



# Tree-Based Adaptive Model Learning

Tiago Ferreira<sup>1</sup>, Gerco van Heerdt<sup>1</sup>, and Alexandra Silva<sup>1,2</sup>

<sup>1</sup> University College London, London, UK

{t.ferreira,gerco.heerdt}@ucl.ac.uk

<sup>2</sup> Cornell University, Ithaca, USA

alexandra.silva@cornell.edu

*A lot of the work the authors did in the last years on learning can be rooted back to learning from Frits. Gerco and Alexandra spent a few years in Nijmegen where Frits' enthusiasm for Angluin's algorithm infected them and made them want to keep learning. Later on, Tiago first learned about automata watching Frits' model learning video for CACM. On this landmark birthday, we thank Frits for his inspiration and wish him many happy returns!*

**Abstract.** We extend the Kearns–Vazirani learning algorithm to be able to handle systems that change over time. We present a new learning algorithm that can reuse and update previously learned behavior, implement it in the LearnLib library, and evaluate it on large examples, to which we make small adjustments between two runs of the algorithm. In these experiments our algorithm significantly outperforms both the classic Kearns–Vazirani learning algorithm and the current state-of-the-art adaptive algorithm.

## 1 Introduction

Formal methods have a long tradition and have had much success in critical applications, e.g. in the medical, space, and hardware industries. The last decade saw a rise in the use of formal methods in the software industry, with dedicated large teams in many companies, notably AWS and Facebook. This has caused a shift in focus to develop techniques that are helpful towards catching the most bugs as quickly as possible instead of performing complete verification [2].

The use of models to analyze system specifications, often pre-production, is common in certain domains, but requires expert knowledge to design the model and, throughout the system's life, update it. Motivated by the difficulties of building models, an automated technique called model learning [15] has been developed and used in the analysis of a range of (black-box system) implementations. One particularly successful application is that of network protocol implementations, e.g. TCP [6], SSH [7], and QUIC [5].

Classic active model learning algorithms like  $L^*$  [1], Kearns–Vazirani [13], and TTT [10] exist for a number of automata types (deterministic, input-output,

---

Full implementation and experiment results available at <https://github.com/UCL-PPLV/learnlib>.

weighted, register) and have enabled analysis of numerous systems by providing faithful models to be used in model checking. However, these algorithms suffer from a common issue: systems often change faster than we can learn their models, as there is an inherent assumption that the learning process is running from scratch every time. As such, keeping a model up-to-date becomes quite challenging, often needing manual intervention from an expert. This does however introduce the chance of producing a model that is not actually correct, while being an extremely laborious task for any moderately sized complex system.

In this paper, we present an *incremental* model learning algorithm that can cope with evolving systems more effectively, and does not have to restart the learning process when a change occurs, providing a gain in efficiency when learning systems that undergo changes at known moments in time.

The work in the present paper can be seen as advancing the state of the art in *adaptive automata learning* [3, 4, 8, 9, 16]. Whereas previous work adapted the  $L^*$  algorithm to reuse part of a data structure from a previous run, our adaptive algorithm is the first to use the more efficient Kearns–Vazirani algorithm. Although we focus on developing an algorithm for deterministic finite automata (DFAs), the techniques we developed can be transferred to other automata models. We expect this work and subsequent developments on adaptive learning will bring us closer to meet the specific needs of employing formal methods in fast moving environments commonly seen in the industry, and in large evolving systems.

The paper is organized as follows. After preliminary material on automata and learning (Sect. 2), in Sect. 3 we introduce and prove correctness of our incremental learning algorithm for DFAs, which targets evolving systems whose mutation points are known during learning. In Sect. 4 we benchmark and evaluate the efficiency of the algorithm, and we demonstrate its effectiveness in learning evolving systems. We conclude in Sect. 5 with a discussion of further directions on adaptive automata learning.

## 2 Preliminaries

Throughout this paper we will make use of standard notation from automata and learning theory. We define some of this notation below for the sake of clarity.

We fix a finite alphabet  $A$  and write  $A^*$  for the set of finite words over this alphabet. The empty word is  $\varepsilon$  and concatenation is written either by juxtaposition or with the operator  $\cdot$ . Given a word  $w \in A^*$  we write  $|w|$  for its length, and  $w[i]$ ,  $0 \leq i \leq |w|$ , for the prefix of  $w$  of length  $i$ . A language over  $A^*$  is a set of strings such that  $L \subseteq A^*$ . We sometimes refer to the characteristic function  $A^* \rightarrow 2 = \{\top, \perp\}$  of  $L$  also by  $L$  ( $L(w) = \top$  if  $w \in L$  and  $L(w) = \perp$  if  $w \notin L$ ).

The formal models generated by our algorithm are deterministic finite automata (DFAs). These are 4-tuples  $\mathcal{A} = \langle Q, q_0, \delta, F \rangle$  where  $Q$  is a finite set of states,  $q_0 \in Q$  is the initial state,  $\delta: Q \times A \rightarrow Q$  is a transition function, and  $F \subseteq Q$  is a set of final states. We define the usual reachability map  $\text{reach}_{\mathcal{A}}: A^* \rightarrow Q$  by  $\text{reach}_{\mathcal{A}}(\varepsilon) = q_0$  and  $\text{reach}_{\mathcal{A}}(ua) = \delta(\text{reach}_{\mathcal{A}}(u), a)$ . The language accepted by  $\mathcal{A}$  is  $\mathcal{L}_{\mathcal{A}} = \{u \in A^* \mid \text{reach}_{\mathcal{A}}(u) \in F\}$ . We write  $A_{\top}$  for the

minimal DFA accepting all of  $A^*$  and  $A_\perp$  for the minimal DFA accepting the empty language  $\emptyset$ .

Our learner makes use of two types of oracles: *membership* and *equivalence* oracles. A membership oracle  $\text{mq}: A^* \rightarrow 2$  is able to answer whether a given input word is accepted in the target system; an equivalence oracle takes a DFA  $\mathcal{A}$  and responds with  $\text{eq}(\mathcal{A}) \in A^* \cup \{\text{null}\}$ , which represents either that the DFA is correct (*null*) or a word  $w$  such that  $\mathcal{L}_{\mathcal{A}}(w) \neq \text{mq}(w)$ .

## 2.1 Learning with a Classification Tree

When designing a learning algorithm, one of the key aspects to consider is how we store the information we acquire over time. Learning then becomes a matter of being able to extend this structure with as little queries as possible, and transforming the data into a hypothesis automaton. The learning algorithm we introduce later, similarly to the classic Kearns–Vazirani algorithm, uses *classification trees* as its base data structure. Formally, the set of classification trees is given by the following grammar:

$$\text{CT} ::= \text{Node } A^* \text{ CT CT} \mid \text{Leaf } A^*$$

Here, a node contains a classifier  $e \in A^*$ , and  $\perp$ -child and  $\top$ -child subtrees, and a leaf contains a single access sequence  $s \in A^*$ . The child subtrees are named this way because of how the classifier  $e$  distinguishes access sequences  $s_\perp$  and  $s_\top$  present in the respective subtrees:  $\text{mq}(s_\perp \cdot e) = \perp$  and  $\text{mq}(s_\top \cdot e) = \top$ . In particular, this holds for every pair of leaves and their *lowest common ancestor* node, the root node of the smallest subtree containing both leaves.

The classification tree is then able to classify every word  $w \in A^*$  into a specific leaf of the tree, depending on how the target accepts or rejects  $w$  concatenated with specific classifiers  $e$  in the tree. This is done through *sifting* (Subroutine 1), where, starting from the root of the tree with classifier  $e$ , we pose the query  $\text{mq}(w \cdot e)$  and, depending on the result, proceed to the  $\perp$ -child or the  $\top$ -child of the node, from which we continue sifting, until we reach a leaf. The access sequence  $s$  of that leaf will then be deemed equivalent to  $w$ .

---

**Subroutine 1:** *sift* returns the leaf in a classification tree equivalent to a provided word w.r.t. the equivalence induced by the tree.

---

**Data:** Classification tree *tree*, membership oracle  $\text{mq}$ , word  $w$ .

**Result:** Equivalent leaf  $l$  in *tree*.

```

1  $n \leftarrow \text{tree}$ ;
2 while  $n = \text{Node } e \text{ left right}$  do
3   |  $n \leftarrow \text{mq}(w \cdot e) ? \text{right} : \text{left}$ ;
4 return  $n$ ;
```

---

The leaves of a classification tree represent the discovered states of the hypothesis, and sifting a word down the tree gives us the state this word should reach in the hypothesis. As such, one can easily retrieve the transitions of the

hypothesis from the tree by, for each leaf with access sequence  $s$  sifting each extended word  $s \cdot a$  to obtain the destination of the transition with symbol  $a \in A$  from  $s$ . The initial state is simply the state we find by sifting the empty word  $\varepsilon$ , and the accepting states are the leaves in the  $\top$ -child subtree of the root node, which will have classifier  $\varepsilon$ . This logic is used by `buildHyp` (Subroutine 2) to extract the DFA represented by a classification tree.

---

**Subroutine 2:** `buildHyp` extracts a hypothesis DFA from a classification tree in the Classic KV algorithm.

---

**Data:** Classification tree  $tree$

**Result:** Updated  $\mathcal{H}$  w.r.t the current tree  $tree$ .

```

1  $q_0 \leftarrow sift(tree, mq, \varepsilon); Q \leftarrow \{q_0\}; F \leftarrow \emptyset;$ 
2 for  $l \in leaves(tree)$  do
3    $Q \leftarrow Q \cup \{l\};$ 
4   if  $l \in leaves(child(tree, \top))$  then
5      $F \leftarrow F \cup \{l\};$ 
6 for  $l \in Q$  do
7   for  $a \in \Sigma$  do
8      $\delta(l, a) \leftarrow sift(tree, mq, label(l) \cdot a);$ 
9 return  $\langle \Sigma, Q, q_0, F, \delta \rangle;$ 
```

---

With a hypothesis extracted from the classification tree, we can now pass this to an equivalence oracle to determine if the hypothesis is correct. If not, we will receive a counterexample word that we can use to improve the classification tree. The algorithm does this by understanding that, given that every hypothesis classifies the empty string  $\varepsilon$  correctly, and by definition the current hypothesis classifies the counterexample  $c$  incorrectly, there must be a prefix of  $c$  for which the classification first diverges. In terms of states, this then means that we are taking a transition into a state that is accepting in the hypothesis, and rejecting in the target, or vice-versa. This is fixed by realising that the state we take this transition from then must actually be two different states. The algorithm then uses this logic in `updateTree` (Subroutine 3) to split the state into two at the tree level, turning a leaf into a node with two leaves, one representing the new state discovered by the counterexample.

---

**Subroutine 3:** `updateTree` with a provided counterexample.

---

**Data:** Classification tree  $tree$ , counterexample  $c \in A^*$ .

**Result:** Updated  $tree$  taking into account  $c$ .

```

1 for  $i \in [0 \dots length(c) - 1]$  do
2    $s_i \leftarrow sift(tree, mq, c[i]); \hat{s}_i \leftarrow reach_{\mathcal{H}}(c[i]);$ 
3   if  $s_i \neq \hat{s}_i$  then
4      $e \leftarrow c_i \cdot LCA(tree, s_i, \hat{s}_i);$ 
5      $tree \leftarrow split(tree, s_{i-1}, c[i - 1], e, mq(c[i - 1] \cdot e));$ 
6   return  $tree;$ 
```

---

We provide the classic Kearns–Vazirani algorithm in Algorithm 4. This algorithm uses the `buildHyp` routine explained above to build a DFA from a classification tree, as well as `updateTree` (Subroutine 3) to extend the classification tree on receipt of a counterexample from an equivalence oracle.

---

**Algorithm 4:** Classic Kearns–Vazirani Algorithm
 

---

**Data:** Alphabet  $A$ , membership oracle `mq` and equiv. oracle `eq` for language  $L$ .

**Result:** The learned DFA  $\mathcal{H}$  accepting  $L$ .

```

1 init  $\leftarrow$  mq( $\varepsilon$ );  $\mathcal{H} \leftarrow$  init ?  $A_{\top}$  :  $A_{\perp}$ ;  $s \leftarrow$  eq( $\mathcal{H}$ );
2 if  $s \neq$  null then
3   tree  $\leftarrow$  init ? Node  $\varepsilon$  (Leaf  $s$ ) (Leaf  $\varepsilon$ ) : Node  $\varepsilon$  (Leaf  $\varepsilon$ ) (Leaf  $s$ );
4    $\mathcal{H} \leftarrow$  buildHyp(tree);
5   cex  $\leftarrow$  eq( $\mathcal{H}$ );
6   while cex  $\neq$  null do
7     tree  $\leftarrow$  updateTree(tree, cex);
8      $\mathcal{H} \leftarrow$  buildHyp(tree);
9     cex  $\leftarrow$  eq( $\mathcal{H}$ );
10 return  $\mathcal{H}$ ;
```

---

### 3 Learning Evolving Systems Incrementally

We now develop a learning algorithm that is able to learn updates to a previous model without having to discard all behavior learned so far, and is also able to detect and remove behavior that no longer holds.

Classic algorithms are partially able to do this with equivalence oracles—they correct the current hypothesis based on a counterexample. *Adaptive model learners* are able to do this with the answers of both membership and equivalence oracles, even if the answer conflicts previous ones. Thus, they can deal with languages that mutate over time, adapting to changes by either trimming outdated behavior or distinguishing new behavior.

Our adaptive learning algorithm is targeted at systems with discrete changes, such as version controlled systems. Specifically, the system evolves at discrete known points, such as every version, forming a stream of target systems. Our *incremental learning algorithm* for DFAs, presented in Algorithm 6, is based on Kearns–Vazirani. As a first crucial difference, the incremental algorithm uses a previous learned model as its starting point. As such, it cannot just acquire new information; it also needs to be able to trim outdated behavior. This is done by `minimizeTree` (Subroutine 5), which prunes an initial tree by removing all leaves that are no longer represented by their reported access sequence. It achieves this by sifting every access sequence down the tree, removing leaves whose access sequences do not sift back into themselves.

This guarantees not only that the leaves left in our tree are correct in a correct automaton for this language, but that every leaf in the tree is unique w.r.t. the Myhill–Nerode congruence. If a pair of leaves were equivalent, then

---

**Subroutine 5:** `minimizeTree` trims the tree from redundant leaves.

---

**Data:** Classification tree  $tree$ , membership oracle  $mq$ .

**Result:** A minimized classification tree.

```

1 for  $l \in leaves(tree)$  do
2    $s \leftarrow sift(tree, mq, label(l));$ 
3   if  $s \neq l$  then
4      $tree \leftarrow removeLeaf(tree, l);$ 
5 return  $tree;$ 

```

---



---

**Algorithm 6:** Incremental Algorithm

---

**Data:** Fixed alphabet  $A$ , optional previous classification tree  $tree$ , membership oracle  $mq$  and equivalence oracle  $eq$  w.r.t the language  $L$ .

**Result:** The learned DFA  $\mathcal{H}$  equivalent to the language  $L$ .

```

1  $init \leftarrow mq(\varepsilon); \mathcal{H} \leftarrow init ? A_{\top} : A_{\perp}; s \leftarrow eq(\mathcal{H});$ 
2 if  $s \neq null$  then
3   if  $tree = null$  then
4      $tree \leftarrow init ? Node \varepsilon (Leaf s) (Leaf \varepsilon) : Node \varepsilon (Leaf \varepsilon) (Leaf s);$ 
5   else
6      $tree \leftarrow minimizeTree(tree, mq);$ 
7    $\mathcal{H} \leftarrow buildHyp(tree);$ 
8    $cex \leftarrow eq(\mathcal{H});$ 
9   while  $cex \neq null$  do
10     $tree \leftarrow updateTree(tree, cex);$ 
11     $\mathcal{H} \leftarrow buildHyp(tree);$ 
12     $cex \leftarrow eq(\mathcal{H});$ 
13 return  $\mathcal{H};$ 

```

---

both their access sequences would sift into only one of the nodes, leaving a leaf whose access sequence does not sift back into itself, and causing it to be removed.

While here we only present the relevant changes made to the classic Kearns–Vazirani algorithm to be able to adapt to changing behavior, we include the full algorithm with all its subroutines in Appendix A.

### 3.1 Correctness and Termination

As the algorithm only terminates with a hypothesis that is correct according to an equivalence query, correctness follows from termination. Termination of the original Kearns–Vazirani algorithm relies on the following key property: for every subtree of the form `Node  $e$  left right`, each leaf  $s \in A^*$  of `left` satisfies  $se \notin \mathcal{L}$  and every leaf  $s \in A^*$  of `right` satisfies  $se \in \mathcal{L}$ . We note that this property also holds in our incremental algorithm as soon as we enter the main loop, as `minimizeTree` removes any leaf violating it.

When a counterexample  $cex$  is found, the procedure is the same as for the original Kearns–Vazirani algorithm, and can only be applied a finite number of times: By the property shown above every pair of leaves corresponds to a

pair of distinct equivalence classes of the Myhill–Nerode congruence for  $\mathcal{L}$ , and therefore the leaves in the tree cannot exceed the number of equivalence classes. Furthermore, every counterexample of length at least 2 leads to an increase of the number of leaves (via `updateTree`, which preserves the above invariant).

## 4 Experiments

We evaluate the efficiency of our new learning algorithm by running experiments over random targets with different types of features. While we would like to evaluate it over a standard set of benchmarks [14], these currently only cover single target automata, and so are not fit for adaptive learners designed to learn automata that are *linked* due to small evolutions in their behavior. We designed two scenarios to benchmark this incremental algorithm. The first scenario takes an initial automaton and applies a series of random mutations: it randomly adds a state, removes a state, diverts a transition, and flips the acceptance of a state. Our second scenario simulates the common occurrence of adding a feature to an existing system. We do this by introducing a small *feature automaton* of 3 states to an original base automaton by diverting 3 random transitions into the start state of the feature automaton.

We perform these benchmarks on different automata of increasing number of states, while maintaining the number of mutations applied to them, and the size of the feature automaton. This way, we create different ratios of change, and simulate applying fixes, or adding features to different systems.


We call the first and second targets  $t_0$  and  $t_1$ , respectively. For our adaptive learning algorithm, the target evolving system starts as  $t_0$  and mutates to  $t_1$  after 10000 queries. For the classic Kearns–Vazirani algorithm, as it cannot learn evolving systems at all, we have to run the algorithm twice, first targeting  $t_0$ , then from scratch targeting  $t_1$ . To ensure repeatable results, each benchmark in question has been run 300 times, each with fresh random inputs of the same parameters, and averaged. The graphs below represent the average run of both benchmarks, using that both have very similar results. Separate graphs per benchmark can be found in Appendix B.

We start by running the benchmark on the classic Kearns–Vazirani algorithm to set a baseline. These results can be seen in Fig. 1. The progress of each instance here is measured according to the following definition.


**Definition 1.** Given  $\alpha \in [0, 1]$ , a stream  $(t_i)_{i \in \mathbb{N}}$  of target automata, and a stream  $(h_i)_{i \in \mathbb{N}}$  of hypothesis automata,<sup>1</sup> the progress is the stream  $(p_i)_{i \in \mathbb{N}}$ , with  $p_i \in [0, 1]$  for all  $i \in \mathbb{N}$ , given by

$$p_i = \sum_{u \in A^*, \mathcal{L}_{t_i}(u) = \mathcal{L}_{h_i}(u)} (1 - \alpha) \cdot \left( \frac{\alpha}{|A|} \right)^{|u|}.$$

<sup>1</sup> Finite streams may be turned into infinite ones by repeating the last element.



**Fig. 1.** Average progress graph of the classic Kearns–Vazirani algorithm.



**Fig. 2.** Average progress graph of the incremental algorithm.

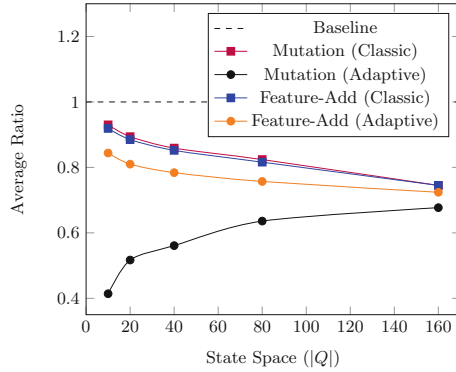
As the learning curves show, the linear process of the classic algorithm for  $t_1$  is very similar to the one for  $t_0$ . The two halves of the lines may not have the exact same gradients due to the targets being different, but we can see they follow a similar pattern, and more importantly, converge to 1.0 using a similar number of queries. This is because no knowledge is reused, and all states, even persisting ones, will have to be relearned.

We perform the same benchmark on the incremental algorithm: see results in Fig. 2. The incremental algorithm has a  $t_0$  run very similar to the classic algorithm, due to the lack of previous knowledge. However, the  $t_1$  segment is already very different. We immediately see that it does not start from such a low similarity value as the run of the classic algorithm. This is because while the classic algorithm always starts from a one state automaton, our incremental algorithm starts from the previous hypothesis, with outdated states pruned.

We can also see that it is not always the case that this line ascends immediately at the mutation point. This is because, as we are not starting from a widely dissimilar automaton, the equivalence oracle actually requires a number of queries to find a counterexample. This can be optimized by using more efficient equivalence oracles, but for the sake of simplicity and comparison all algorithms use the same random word search algorithm for equivalence testing.

These learning progress graph representations are great at demonstrating the overall behavior and approach of the learning algorithm, but we now want to evaluate whether this algorithm does indeed provide a benefit over learning systems classically, or with the current state-of-the-art adaptive algorithms, in terms of the number of queries it takes them to fully learn mutated targets. As such, we have computed and averaged the number of queries it takes to reach the final hypothesis while learning the  $t_1$  on each run. We present the results as two ratios relative to the incremental algorithm: one comparing the classic Kearns–Vazirani algorithm, and another comparing the current state-of-the-art adaptive algorithm, Partial Dynamic  $L^*$  [4]. We plot these in Fig. 3, and show how the ratio changes with the size of the state space of  $t_0$ .





**Fig. 3.** Ratio of number of queries needed to learn the automata.

In this representation, a value of 1.0 would indicate that the incremental algorithm requires as many queries as the algorithm in comparison. Any value below indicates a benefit in running the incremental algorithm over such algorithm, and vice versa (the lower the ratio, the better the incremental algorithm).

As we can see, **the incremental algorithm consistently outperforms all other algorithms in all benchmarks**, even with small targets where a relatively big portion of the system has changed. When compared to the adaptive algorithm from the literature, we can see the incremental algorithm still consistently outperforms in terms of queries, with a tendency to plateau between a 0.75 and 0.7 ratio—a reduction by 25–30% in the number of queries.

## 5 Conclusion

We introduced a new state-of-the-art algorithm for *adaptive learning* which provides, to our knowledge, the most efficient adaptive learner to date, allowing us to learn systems that were before too big or evolved too quickly to be learned classically. We evaluated the fitness of this algorithm through a set of realistic experiments, demonstrating their benefit over classic and existing adaptive learners. Adaptive learning could provide a cooperative relationship between active and passive learning, due to its flexibility towards information that changes over time and becomes available at different stages. We want to explore this in future work, as well as developing a formal framework of adaptive learning, where other algorithms can be easily adapted, e.g. efficient classic algorithms such as TTT [11]. Finally, although our incremental algorithm allows us to learn systems that evolve at known points, this would not work in a true black-box scenario, where we cannot know if or when the system changes. In the future we would like to then develop a *continuous* learning algorithm for such evolving systems.

**Related Work.** Our algorithm contributes to the field of adaptive learning first introduced in [8], where information learned in previous models was used as

guidance for which states to check first, instead of blindly looking for new ones. This was done by slightly modifying the  $L^*$  algorithm to start from a previous set of access sequences. Chaki et al. [3] use a similar algorithm in combination with assume-guarantee reasoning [12] to provide a framework where model checking is used to find counterexamples in the model, and thus make progress in learning. Their algorithm, Dynamic  $L^*$ , reuses not only the starting prefix/suffix sets, but also their computed values. However, these must still be validated on the new target. Finally, Damasceno et al. [4] introduce Partial Dynamic  $L^*$ , which improves Dynamic  $L^*$  by analysing the start prefix/suffix sets to trim them where possible, reducing the amount of information that needs to be validated.

These previous algorithms, however, suffer from using an observation table as their data structure, which increases the amount of redundant data to be acquired. As the tables grow, these redundancies significantly increase the number of queries that need to be asked.

## A Omitted Incremental Subroutines

---

**Subroutine A.1:** `child` returns the ( $\perp/\top$ )-child of a provided tree node.

---

**Data:** Tree node  $p$ ,  $b \in 2$

**Result:**  $b$ -child of  $p$ .

```

1 if  $p = \text{Node } e \text{ left right}$  then
2 |   return  $b ? \text{right} : \text{left}$ ;
3 return null;

```

---



---

**Subroutine A.2:** `children` returns the children of a provided tree node.

---

**Data:** Tree node  $p$ .

**Result:** Set of children of  $p$ .

```

1  $C \leftarrow \emptyset$ ;
2 if  $\text{child}(p, \perp) \neq \text{null}$  then
3 |    $C \leftarrow C \cup \{\text{child}(p, \perp)\}$ ;
4 if  $\text{child}(p, \top) \neq \text{null}$  then
5 |    $C \leftarrow C \cup \{\text{child}(p, \top)\}$ ;
6 return  $C$ ;

```

---

---

**Subroutine A.3:** `setChild` updates a child of a given inner node in a classification tree.

---

**Data:** Classification tree  $tree$ , parent node  $p$ , child outcome  $b$ , new child  $n$ .

**Result:** Updated classification tree  $tree$  where  $n$  is the  $b$ -child of  $p$ .

```

1 if  $tree = \text{Node } e \text{ left right}$  then
2   | if  $tree = p$  then
3     | if  $b$  then
4       |   return  $\text{Node } e \text{ left } n$ ;
5     |   else
6       |   return  $\text{Node } e \text{ } n \text{ right}$ ;
7   | return  $\text{Node } e \text{ setChild}(left, p, b, n) \text{ setChild}(right, p, b, n)$ ;
8 return  $tree$ ;

```

---



---

**Subroutine A.4:** `nodes` returns the set of all nodes in a given tree.

---

**Data:** Classification tree  $tree$ .

**Result:** Set of nodes in  $tree$ .

```

1 if  $tree = \text{Leaf } s$  then
2   | return  $\{tree\}$ ;
3 return  $tree \cup \text{nodes}(\text{child}(tree, \perp)) \cup \text{nodes}(\text{child}(tree, \top))$ ;

```

---



---

**Subroutine A.5:** `leaves` returns the set of leaves in a given tree.

---

**Data:** Classification tree  $tree$ .

**Result:** Set of leaves in  $tree$ .

```

1 if  $tree = \text{Leaf } s$  then
2   | return  $\{tree\}$ ;
3 return  $\text{leaves}(\text{child}(tree, \perp)) \cup \text{leaves}(\text{child}(tree, \top))$ ;

```

---



---

**Subroutine A.6:** `label` returns the label of a given node, be it a classifier or an access sequence.

---

**Data:** Node  $n$ .

**Result:** Label  $\in A^*$ .

```

1 if  $n = \text{Leaf } s$  then
2   | return  $s$ ;
3 if  $n = \text{Node } e \text{ left right}$  then
4   | return  $e$ ;

```

---

---

**Subroutine A.7:** `setLabel` replaces the label in a specific leaf.

---

**Data:** Classification tree  $tree$ , leaf  $l$ , new label  $w$ .

**Result:** Updated classification tree  $tree$  with new label  $w$  in  $l$ .

```

1 if  $tree = l$  then
2   | return Leaf  $w$ ;
3 if  $tree = \text{Node } e \text{ left right}$  then
4   | return Node  $e \text{ setLabel(left, l, w) setLabel(right, l, w)$ ;
5 return  $tree$ ;

```

---

**Subroutine A.8:** `outcome` returns whether a provided node  $n$  is a  $\perp$ -child or a  $\top$ -child, or neither.

---

**Data:** Classification tree  $tree$ , provided node  $n$

**Result:**  $\perp$ ,  $\top$  or  $null$  if the node does not have a parent in  $tree$

```

1 if  $n = tree$  then
2   | return  $null$ ;
3 return  $child(\text{parent}(tree, n), \top) = n$ ;

```

---

**Subroutine A.9:** `parent` returns the parent of a provided node in the classification tree.

---

**Data:** Classification tree  $tree$ , child leaf  $l$ .

**Result:** Parent node  $p$ .

```

1  $q \leftarrow \text{Queue}(tree)$ ;
2 while  $|q| \neq 0$  do
3   |  $p \leftarrow \text{pop}(q)$ ;
4   | if  $l \in \text{children}(p)$  then
5     | return  $p$ ;
6   | for  $c \in \text{children}(p)$  do
7     |  $q \leftarrow \text{push}(q, c)$ ;
8 return  $null$ ;

```

---

**Subroutine A.10:** `removeLeaf`

---

**Data:** Classification tree  $tree$ , leaf  $l \in \text{leaves}(tree)$  to be removed.

**Result:** A valid classification tree  $tree$  with  $l$  removed.

```

1  $node \leftarrow \text{parent}(tree, l)$ ;
2  $sibling \leftarrow \text{child}(node, \neg \text{outcome}(tree, l))$ ;
3 return  $\text{setChild}(tree, \text{parent}(tree, node), \text{outcome}(tree, node), sibling)$ ;

```

---

**Subroutine A.11:** `LCA` returns the lowest common ancestor node of two provided leaves in a classification tree.

---

**Data:** Classification tree  $tree$ , first leaf  $l_a$ , second leaf  $l_b$ .

**Result:** LCA node  $n$  in the classification tree  $tree$ .

```

1  $n \leftarrow l_a$ ;
2 while  $l_b \notin \text{leaves}(n)$  do
3   |  $n \leftarrow \text{parent}(tree, n)$ ;
4 return  $n$ ;

```

---

**Subroutine A.12: split** splits a leaf in the tree into a node with 2 child leaves, one of them a new leaf introduced to the tree.

**Data:** Classification tree  $tree$ , current leaf  $l$  being split, label  $w \in A^*$  for the new leaf, classifier  $e \in A^*$  for the new node,  $b \in 2$  indicating whether the new leaf is a  $\top$ -child.


**Result:** Updated classification tree  $tree$ .

- 1  $node \leftarrow b ? \text{Node } e \ l \ (\text{Leaf } w) : \text{Node } e \ (\text{Leaf } w) \ l;$
- 2 **return**  $setChild(tree, parent(tree, l), outcome(tree, l), node);$


## B Additional Experiment Graphs

### B.1 Mutation Benchmark

(See Figs. 4 and 5)




**Fig. 4.** Average progress graph of the classic Kearns–Vazirani algorithm.




**Fig. 5.** Average progress graph of the incremental algorithm.

## B.2 Feature-Add Benchmark

(See Figs. 6 and 7)



**Fig. 6.** Average progress graph of the classic Kearns–Vazirani algorithm.



**Fig. 7.** Average progress graph of the incremental algorithm.

## References

1. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**, 87–106 (1987). [https://doi.org/10.1016/0890-5401\(87\)90052-6](https://doi.org/10.1016/0890-5401(87)90052-6)
2. Bornholt, J., et al.: Using lightweight formal methods to validate a key-value storage node in Amazon S3. In: SIGOPS, pp. 836–850. ACM (2021). <https://doi.org/10.1145/3477132.3483540>
3. Chaki, S., Clarke, E.M., Sharygina, N., Sinha, N.: Verification of evolving software via component substitutability analysis. *Formal Methods Syst. Des.* **32**(3), 235–266 (2008). <https://doi.org/10.1007/s10703-008-0053-x>
4. Damasceno, C.D.N., Mousavi, M.R., da Silva Simao, A.: Learning to reuse: adaptive model learning for evolving systems. In: Ahrendt, W., Tapia Tarifa, S.L. (eds.) IFM 2019. LNCS, vol. 11918, pp. 138–156. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-34968-4\\_8](https://doi.org/10.1007/978-3-030-34968-4_8)
5. Ferreira, T., Brewton, H., D’Antoni, L., Silva, A.: Prognosis: closed-box analysis of network protocol implementations. In: SIGCOMM, pp. 762–774. ACM (2021). <https://doi.org/10.1145/3452296.3472938>
6. Fiterău-Broștean, P., Janssen, R., Vaandrager, F.: Combining model learning and model checking to analyze TCP implementations. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 454–471. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-41540-6\\_25](https://doi.org/10.1007/978-3-319-41540-6_25)
7. Fiterău-Broștean, P., Lenaerts, T., Poll, E., de Ruiter, J., Vaandrager, F.W., Verleg, P.: Model learning and model checking of SSH implementations. In: SPIN, pp. 142–151. ACM (2017). <https://doi.org/10.1145/3092282.3092289>
8. Groce, A., Peled, D., Yannakakis, M.: Adaptive model checking. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, pp. 357–370. Springer, Heidelberg (2002). [https://doi.org/10.1007/3-540-46002-0\\_25](https://doi.org/10.1007/3-540-46002-0_25)

9. Huistra, D., Meijer, J., van de Pol, J.: Adaptive learning for learn-based regression testing. In: Howar, F., Barnat, J. (eds.) FMICS 2018. LNCS, vol. 11119, pp. 162–177. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-00244-2\\_11](https://doi.org/10.1007/978-3-030-00244-2_11)
10. 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](https://doi.org/10.1007/978-3-319-11164-3_26)
11. Isberner, M., Howar, F., Steffen, B.: The open-source LearnLib. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 487–495. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-21690-4\\_32](https://doi.org/10.1007/978-3-319-21690-4_32)
12. Jones, C.B.: Tentative steps toward a development method for interfering programs. *ACM Trans. Program. Lang. Syst.* **5**, 596–619 (1983). <https://doi.org/10.1145/69575.69577>
13. Kearns, M.J., Vazirani, U.V.: *An Introduction to Computational Learning Theory*. MIT Press, Cambridge (1994). <https://mitpress.mit.edu/books/introduction-computational-learning-theory>
14. Neider, D., Smetsers, R., Vaandrager, F., Kuppens, H.: Benchmarks for automata learning and conformance testing. In: Margaria, T., Graf, S., Larsen, K.G. (eds.) *Models, Mindsets, Meta: The What, the How, and the Why Not?* LNCS, vol. 11200, pp. 390–416. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-22348-9\\_23](https://doi.org/10.1007/978-3-030-22348-9_23)
15. Vaandrager, F.W.: Model learning. *Commun. ACM* **60**, 86–95 (2017). <https://doi.org/10.1145/2967606>
16. Windmüller, S., Neubauer, J., Steffen, B., Howar, F., Bauer, O.: Active continuous quality control. In: CBSE, pp. 111–120. ACM (2013). <https://doi.org/10.1145/2465449.2465469>