the empty string to the initial location of A (Eq. 1). Second, we assert that a register is updated if and only if its valuation changes:

$$\forall \sigma l(v) \in Q^T \quad \forall r \in R^A \quad \text{val}(\sigma l(v), r) \neq \text{val}(\sigma, r) \Leftrightarrow \pi^A(\text{map}(\sigma), l) = r \quad (8)$$

We proceed by formulating how a register's valuation changes:

$$\forall \sigma l(v) \in Q^T \quad \forall r \in R^A$$

$$\text{val}(\sigma l(v), r) = \begin{cases} v & \text{if } \delta^A(\text{map}(\sigma), l, r_{\perp}) = \text{map}(\sigma l(v)) \\ & \wedge \pi(\text{map}(\sigma), l) = r \\ \text{val}(\sigma, r) & \text{otherwise} \end{cases} \quad (9)$$

We then encode the observed transitions:

$$\forall \sigma l(v) \in Q^T$$

$$\text{map}(\sigma l(v)) = \begin{cases} \delta^A(\text{map}(\sigma), l, r) & \text{if } \exists! r \in R : \tau^A(\text{map}(\sigma), r) = \text{true} \\ & \wedge \text{val}(\sigma, r) = v \\ \delta^A(\text{map}(\sigma), l, r_{\perp}) & \text{otherwise} \end{cases} \quad (10)$$

Finally, we encode the observed outputs, which can be done in the same way as for DFAs (see Eq. 3).

The task for the SMT solver is to find a solution that is consistent with these constraints. Obviously, we are interested in an RA with the minimal number of locations and registers. The number of locations can be limited in the same way as we limited states in DFAs (see Eq. 4). The number of registers is defined by the variables r that we quantify over in the presented equations. Therefore, they can be limited as such. In our case, the number of registers is never higher than the number of locations (because we can only update a single register from each location). Hence, the learning problem can be solved by iteratively incrementing the number of locations n , and for each n incrementing the number of registers from 1 to n , until a satisfiable encoding is found.

2.3 An Extension for Input-Output Register Automata

An *input-output register automaton* (IORA) is a register automaton transducer that generates an output action (i.e. label and value) after each input action. As in the RA-case, we restrict both input and output labels to a single parameter. Input and output values may update registers. Input values may be tested for (dis)equality with values in registers. Output values can be equal to the values stored, or may be fresh. As such, an input-output register automaton can be used for modeling software that produces parameterized outputs.

For a formal description of IORAs we refer to [1]. We define an IORA in Definition 4. Again, in the interest of our encoding, our definition is very different from that in [1]. Despite this, the semantics are similar.

Definition 4. *An IORA is a tuple $(I, O, R, Q, q_0, \delta, \lambda, \tau, \pi, \omega)$. Therein, I and O are finite, disjoint sets of input and output labels. R, Q, q_0, τ and π are the same*

as for an RA (Definition 3). $\delta : (Q \cup \{q_\perp\}) \times (I \cup O) \times (R \cup \{r_\perp\}) \rightarrow (Q \cup \{q_\perp\})$ is a register transition function with a sink location. $\lambda : (Q \cup \{q_\perp\}) \rightarrow \mathbb{B}$ is a location output function with a sink location while $\omega : Q \rightarrow \mathbb{B}$ is a location type function which returns **true** if a location is an input location, and **false** if it is an output location.

A set of observations S for an IORA consists of *action traces*, which are pairs (σ^I, σ^O) where $\sigma^I \in (I \times \mathbb{Z})^*$ is an *input action string*, and $\sigma^O \in (O \times \mathbb{Z})^*$ is an *output action string* with $|\sigma^I| = |\sigma^O|$. An IORA is consistent with a set S if for each pair $(\sigma^I, \sigma^O) \in S$ it generates σ^O when provided with σ^I .

Despite that semantically an IORA is a transducer, we define it as an RA (Definition 3) which distinguishes between input and output labels, and which defines an additional function ω for the location type. From an *input location* transitions are allowed only for input actions. After an input action the IORA reaches an *output location*, in which a *single* transition is allowed. This transition determines the output action generated in response, and the input location the IORA will transition to. Transitions that are not allowed lead to a designated *sink location*, which is denoted q_\perp .

Using this definition we incorporate the axioms defined for our RA encoding (Eqs. 5–7) also in our IORA encoding. To these, we add the following axioms for an IORA to behave as intended.

First, observe that we do not use λ as an output function for an IORA. Instead, we use it to denote which locations are allowed. Hence, we require that the sink location q_\perp is the only rejecting location:

$$\forall q \in (Q \cup \{q_\perp\}) \quad \lambda(q) = \begin{cases} \mathbf{false} & \text{if } q = q_\perp \\ \mathbf{true} & \text{otherwise} \end{cases} \quad (11)$$

Second, we require that transitions do not lead to the sink location:

$$\forall q \in Q \quad \forall o \in O \quad \forall r \in (R \cup \{r_\perp\}) \quad \omega(q) = \mathbf{true} \implies \delta(q, o, r) = q_\perp \quad (12)$$

$$\forall q \in Q \quad \forall i \in I \quad \forall r \in (R \cup \{r_\perp\}) \quad \omega(q) = \mathbf{false} \implies \delta(q, i, r) = q_\perp \quad (13)$$

$$\forall l \in I \cup O \quad \forall r \in (R \cup \{r_\perp\}) \quad \delta(q_\perp, l, r) = q_\perp \quad (14)$$

Finally, we require that input locations are *input enabled* (Eq. 15), and that there is only one transition possible in an output location (Eq. 16):

$$\forall q \in Q \quad \forall i \in I \quad \forall r \in (R \cup \{r_\perp\}) \quad \omega(q) = \mathbf{true} \implies \delta(q, i, r) \neq q_\perp \quad (15)$$

$$\forall q \in Q \quad \exists! o \in O \quad \exists! r \in (R \cup \{r_\perp\}) \quad \omega(q) = \mathbf{false} \implies \delta(q, o, r) \neq q_\perp \quad (16)$$

Our goal is to learn an IORA $A = (I, O, R^A, Q^A, q_0, \delta^A, \lambda^A, \tau^A, \pi^A, \omega^A)$ that is consistent with a set of action traces S . Because of the nature of our encoding, we consider each action trace $\sigma = (\sigma^I, \sigma^O)$ in S as an interleaving of the input action string σ^I and the output action string σ^O , i.e. $\sigma = \sigma_1^I \sigma_1^O \dots \sigma_{|\sigma^I|}^I \sigma_{|\sigma^I|}^O$. Let $T = ((I \cup O) \times \mathbb{Z}, Q^T, \lambda^T)$ be an OT for such strings.

The constraints for an IORA can now be constructed in the same way as for an RA (Eqs. 1 and 8–10). Observe that we do not use λ to encode observed outputs (this is already done by encoding the transitions of the OT). Instead, λ is used to denote which locations are allowed. All the locations in Q are allowed (because we have observed them) and q_{\perp} is the only location that is not allowed ($\lambda(q_{\perp}) = \mathbf{false}$ by Eq. 11). As such, we add the following constraint:

$$\forall \sigma \in Q^T \quad \text{map}(\sigma) \neq q_{\perp} \quad (17)$$

We can now determine if there is an IORA with at most n locations and m registers in the same way as for RAs.

3 Implementation and Evaluation

We implemented our encodings using Z3Py, the Python front-end of Z3 [9]². Our tool can generate an automaton model from a given set of observations (passive learning), or a reference to the system and a tester implementation (active learning), also when this system cannot be reset. We have also implemented a tester for the classes of automata supported. The tester generates test queries (or *tests*) each test consisting of an access sequence to an arbitrary state in the current hypothesis, and a sequence generated by a random walk from that state. In experiments, we configure the tester to build shorter tests. Longer tests worsen the scalability of our tool, and are unneeded for learning small models. Validity of the learned models was ensured by running a large number of tests on the last hypothesis and checking the number of states. We conducted a series of experiments to assess the scalability and effectiveness of our approach.

Our first experiment assesses the scalability of our encodings by adapting the scalable *Login*, *FIFO set* and *Stack* benchmarks of [2] to DFAs, Mealy machines, RAs and IORAs. To generate tests, we used the testing algorithm described earlier. The maximum length of the random sequence is $3 + \text{size}$, where *size* is the number of users or elements in the system. The *solver timeout* – the amount of time the solver was provided to compute a solution or indicate its absence – was set to 10 seconds for the DFA and Mealy systems, and to 10 min for the RA and IORA systems whose constraints could take considerably longer to process. Each learning run ends when either a solution is found (indicating success), or when a maximum bound for n is reached (indicating failure). We hence continue incrementing n even after encountering solver timeouts. By doing so, we lose the minimality guarantee yet may learn larger systems. For each system we performed 5 learning runs and collated the resulting statistics.

The results are shown in Table 1. Columns describe the system, the number of successful learning runs, the number of states/locations (which may vary due to loss of minimality) and registers (where applicable), average and standard deviation for the number of tests and inputs used in learning except for validating the last hypothesis, and for the amount of time learning took. The table only

² See <https://gitlab.science.ru.nl/rick/z3gi/tree/lata>.

Table 1. Scalable experiments. The model name encodes the formalism, type and size.

| Model | Succ | States loc | Regs | Tests | | Inputs | | Time (sec) | |
|-------------------|------|---------------|------|------------|------------|------------|---------|------------|---------|
| | | | | <i>Avg</i> | <i>Std</i> | <i>Avg</i> | Std | Avg | Std |
| DFA_FIFOSet(8) | 5 | 10.0 | | 228.0 | 81.11 | 2156.0 | 832.02 | 76.28 | 42.49 |
| DFA_FIFOSet(9) | 2 | 11.5 | | 413.5 | 161.93 | 4194.0 | 2104.35 | 199.99 | 26.1 |
| DFA_Login(3) | 5 | 11.0 | | 446.0 | 100.18 | 3092.0 | 781.73 | 131.84 | 39.26 |
| Mealy_FIFOSet(11) | 5 | 12.0 | | 280.0 | 110.65 | 3694.0 | 1722.57 | 79.75 | 23.69 |
| Mealy_FIFOSet(12) | 4 | 15.5 | | 370.0 | 200.12 | 5021.0 | 3312.85 | 227.15 | 290.99 |
| Mealy_FIFOSet(13) | 2 | 14.5 | | 526.5 | 318.91 | 8021.5 | 5608.06 | 190.36 | 51.96 |
| Mealy_Login(3) | 5 | 10.0 | | 104.0 | 10.03 | 726.0 | 69.93 | 52.4 | 4.83 |
| Mealy_Login(4) | 1 | 16.0 | | 241.0 | 0.0 | 2094.0 | 0.0 | 370.19 | 0.0 |
| RA_Stack(1) | 5 | 3.0 | 1 | 32.0 | 21.18 | 109.0 | 90.48 | 3.18 | 0.86 |
| RA_Stack(2) | 5 | 5.0 | 2 | 202.0 | 71.88 | 1018.0 | 394.58 | 124.72 | 53.41 |
| RA_FIFOSet(1) | 5 | 3.0 | 1 | 49.0 | 12.92 | 180.0 | 52.56 | 4.94 | 6.07 |
| RA_FIFOSet(2) | 5 | 6.0 | 2 | 365.0 | 88.41 | 2025.0 | 578.33 | 333.09 | 334.12 |
| RA_Login(1) | 5 | 4.0 | 1 | 306.0 | 163.18 | 1336.0 | 765.23 | 54.96 | 9.99 |
| RA_Login(2) | 3 | 8.0 | 2 | 1606.0 | 345.22 | 9579.0 | 2163.67 | 6258.11 | 1179.27 |
| IORA_Stack(1) | 5 | 7.0 | 1 | 7.0 | 1.58 | 24.0 | 6.63 | 8.77 | 1.92 |
| IORA_FIFOSet(1) | 5 | 7.0 | 1 | 8.0 | 3.27 | 31.0 | 9.36 | 6.45 | 0.84 |
| IORA_Login(1) | 5 | 9.0 | 1 | 33.0 | 6.65 | 152.0 | 29.89 | 1509.18 | 477.04 |

includes entries for the largest systems we could learn. A table with entries for all the systems learned is featured in the extended version of this work (See footnote 1).

As part of our second experiment we applied our tool on models learned in three case-studies [3, 4, 22]. These models are Mealy machines detailing aspects of the behavior of bankcard protocols, biometric passports and power control services (PCS). We used the tester described in [23]. This produces tests similar to our own, but extended by distinguishing sequences. These tests are parameterized by the length of the random sequence and a factor k . We set both to 1. The solver timeout is set to 1 min.

We use this experiment to also draw comparison between our approach and learners in the classical framework. To that end, to learn the case study models we also use the TTT [15] and classical L* [5] algorithms. Both of these algorithms are provided by the LearnLib learning tool (v0.12.1) [21]. To draw further comparison we additionally include scalable systems in this experiment. This allows us to compare the SMT-based approach with Tomte (v0.41) [2], a state-of-the-art learning tool for register automata. We note that the test configurations are similar for all learners. Due to the high standard deviation, we ran 20 experiments for each benchmark.

Table 2 shows the results of this experiment. We include the alphabet size and number of states/locations in the model, as well as the average time it took for the SMT approach to learn (this statistic is not useful for the other learners, which took at most a couple of seconds). We remark that the SMT-based approach is able to successfully learn all models every time, though it takes a long time for the larger ones. We also note that increasing the length of the random sequence used in tests even by 1 meant the SMT-based approach occasionally failed to infer the larger VISA model.

Analyzing Table 2, we remark that the SMT-based approach largely outperforms L^* and Tomte, while staying competitive with TTT. The approach performs better on scalable systems than on the case study models. This can be attributed to scalable systems requiring more counterexamples. To process these counterexamples, active learning algorithms require additional queries. Performance is further hindered if counterexamples are not optimal (e.g. they are unnecessarily long). Such problems do not affect the SMT-based approach.

Table 2. Comparison with other learners

| Model | States loc | Alph size | SMT | | | TTT | | L* | | Tomte | |
|--------------------|---------------|--------------|-------|--------|------|-------|--------|-------|--------|-------|--------|
| | | | Tests | Inputs | Time | Tests | Inputs | Tests | Inputs | Tests | Inputs |
| Biometric passport | 6 | 9 | 220 | 1057 | 26 | 220 | 941 | 333 | 1143 | | |
| MAESTRO | 6 | 14 | 835 | 4375 | 359 | 860 | 4437 | 1190 | 4718 | | |
| MasterCard | 6 | 14 | 839 | 4379 | 353 | 996 | 5260 | 1190 | 4718 | | |
| PIN | 6 | 14 | 757 | 3945 | 338 | 911 | 4769 | 1190 | 4718 | | |
| SecureCode | 4 | 14 | 313 | 1485 | 90 | 194 | 682 | 798 | 2758 | | |
| VISA | 9 | 14 | 796 | 4770 | 2115 | 750 | 4094 | 2040 | 9015 | | |
| PCS.1 | 8 | 9 | 629 | 3530 | 189 | 417 | 2179 | 657 | 2682 | | |
| PCS.2 | 3 | 9 | 71 | 279 | 9 | 75 | 196 | 252 | 657 | | |
| PCS.3 | 7 | 9 | 508 | 2651 | 154 | 476 | 2472 | 576 | 2196 | | |
| PCS.4 | 7 | 9 | 559 | 3024 | 154 | 451 | 2297 | 576 | 2196 | | |
| PCS.5 | 9 | 9 | 1120 | 6260 | 778 | 417 | 1753 | 1308 | 5340 | | |
| PCS.6 | 9 | 9 | 1158 | 6442 | 704 | 457 | 1977 | 1308 | 5340 | | |
| Mealy_FIFOSet(2) | 3 | 2 | 6 | 27 | 0 | 12 | 38 | 14 | 38 | | |
| Mealy_FIFOSet(7) | 8 | 2 | 52 | 481 | 7 | 71 | 588 | 235 | 2494 | | |
| Mealy_FIFOSet(10) | 11 | 2 | 179 | 2152 | 63 | 163 | 1822 | 486 | 6743 | | |
| Mealy_Login(2) | 6 | 3 | 37 | 214 | 7 | 57 | 242 | 57 | 219 | | |
| Mealy_Login(3) | 10 | 3 | 89 | 644 | 64 | 120 | 704 | 240 | 1720 | | |
| IORA_FIFOSet(1) | 7 | 2 | 9 | 31 | 7 | | | | | 21 | 36.5 |
| IORA_Stack(1) | 7 | 2 | 8.5 | 33 | 8 | | | | | 19 | 34 |
| IORA_Login(1) | 9 | 3 | 33 | 152 | 849 | | | | | 157 | 580 |

Our final experiment assesses our extension for learning systems without resets using benchmarks from recent related work [20]. For the sake of space we

summarize results and refer the reader to the extended version (See footnote 1) for details. Despite the simplicity of the test setup, the SMT-based approach performed only slightly worse than the sophisticated SAT-based approach of [20].

4 Conclusions

We have experimented with an approach for model learning which uses SMT solvers. The approach is highly versatile, as shown in its adaptations for learning FSMs and register automata, and for learning without resets. We provide an open source tool implementing these adaptations. Experiments indicate that our approach is competitive with the state-of-the-art. While the approach does not scale well, we have shown that it can be used for learning small models in practice. In the future we wish to improve the scalability of the approach via more efficient encodings. We hope this paper gives rise to a broader direction of future work, since the presented approach has several advantages over traditional model learning algorithms. Notably, it appears to be quite effective for rapid prototyping of learning algorithms for new formalisms and settings.

References

1. Aarts, F., et al.: Generating models of infinite-state communication protocols using regular inference with abstraction. *FMSD* **46**(1), 1–41 (2015)
2. Aarts, F., Fiterau-Brostean, P., Kuppens, H., Vaandrager, F.: Learning register automata with fresh value generation. In: Leucker, M., Rueda, C., Valencia, F.D. (eds.) *ICTAC 2015*. LNCS, vol. 9399, pp. 165–183. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25150-9_11. <http://www.sws.cs.ru.nl/publications/papers/fvaan/TomteFresh/>
3. Aarts, F., de Ruitter, J., Poll, E.: Formal models of bank cards for free. In: *ICST Workshops*, pp. 461–468. IEEE (2013)
4. Aarts, F., Schmaltz, J., Vaandrager, F.: Inference and abstraction of the biometric passport. In: Margaria, T., Steffen, B. (eds.) *ISoLA 2010*. LNCS, vol. 6415, pp. 673–686. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16558-0_54
5. Angluin, D.: Learning regular sets from queries and counterexamples. *I&C* **75**(2), 87–106 (1987)
6. Angluin, D.: Negative results for equivalence queries. *Mach. Learn.* **5**, 121–150 (1990)
7. Bruynooghe, M., et al.: Predicate logic as a modeling language: modeling and solving some machine learning and data mining problems with IDP3. *TPLP* **15**(6), 783–817 (2015)
8. Cassel, S., Howar, F., Jonsson, B., Steffen, B.: Active learning for extended finite state machines. *FAOC* **28**(2), 233–263 (2016)
9. De Moura, L., Bjørner, N.: Satisfiability modulo theories: introduction and applications. *CACM* **54**(9), 69–77 (2011)
10. Florêncio, C.C., Verwer, S.: Regular inference as vertex coloring. *TCS* **558**, 18–34 (2014)
11. Gold, E.: Language identification in the limit. *I&C* **10**(5), 447–474 (1967)

12. Heule, M., Verwer, S.: Software model synthesis using satisfiability solvers. *Empir. Softw. Eng.* **18**(4), 825–856 (2013)
13. De la Higuera, C.: *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press, Cambridge (2010)
14. Howar, F., Steffen, B., Jonsson, B., Cassel, S.: Inferring canonical register automata. In: Kuncak, V., Rybalchenko, A. (eds.) *VMCAI 2012*. LNCS, vol. 7148, pp. 251–266. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-27940-9_17
15. Isberner, M., Howar, F., Steffen, B.: The TTT algorithm: a redundancy-free approach to active automata learning. In: Bonakdarpour, B., Smolka, S.A. (eds.) *RV 2014*. LNCS, vol. 8734, pp. 307–322. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11164-3_26
16. Lee, D., Yannakakis, M.: Testing finite-state machines: state identification and verification. *IEEE Trans. Comput.* **43**(3), 306–320 (1994)
17. Lee, D., Yannakakis, M.: Principles and methods of testing finite state machines—a survey. *Proc. IEEE* **84**(8), 1090–1123 (1996)
18. Neider, D.: Computing minimal separating DFAs and regular invariants using SAT and SMT solvers. In: Chakraborty, S., Mukund, M. (eds.) *ATVA 2012*. LNCS, pp. 354–369. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33386-6_28
19. Peled, D., Vardi, M., Yannakakis, M.: Black box checking. In: *FORTE*, pp. 225–240. Kluwer (1999)
20. Petrenko, A., Avellaneda, F., Groz, R., Oriat, C.: From passive to active FSM inference via checking sequence construction. In: Yevtushenko, N., Cavalli, A.R., Yenigün, H. (eds.) *ICTSS 2017*. LNCS, vol. 10533, pp. 126–141. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67549-7_8
21. Raffelt, H., Steffen, B., Berg, T., Margaria, T.: LearnLib: a framework for extrapolating behavioral models. *STTT* **11**(5), 393–407 (2009)
22. Schuts, M., Hooman, J., Vaandrager, F.: Refactoring of legacy software using model learning and equivalence checking: an industrial experience report. In: Ábrahám, E., Huisman, M. (eds.) *IFM 2016*. LNCS, vol. 9681, pp. 311–325. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-33693-0_20
23. Smeenk, W., Moerman, J., Vaandrager, F., Jansen, D.N.: Applying automata learning to embedded control software. In: Butler, M., Conchon, S., Zaïdi, F. (eds.) *ICFEM 2015*. LNCS, vol. 9407, pp. 67–83. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25423-4_5
24. Smetsers, R.: *Grammatical Inference as a Satisfiability Modulo Theories Problem*. arXiv preprint [arXiv:1705.10639](https://arxiv.org/abs/1705.10639) (2017)
25. Vaandrager, F.: Model learning. *CACM* **60**(2), 86–95 (2017)
26. Walkinshaw, N., Derrick, J., Guo, Q.: Iterative refinement of reverse-engineered models by model-based testing. In: Cavalcanti, A., Dams, D.R. (eds.) *FM 2009*. LNCS, vol. 5850, pp. 305–320. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-05089-3_20