Combining Black-Box and White-Box Techniques for Learning Register Automata

Falk Howar¹(✉), Bengt Jonsson², and Frits Vaandrager³

¹ Dortmund University of Technology and Fraunhofer ISST, Dortmund, Germany
`falk.howar@tu-dortmund.de`

² Department of Information Technology, Uppsala University, Uppsala, Sweden

³ Institute for Computing and Information Sciences,
Radboud University, Nijmegen, The Netherlands

Abstract. Model learning is a black-box technique for constructing state machine models of software and hardware components, which has been successfully used in areas such as telecommunication, banking cards, network protocols, and control software. The underlying theoretic framework (active automata learning) was first introduced in a landmark paper by Dana Angluin in 1987 for finite state machines. In order to make model learning more widely applicable, it must be further developed to scale better to large models and to generate richer classes of models. Recently, various techniques have been employed to extend automata learning to extended automata models, which combine control flow with guards and assignments to data variables. Such techniques infer guards over data parameters and assignments from observations of test output. In the black-box model of active automata learning this can be costly and require many tests, while in many application scenarios source code is available for analysis. In this paper, we explore some directions for future research on how black-box model learning can be enhanced using white-box information extraction methods, with the aim to maintain the benefits of dynamic black-box methods while making effective use of information that can be obtained through white-box techniques.

1 Introduction

Model learning, also known as active automata learning, is a black-box technique which constructs state machine models of software and hardware components from information obtained by providing inputs and observing the resulting outputs. Model learning has been successfully used in several different application domains, including

- generating conformance test suites of software components, a.k.a. *learning-based testing* (e.g., [38, 39]),
- finding mistakes in implementations of security-critical protocols (e.g., [3, 16, 28–30, 62]),

- learning interfaces of classes in software libraries (e.g., [42]),
- checking that a legacy component and a refactored implementation have the same behavior (e.g., [63]).

Active automata learning techniques are partly based on a landmark paper by Angluin [6] showing that finite automata can be learned in the so-called *Minimally Adequate Teacher* (MAT) framework [6], using two types of queries. A *membership* (or *I/O*) *query* asks what the output is in response to an input sequence. An *equivalence query* asks whether a hypothesized state machine is equivalent to the sought finite automaton; it is answered by *yes* if this is the case, otherwise by a *counterexample* that distinguishes the hypothesized and sought state machines.

Steffen et al. [38] made the important observation that the MAT framework can be used for black-box learning of abstract and concise state machine models of software components. They assume a software component, called the *System Under Learning* (SUL), whose behavior can be described by (an unknown) state machine, and which can always be brought back to its initial state. An I/O query can now be implemented by bringing the SUL to its initial state and then observing the outputs generated in response to the given input sequence. Equivalence queries can be approximated using a conformance testing tool [9, 52, 68] via a finite number of *test queries* to the SUL. Peled et al. [37, 60] observed that learning models can be used as a basis for model checking of black-box components.

The most widely known algorithm for model learning of finite automata is L^* [6], which has, e.g., been implemented in the LearnLib framework [48]. A key strength of model learning is that it aims to produce succinct models of the externally observable behavior of the SUL. This allows it to extract simple models of complex software, especially if we choose the right perspective (e.g., focus on a subset of a component's functionality) and apply appropriate abstractions. Examples are implementations of network protocols, which typically consist of many thousands of lines of code, but after appropriate abstraction induce state diagrams with at most a few dozen states, see e.g. [29, 62].

There is certainly a large potential for application of model learning to many different aspects of software development, maintenance and refactoring, especially when it comes to handling legacy software. We survey examples of existing such applications in Sect. 3. To realize this potential, two major challenges must be addressed: (1) currently, techniques do not scale well, and (2) they are not yet satisfactorily developed for richer classes of models. Let us elaborate on these challenges.

1. Concerning scaling, we note that the complexity of the currently most efficient algorithm for active learning of finite-state models [50] has a cubic worst-case time complexity in the size of the learned model. Another expensive component of model learning is that sufficiently precise approximation of equivalence queries in a black-box setting may require a number of membership queries that is exponential in the number of states of the SUL.

2. Concerning richness of models, in many situations it is crucial for models to also be able to describe *data flow*, i.e., constraints on data parameters that are passed when the component interacts with its environment, as well as the mutual influence between control flow and data flow. For instance, models of protocol components must describe how different parameter values in sequence numbers, identifiers, etc. influence the control flow, and vice versa. Such models often take the form of *extended finite state machines* (EFSMs). Recently, various techniques have been employed to extend automata learning to EFSM models, which combine control flow with guards and assignments to data variables [3, 15]. Such techniques either rely on manually constructed mappers that abstract the data aspects of input and output symbols into a finite alphabet, or otherwise infer guards and assignments from observations of test outputs. The latter can be costly, especially for models where control flow depends on test on data parameters in input: in this case, learning an exact guard that separates two control flow branches may require a large number of queries. For instance, to infer that a branch is taken if an input parameter is greater than 42 may take a number of membership queries.

One way to address these challenges is to augment model learning with white-box information extraction methods, which are able to obtain information about the SUL at lower cost than black-box techniques. When dealing with computer-based systems, there is a spectrum of how much information we have about the code. For third party components that run on separate hardware, we may not have access to the code at all. Frequently we will have access to the executable, but not anymore to the original code. Or we may have access to the code, but not to adequate tools for analyzing it (this often happens with legacy components). If we can construct a good model of a component using black-box learning techniques, we do not need to worry about the code. However, in cases where black-box techniques do not work and/or the number of queries becomes too high, it makes sense to exploit information from the code during model learning.

In this article, we explore how existing approaches for model learning can be improved by effective use of available white-box information about the SUL (the full code, the executable,...), with the aim to maintain the benefits of black box methods.¹ We will develop three promising directions for future research:

1. In a black-box setting equivalence queries are approximated using a finite number of test queries to the SUL. This is time consuming and often constitutes a major bottleneck in model learning [65]. Moreover, the approximation may be incorrect because, as Dijkstra observed, testing can be used to show the presence of bugs, but never to show their absence. We have no guarantees that a learned model is correct. In Sect. 6.1, we review a number of techniques

¹ Of course, dynamic white-box techniques (for instance, based on symbolic execution) or other static analysis white-box techniques can also be used to generate models directly from code without using any active learning. Such models, however, will typically depend heavily on the internal structure of the SUL's program, and programs with the same observable behavior do not necessarily induce equivalent models.

that use white-box information about the SUL to reduce the time required to find discrepancies in a hypothesis model, or to prove the absence of such discrepancies.

2. In Sect. 6.2, we discuss extensions of Angluin’s MAT framework with new types of queries. A learner may for instance ask which previous values and operations have been used by the SUL to compute some output value. Or she may ask if some previous input value may subsequently be tested or output by the SUL. Such queries may dramatically simplify the task for the learner, but can often be simply answered by the teacher using off-the-shelf code analysis tools. An example would be a query about which registers are needed in a specific state or location.
3. Finally, access to the source code of a program or component can be used to compute information about the component that can help saving queries during learning. Information about which methods access internal variables and read or write those variables, e.g., can be used to decide whether queries can be reduced or even skipped. We discuss possible use cases for static code analysis in Sect. 6.3.

In order to make the discussion of the paper concrete, we will place it in a simple setting that is well-understood, namely *register automata*. Register automata (and the related *nominal automata*) constitute an extension of finite automata in which data values may be communicated, stored and manipulated. Recently, black-box learning algorithms have been generalized from finite automata to register automata [1, 15, 57].

Outline. We start by giving a brief overview to the field of model learning: In the next section, we provide a short introduction to the underlying learning setting and the practical challenges that arise when using automata learning for generating models of program behavior. Section 3 discusses related work and highlights cases in which learning has been applied successfully in practice. In the second half of the paper we develop several proposals for leveraging white-box analysis techniques in the concrete setting of learning register automaton models: We introduce register automata in Sect. 4 and we discuss existing learning algorithms for register automata in Sect. 5. We conclude by presenting our proposals in Sect. 6.

2 Inferring Models of Program Behavior

The general setting that we consider is illustrated in Fig. 1. We assume a SUL that execute some program in the set **Programs**. The semantics of programs is given by a function $beh : \mathbf{Programs} \rightarrow \mathbf{Behaviors}$ that describes the external, observable behavior of the SUL when it runs a program. Two programs $P, P' \in \mathbf{Programs}$ are deemed equivalent if they induce the same behavior: $P \equiv P' \Leftrightarrow beh(P) = beh(P')$. We postulate that each program $P \in \mathbf{Programs}$ can be described by a model $model(P)$ from some universe **Models**. In general, a program can be described by several models. The semantics of models is specified by a

Fig. 1. Learning setting.

function $beh_M : \text{Models} \rightarrow \text{Behaviors}$ and we assume $beh(P) = beh_M(model(P))$. Within the area of model-based testing, this assumption is commonly referred to as the *test hypothesis* [11, 31]. Two models M and M' are equivalent if they induce the same observable behavior: $M \equiv M' \Leftrightarrow beh_M(M) = beh_M(M')$.

Many instantiations of this general framework are possible. In the case of reactive systems, for instance, the set **Behaviors** may consist of functions $\lambda : \Sigma^* \rightarrow \Omega^*$ from sequences of inputs to sequences of outputs that preserve the prefix ordering (here denoted $\leq$) and the length of sequences, that is, for all $w, w' \in \Sigma^*$, $w \leq w' \Rightarrow \lambda(w) \leq \lambda(w')$ and $|\lambda(w)| = |w|$. In this case, the set **Models** naturally consists of (deterministic) Mealy machines with inputs Σ and outputs Ω . For reactive systems in which inputs and outputs do not alternate strictly, **Behaviors** may be defined as the class of prefix closed sets of *suspension traces* and **Models** as the class of *I/O transition systems* [69, 72].²

2.1 Model Learning

Model learning is relevant in situations where we do not know $model(P)$, either because we do not even know P , or because we know P but are somehow unable to compute $model(P)$ from P . In order to learn a model M with $M \equiv model(P)$, we postulate a collection **Queries** of *queries* that can be applied to the SUL. A function $resp : \text{Queries} \times \text{Programs} \rightarrow \text{Responses}$ specifies the result of applying a query to the SUL that is running some program. In the setting of Mealy machines, for instance, Angluin’s MAT framework uses I/O queries and equivalence queries. An I/O query consists of a sequence of inputs $w \in \Sigma^*$ and the response is given by

$$resp(w, P) = beh(P)(w).$$

² In fact, in the case of nondeterministic reactive systems, instantiations of our framework may be defined for each equivalence from the linear time—branching time spectrum of Van Glabbeek [33]. However, as far as we know, such instantiations have not yet been studied in the literature on model learning.

An equivalence query consists of an hypothesis $H \in \text{Models}$ and the response satisfies

$$\text{resp}(H, P) = \begin{cases} \text{yes} & \text{if } \text{beh}_M(H) = \text{beh}(P) \\ \text{no}, w & \text{otherwise, where } w \in \Sigma^* \text{ and } \text{beh}_M(H)(w) \neq \text{beh}(P)(w). \end{cases}$$

The queries of the MAT framework are *extensional*: in order to answer them we do not need access to the program, but only to its observable behavior. Formally, for all programs P and P' and queries Q ,

$$P \equiv P' \Rightarrow \text{resp}(Q, P) = \text{resp}(Q, P').$$

2.2 Abstraction

As stated in the introduction, the strength of model learning is that it can produce simple models of complex systems. This, of course, depends on the application of an appropriate abstraction. In the above description of model learning, such an abstraction is hidden in functions beh and beh_M . While in practice, beh_M usually is the semantics associated with the class of models that is inferred by some learning algorithm, the function beh abstracts the actual observable behavior of a program to the level of this semantics. Angluin's MAT framework, e.g., has been implemented for Mealy machine models over finite sets of inputs and outputs [46], where beh_M is a mapping from sequences of inputs to outputs. On the other hand, learning Mealy machine models of realistic software components requires a test harness which translates the abstract sequences of inputs to concrete sequences of method invocations on the component interface, and abstracts concrete return values of invocations to abstract outputs.

The choice of a class of models requires the existence of a learning algorithm for this class of models as well as the definition of a function beh that abstracts concrete program executions to traces in the semantics of this class of models. Defining such an appropriate abstraction beh oftentimes is not trivial as it is required to be deterministic and determines the aspects of a component's behavior that becomes observable.

The extension of learning algorithms to richer classes of models is an effort that has two positive impacts in this scenario: On the one hand, using more expressive classes of models can help representing more interesting aspects of a component's behavior in a model. On the other hand, using more expressive models can mitigate the laborious and often error-prone burden of defining appropriate functions beh .

This has led to multiple lines of works that extend Angluin's MAT framework to richer classes of models — most notably classes that can describe control-flow as well as data-flow or timing information. Extensions require finding right-congruences for more expressive classes of automata. One principal challenge that all these works face is that in a black-box setting, models can only be learned from observable behavior. Inferring complex causal relations like data manipulations or timed behavior quickly requires many queries and often has principle

Fig. 2. Models (left) and code (right) for a fictitious protocol component

limitations (e.g., the absence of a right-congruence) as has been shown for learning timed automata [36] models which cannot in general be determinized.

White-box access to a component can be beneficial for defining adequate *beh* functions, for reducing the number of required tests as well as for alleviating limitations on the expressivity of inferred models.

2.3 Example: A Protocol Component

We illustrate the above concepts using a small fictitious protocol component as an example. The C-code for the component is shown in the right half of Fig. 2. The component has two methods `msg(uint32_t s, int d)` and `ack()`, and one internal field `seq` that is initialized to 0. The `msg` method is guarded to accept only sequence numbers `s` in a certain window relative to `seq`. If the guard is satisfied, the internal field `seq` is updated to the value of `s`. The method then performs some operation on the payload data `d` that is not observable from the outside of the component (values of `d` are not stored and no observable error can occur while operating on `d`). The `ack` method returns the current value of `seq`.

When inferring a model of the behavior of the `WindowProtocol` component, a class of models and a corresponding function *beh* has to be fixed that abstracts the observable behavior of the component to the semantics of this class. The left half of Fig. 2 shows four models of the behavior of this component at different levels of abstraction, i.e., for different classes of models and different functions *beh*. For the purpose of illustration, the models only cover the behavior of the component under the `msg(uint32_t s, int d)` method. In each model, we consider a behavior to be a sequence of calls to the `msg(uint32_t s, int d)` method (at some level of abstraction), which does not trigger any failing assertion. The model itself represents the set of such behaviors.

The first model is the (huge) finite state automaton that results from using (abstract) inputs $msg(i)$ for all `ints` i and translating those to calls `msg(i, d)` with some random fixed `d`. Each state of the model represents one concrete valuation of the variable `seq` of the component. All states can be distinguished since they accept pairwise different sets of invocations of $msg(i)$ for Integer data values i . Only accepting locations are shown. While this model is finite and faithful to the behavior of the component, it is expensive to infer (due to the size) and of limited explanatory value.

The second model is a much more concise finite state machine that can be obtained by using the same learning algorithm as for the first model but with a much more involved function *beh* that basically models the implementation of the `msg` method and can concretize sequences of abstract $msg(IN_WINDOW)$ inputs to consecutive concrete method invocations with sequence numbers in the accepted respective windows. While the model is a perfect and concise documentation of the behavior of the component, it can only be inferred and interpreted with the help of an involved and state-dependent *beh* function.

The third model is a register automaton that can encode storing of method parameters (p) in registers (x) and compares parameters and sums of registers and constants for equality. We define register automata formally in Sect. 4. For this preliminary discussion, please note that the register automaton has a single accepting location that corresponds to all concrete accepting states in the first model. The model has one hundred transitions that loop from this location for all guards $x + i = p$ where $1 \leq i \leq 100$. Current learning algorithms for register automata models [15] would need two inputs for producing such a model: a grammar for terms allowed in equalities, e.g., $t ::= x + c \mid p$, and the set of allowed constants c .

Finally, in the fourth model, the one hundred transitions of the third model are merged into a single transition with a slightly more expressive guard, using inequalities instead of equalities and the single constant 101. As in the above case, a learning algorithm would need to be capable of inferring guards with inequalities and receive the constant as an input.

The example shows the potential for application of white-box information during model inference. Access to the source code of the component can be used, e.g., to identify methods that do not change the state of the component, to identify the required expressivity of guards or the necessary constants in an automated fashion. White-box analyses can also be used to determine, which parameters are stored in fields or to compute symbolic guards with fewer executions.

3 Related Work and Applications

Active automata learning gained a lot of traction over the past few years as a technique for inferring models of software components. One indication of the growing attention is a recent article on model learning in the Communications of the ACM [70]. In fact, the field and its applications have grown and diversified

to an extent that makes it impossible to provide a complete and comprehensive overview here. Instead, we try to sketch the lines of work that target increased expressivity of inferred models or the integration of white-box approaches. Additionally, we provide some examples of works that have shown positive results of using learned models for the (formal) analysis of components and systems.

From DFA to More Expressive Models. Dana Angluin presented the MAT model and a first learning algorithm (named L^*) in her seminal 1987 paper [6]. The L^* algorithm infers deterministic finite automata models of unknown regular languages. Hungar et al. presented the first learning algorithm for Mealy machine models in the MAT model [46]. Their work was motivated by the goal of producing more natural models of input/output behavior as well as reducing the cost of learning models.

Learning algorithms for Mealy machine models have been the basis for a line of works that construct beh_P functions (so-called *mappers*) for inferring models of infinite-state components. This approach is described explicitly for the generation of models from protocol entities in [3].

However, defining mappers is an error-prone and laborious manual effort. In [44] automated alphabet abstraction refinement is integrated with active learning to overcome this problem. More recent works extend this approach and combine automata learning with learning symbolic descriptions of transition structures, e.g., for models with large or infinite structured alphabets [56], or for alphabets that expose an algebraic structure [25].

The above approaches essentially still infer finite state machine models. Another different line of work aims at extending model learning to infinite state models in which states are defined over sets of variables. The authors of [10] present a technique for inferring *symbolic Mealy machines*, i.e., automata with guarded transitions and state-local sets of registers. The presented technique learns a Mealy machine over a large enough finite domain. In a post-processing step, from this Mealy machine a symbolic version is constructed.

Howar et al. extend active automata learning in the MAT model to register automata, which model control-flow as well as data-flow between data parameters of inputs, outputs, and a set of registers [43]. Registers and data parameters can be compared for equality. The authors demonstrate the effectiveness of their approach by inferring models of data structures [42] and extend the expressivity to allow for arbitrary data relations that meet certain learnability criteria [14, 15]. Aarts et al. develop a slightly different approach for inferring register automata models that can compare registers and data parameter for equality [1, 5]. The two approaches are compared in [2].

A more in-depth overview of works that extend active automata learning to infinite state models is provided in [49].

Applications. After Peled et al. suggested active automata learning for making black-box systems amenable to the application of formal methods [60], Hagerer et al. pioneered the application of active automata learning for generating models of components; components of a computer telephony system in their particular case [38, 39]. The models were used as a basis for testing the system. In recent

years, generating models for testing has been continued by different authors for several types of systems, e.g., for Event-B models [23,24], for graphical user interfaces of android applications [18,19], and for integration testing [64] of automotive components. Meinke and Sindhu present LBTest, a tool for learning-based testing for reactive systems, integrating model checking, active automata learning, and random testing [55].

Other applications target generating behavioral specifications of Web applications [61], the new biometric European passport [4], bot nets [12], and enterprise applications [73]. Margaria et al. showed that model learning may help to increase confidence that a legacy component and a refactored implementation have the same behavior [53]. Inspired by this work, Schuts et al. use inferred specifications and equivalence checking to assist re-engineering of legacy software in an industrial context at Philips [63]. Sun et al. use active automata learning in combination with automated abstraction refinement and random testing for finding abstract behavioral models of Java classes [67].

An emerging area of applications is the learning-based analysis of safety or security of components and systems: De Ruiter and Poll use active automata learning for inferring models of TLS implementations and discover previously unknown security flaws in the inferred models [62]. Xue et al. use active automata learning for inferring behavioral models of JavaScript malware [74]. Fiterau et al. use learning and model checking to analyze the behavior of different implementations of the TCP protocol stack and document several instances of implementations violating RFC specifications [29]. Using a similar approach, Fiterau et al. show that also three different implementations of the SSH protocol violate the RFC specifications [30]. Khalili et al. [51] use active automata learning to obtain behavioral models of the middleware of a robotic platform. The models are used during verification of control software for this platform.

Integration of White-Box Techniques. In [54], Margaria et al. investigate the potential of what they call “domain-specific knowledge” for reducing the cost of learning models. Their domain-specific knowledge, e.g., assumptions about prefix-closedness of an unknown target language or the independence of inputs, is a first example of the kind of information about a system that can be computed by white-box techniques.

The first works that actually used white-box techniques to implement more powerful queries explore combinations of active automata learning and different forms of symbolic execution for producing expressive models of components. Giannakopoulou et al. develop an active learning algorithm that infers safe interfaces of software components with guarded actions. In their model, the teacher is implemented using concolic execution [32]. Cho et al. present MACE an approach for concolic exploration of protocol behavior. The approach uses active automata learning for discovering so-called deep states in the protocol behavior. From these states, concolic execution is employed in order to discover vulnerabilities [17]. Botinčan and Babić present a learning algorithm for inferring models of stream transducers that integrates active automata learning with symbolic execution and counterexample-guided abstraction refinement [13]. They show how

the models can be used to verify properties of input sanitizers in Web applications. Their work is the first in this line of works that produces infinite-state models that can store data values in a set of registers. Finally, Howar et al. extend the work of [32] and integrate knowledge about the potential effects of component method invocations on a component's state to improve the performance during symbolic queries [45]. The knowledge is obtained through static code analysis.

Another (earlier) line of work uses active learning in model-checking contexts. The moderate style of exploration that achieved by learning is used to mitigate the problem of state space explosion (e.g. [22]). Recent advances in this area have been made by finding active automata learning for expressive classes of models. Learning algorithms are usually based quite directly on the classic L^* algorithm. The required extensions in expressivity of models are usually realized through powerful teachers. For instance, Feng et al. present an algorithm for inferring assumptions for probabilistic assume/guarantee reasoning [26, 27].

4 Register Automata

In order to make the discussion of the paper concrete, we will place it in a setting that is well-understood, namely *register automata*. These extend finite automata with data values that may be communicated, stored and manipulated. In this section, we introduce basic definitions of data languages and register automata that generalize corresponding concepts for languages and finite automata.

In our setting, data languages and register automata are parameterized on a vocabulary that determines how data can be examined, which in our setting is called a *theory*. A *theory* is a pair $\langle \mathcal{D}, \mathcal{R} \rangle$ where $\mathcal{D}$ is an unbounded domain of *data values*, and $\mathcal{R}$ is a set of *relations* on $\mathcal{D}$. The relations in $\mathcal{R}$ can have arbitrary arity. Known constants can be represented by unary relations.

Examples of simple theories include

- $\langle \mathbb{N}, \{=\} \rangle$, the theory of natural numbers with equality; instead of the set of natural numbers, we could consider any other unbounded domain, e.g., the set of strings (representing passwords or usernames).
- $\langle \mathbb{R}, \{<\} \rangle$, the theory of real numbers with inequality: this theory also allows to express equality between elements.

Operations, such as increments, addition and subtraction, can in this framework be represented by relations. For instance, addition can be represented by a ternary relation $p_1 = p_2 + p_3$. In the following, we assume that some theory $\langle \mathcal{D}, \mathcal{R} \rangle$ has been fixed.

Data Languages. We assume a set Σ of *actions*, each with an arity that determines how many parameters it takes from the domain $\mathcal{D}$. For notational convenience, we will here assume that all actions have arity 1. A *data symbol* is a term of form $\alpha(d)$, where α is an action and $d \in \mathcal{D}$ is a data value. A *data word* is a sequence of data symbols. The concatenation of two data words w and w' is

denoted ww' . Two data words $w = \alpha_1(d_1) \dots \alpha_n(d_n)$ and $w' = \alpha_1(d'_1) \dots \alpha_n(d'_n)$ are $\mathcal{R}$ -indistinguishable, denoted $w \approx_{\mathcal{R}} w'$, if they have the same sequence of actions, and $R(d_{i_1}, \dots, d_{i_j}) \Leftrightarrow R(d'_{i_1}, \dots, d'_{i_j})$ whenever R is a j -ary relation in $\mathcal{R}$ and $i_1, \dots, i_j$ are indices among $1 \dots n$, i.e., the sequence of data values cannot be distinguished by any of the relations in $\mathcal{R}$.

A *data language* $\mathcal{L}$ is a set of data words that respects $\mathcal{R}$ in the sense that $w \approx_{\mathcal{R}} w'$ implies $w \in \mathcal{L} \Leftrightarrow w' \in \mathcal{L}$. We will often represent data languages as mappings from the set of data words to $\{+, -\}$, where $+$ stands for ACCEPT and $-$ for REJECT.

Register Automata. We assume a set of *registers* $x_1, x_2, \dots$. A *parameterized symbol* is a term of form $\alpha(p)$, where α is an action and p a formal parameter. A *guard* is a conjunction of negated and unnegated relations (from $\mathcal{R}$) over the formal parameter p and registers. An *assignment* is a simple parallel update of registers with values from registers or the formal parameter p . We represent an assignment which updates the registers $x_{i_1}, \dots, x_{i_m}$ with values from the registers $x_{j_1}, \dots, x_{j_n}$ or p as a mapping π from $\{x_{i_1}, \dots, x_{i_m}\}$ to $\{x_{j_1}, \dots, x_{j_n}\} \cup \{p\}$, meaning that the value of the register or formal parameter $\pi(x_{i_k})$ is assigned to the register x_{i_k} , for $k = 1, \dots, m$. Using multiple-assignment notation, this would be written as $x_{i_1}, \dots, x_{i_m} := \pi(x_{i_1}), \dots, \pi(x_{i_m})$.

Definition 1 (Register automaton). A register automaton (RA) is a tuple $A = (L, l_0, \mathcal{X}, \Gamma, \lambda)$, where

- L is a finite set of locations, with $l_0 \in L$ as the initial location,
- $\mathcal{X}$ maps each location $l \in L$ to a finite set $\mathcal{X}(l)$ of registers, and
- Γ is a finite set of transitions, each of form $\langle l, \alpha(p), g, \pi, l' \rangle$, where
 - $l \in L$ is a source location,
 - $l' \in L$ is a target location,
 - $\alpha(p)$ is a parameterized symbol,
 - g is a guard over p and $\mathcal{X}(l)$, and
 - π (the assignment) is a mapping from $\mathcal{X}(l')$ to $\mathcal{X}(l) \cup \{p\}$, and
- λ maps each $l \in L$ to $\{+, -\}$. □

Register automata are required to be *completely specified* in the sense that for each location $l \in L$ and action α , the disjunction of the guards on the α -transitions from l is equivalent to *true*.

A restriction of register automata, as defined by Definition 1, is that transitions do not allow to assign arbitrary expressions to registers, only the value of a formal parameter or a register. A main reason for this restriction is to limit the number of possibilities for inferring guards and assignments that match the result of membership queries. As an example, suppose that a SUL accepts sequences with increasing parameter values, e.g., *offer*(4) *offer*(5) *offer*(6) *offer*(7). We could then learn a RA if the theory includes, e.g., the relation *issucc*, defined by *issucc*(x, y) iff $x + 1 = y$. If assignments to registers would allow expressions that include e.g., the $+1$ operator, or even arbitrary addition, then the learning algorithm would have to choose between a potentially large number of different

guards and assignments on each transition, This would complicate the design of a learning algorithm. On the other hand, we do not foresee any fundamental difficulty in extending the theory for learning RAs in order to produce more expressive classes of RAs; we conjecture that this could be done by making the implementation of tree queries more advanced and extending the Nerode equivalence (cf. Sect. 5.2). However, in order to focus on the conceptual extensions needed to learn RAs, we have so far excluded expressions in assignments of RAs.

The semantics of register automata is defined in the standard way. Let $\mathcal{A}$ be an RA $\mathcal{A} = (L, l_0, \mathcal{X}, \Gamma, \lambda)$. A *state* of $\mathcal{A}$ is a pair $\langle l, \nu \rangle$ consisting of a location l and a *valuation* $\nu : \mathcal{X}(l) \rightarrow \mathcal{D}$ of the registers $\mathcal{X}(l)$ of that location. A *step* of $\mathcal{A}$, denoted $\langle l, \nu \rangle \xrightarrow{\alpha(d)} \langle l', \nu' \rangle$, transfers $\mathcal{A}$ from state $\langle l, \nu \rangle$ to state $\langle l', \nu' \rangle$ on input of the data symbol $\alpha(d)$ if $\mathcal{A}$ has a transition $\langle l, \alpha(p), g, \pi, l' \rangle \in \Gamma$, whose guard is satisfied by d under valuation ν (i.e., $\nu \models g[d/p]$), and whose assignment produces the updated valuation ν' (i.e., $\nu'(x_i) = \nu(x_j)$ if $\pi(x_i) = x_j$, otherwise $\nu'(x_i) = d$ if $\pi(x_i) = p$). A *run* of $\mathcal{A}$ over a data word $w = \alpha(d_1) \dots \alpha(d_n)$ is a sequence of steps of $\mathcal{A}$

$$\langle l_0, \nu_0 \rangle \xrightarrow{\alpha_1(d_1)} \langle l_1, \nu_1 \rangle \quad \dots \quad \langle l_{n-1}, \nu_{n-1} \rangle \xrightarrow{\alpha_n(d_n)} \langle l_n, \nu_n \rangle$$

for some initial valuation ν_0 . The run is *accepting* if $\lambda(l_n) = +$ and *rejecting* if $\lambda(l_n) = -$. The word w is *accepted (rejected)* by $\mathcal{A}$ under ν_0 if $\mathcal{A}$ has an accepting (rejecting) run over w which starts in $\langle l_0, \nu_0 \rangle$.

Existing techniques for active learning of register automata only consider RAs that are *determinate*, meaning that there is no data word over which it has both accepting and rejecting runs. A nondeterministic but determinate RA can be easily transformed into a deterministic RA by strengthening its guards. Note that, unlike for finite automata, nondeterministic (and nondeterminate) RAs are strictly more expressive than deterministic RAs (for instance, the universality problem for nondeterministic RAs is undecidable [59]).

Fig. 3. Priority queue with capacity 2.

Example. As a simple example of a determinate register automaton, let us consider a priority queue with capacity two. A priority queue stores a set of keys from some totally ordered set. We will use the set of rational numbers as the set of keys. We abstract from values that are stored along with keys. The interface of the priority queue supports two operations:

- *offer* inserts a given key into the priority queue. It succeeds if the queue is not full;
- *poll* asks for the smallest key in the queue; the operation returns that key and removes it; if the queue contains several copies of the smallest key only one is removed. If the queue is empty, the operation does not succeed.

The interface consists of operations with input parameters and return values. In order to represent it as a data language, we model sequences of successful operations as data words. A successful *offer* is represented by the data symbol $offer(d)$, where d is the inserted key. A successful *poll* operation is represented by the data symbol $poll(d)$, where d is the returned value. A valid sequence of operations is represented by the sequence of data symbols that represent its successful operations.

The RA in Fig. 3 accepts the language which models a priority queue with capacity two. The two keys are stored in registers x_1 and x_2 , respectively, with $x_1 \leq x_2$, so that a successful *poll* operation always returns the value of x_1 .

For conciseness, we have omitted nonaccepting locations. Thus the RA in Fig. 3 should be extended with a terminal non-accepting location; from each location, there should be transitions to the non-accepting location for data symbols that do not satisfy any of the existing guards. For instance, from l_1 there is a transition to the non-accepting location for $poll(p)$ symbols where $p \neq x_1$.

5 Black-Box Learning of Register Automata

In this section, we summarize some of the proposed algorithms for learning register automata. A number of such algorithms have been proposed, which generalize the classical L^* algorithm in some way. The obvious challenge for such algorithms is that register automata is a much richer formalism than finite automata. It is a challenge to devise techniques that can infer all the features of an RA, including locations, registers, guards, and assignments, in a black-box context, using only membership queries and counterexamples returned by equivalence queries. The only *a priori* information available is the static interface of the SUL, i.e., the set of actions that it can process, and a theory (i.e., set of relations on the data domain) which is assumed to be expressive enough to model the behavior of the SUL.

5.1 Learning Symbolic Automata

Let us first consider the subclass of *symbolic automata*, which are essentially register automata without registers and assignments. Symbolic automata have been studied in recent years as a tool for string processing, where transitions are guarded by predicates of various theories [40, 71]. They allow automata over large or infinite alphabets to be expressed compactly.

For large finite domains of data values, symbolic automata are equivalent to ordinary finite automata. However, a naive application of L^* to such automata

will lead to an excessive number of membership queries, since queries must be performed for each data value in the data domain. More efficient approaches have been presented by Isberner et al. [47] and by Mens and Maler [56]. A main idea is to associate to each transition a *representative* data value which satisfies its guard. The hypothesis is that all data values that satisfy the guard induce the same behavior in the SUL; formally this means that they lead to the same residual languages. When such a hypothesis is refuted, typically as a result of an equivalence query (i.e., two data values that satisfy the guard lead to different residual languages), the algorithm splits the guard accordingly. In [47, 56], some predefined structure for splitting guards by need is assumed. For instance, when the data domain is the set of integers, the algorithm could prescribe that all membership queries initially be performed using a specific integer (e.g., 0). If later, it turns out that different behavior is induced by a negative number, a transition may be split into two, one for nonnegative and one for negative numbers.

The approach of [47] has been shown to improve over naive L^* by several orders of magnitude on a set of benchmarks of moderate size. So far, the approach has not been applied to learn symbolic automata of the form considered in, e.g., [40, 71] with its rich collection of predicates.

5.2 Learning Register Automata: The SL^* Algorithm

The class of register automata with registers and assignments brings additional challenges to the design of learning algorithms. For symbolic automata, which do not have registers, the learning algorithm can still be based on the classical definition of Nerode congruence for identifying locations. That is, two data words are Nerode congruent if they induce the same future behavior (i.e., residual language), and each congruence class corresponds to a location. In this case, the main problem is to infer, for each state, a suitable partitioning of data symbols that processed as input in that state, and map the partitioning onto guards on outgoing transitions. For register automata, the future behavior from a location (i.e., its residual language) depends also on the data values assigned to its registers. A learning algorithm must thus infer both (i) which registers are needed in a location, and (ii) how to partition the set of data values in input data symbols to produce guards on different outgoing transitions, in a way that depends on the register valuation. Thus, the concept of residual language must be generalized to a mapping from register valuations to future behaviors. Furthermore, since assignments can permute registers, the concept of Nerode congruence must be defined in such a way that it allows permutations of registers: different permutations will result in different assignments to registers on incoming transitions.

As an illustration, the future behavior from location l_2 in Fig. 3 depends on two data values. A learning algorithm will infer that a word leading to l_2 has two *memorable* data values; intuitively, an input value d is memorable if it has an impact on the future behavior of the SUL: either d occurs in a future output, or a future output depends on d (e.g., if d appears in a guard). A learning algorithm will therefore create two registers: x_1 and x_2 . Assume that location l_1 is

represented by the word *offer*(3). Initially, a learning algorithm may assume that all outgoing *offer*-transitions from l_1 can be represented by a single symbol, say *offer*(5), and generate a single outgoing transition from l_1 with the guard *true*. A subsequent counterexample, e.g., of form, *offer*(3)*offer*(1)*poll*(1) will make the learning algorithm realize that the future behavior after *offer*(3)*offer*(1) is equivalent to that after *offer*(3)*offer*(5), if we are allowed to adapt the contents of registers so that x_1 is assigned the smaller value and x_2 the larger value. The learning algorithm will therefore split the transition guarded by *true* into two different transitions, with a guard that compares the parameter of the current data symbol to the value of the register, as in Fig. 3, and equip each transition with a different assignment.

The above concepts are a basis for the SL^* algorithm for learning register automata [15]. It extends predecessor algorithms such as [6, 47, 56] by the concept of *tree queries*, which are used in place of membership queries. The arguments of a tree query is a prefix data word, and a set of so-called *symbolic suffixes*, i.e., data words with uninstantiated data parameters. The tree query returns a so called *symbolic decision tree* (SDT), which has the form of an “RA-fragment” which accepts/rejects suffixes obtained by instantiating data parameters in one of the symbolic suffixes.

Fig. 4. Symbolic decision tree returned by tree query for prefix *offer*(3) and symbolic suffixes $\{\textit{offer}(p), \textit{poll}(p)\}$

Let us illustrate this on the priority queue example for the prefix *offer*(3) and the set of symbolic suffixes $V = \{\textit{offer}(p), \textit{poll}(p)\}$. The acceptance/rejection of suffixes obtained by instantiating data parameters in V after *offer*(3) can be represented by the SDT in Fig. 4. In the initial location, the value 3 from *offer*(3) is stored in a register. We use the convention that register x_i stores the i th data value from the prefix. Thereafter, suffixes of form *poll*(d) are accepted if the data value d equals the value stored in the register, and rejected otherwise (for each action there is an implicit transition to a rejecting location with a guard that is the conjoined negations of all accepting transitions). Suffixes of form *offer*(d) are always accepted.

Tree queries can be implemented by performing several membership queries and combining their results. A straightforward implementation of a tree query for a prefix u and an uninstantiated suffix of form $\alpha(p)$ produces a set of maximally refined but still satisfiable guards over registers x_i , storing values from u , and

suffix parameter p without introducing additional constraints between registers. For each such maximally refined satisfiable guard in the theory a membership query for the concatenation of u and $\alpha(d)$ for some d that satisfies the guard is sufficient to characterize the behavior for all data values satisfying this guard. For instance, the SDT of Fig. 4 can be formed by combining the results of the membership queries $offer(3)offer(1)$ for guard $(x_1 > p)$, $offer(3)offer(3)$ for guard $(x_1 = p)$, $offer(3)offer(5)$ for guard $(x_1 < p)$, and analogous membership queries for suffix $poll(p)$.

One reason for using the tree queries instead of simple membership queries is that they allow to properly approximate a Nerode equivalence under which words can be equivalent after permutation of data values. Suppose, for instance, that we would only pose the membership query $offer(3)offer(5)poll(3)$, which is accepted; we then use the principles of [47, 56] to infer that after $offer(3)offer(5)$ all symbols of form $poll(p)$ are accepted. If then the membership query $offer(3)offer(1)poll(3)$ is rejected, we infer that after $offer(3)offer(1)$ all symbols of form $poll(p)$ are rejected. We would then conclude that the prefixes $offer(3)offer(5)$ and $offer(3)offer(1)$ are not Nerode equivalent and represent different locations, although in the final automaton (see Fig. 3) they are represented by the same location. Thus, our preliminary approximation of the Nerode equivalence would *not* overapproximate the “actual” Nerode equivalence. This would destroy the partition-refinement structure of the L^* algorithm, which in turn is the basis for establishing convergence guarantees of the learning algorithm.

5.3 Limitations and Extensions

Section 5.2 outlined the principles of a framework for extending active automata learning to register automata [15]. This framework has been implemented in RALib [14], and used to infer register automata and register mealy machines with simple theories, such as theories of numbers with equality and inequality. In order to make it more generally usable, some limitations must be overcome, of which we list some here.

- **Richer theories with structured data types:** The simple framework outlined in the preceding sections can be used to infer register automata models, whose registers are assigned scalar values, such as the priority queue model in Fig. 3. Obviously, this model structure does not scale to modeling priority queues of arbitrary size. For this, we need to work with theories that can describe structure data, such as sequences, and operations on such sequences. It remains to be seen whether the overall framework for learning RAs outlined in this section will also function well with such structured data type.
- **Scalability:** The realization of the ideas is still somewhat naive. For instance, the number of membership queries used to realize a tree query is exponential in the length of the prefix (and suffix) of the query. As explained in Sect. 5.2 this is motivated by the desire to stay within a partition refinement framework, but there is still much room for optimizing this aspect of the algorithm.

One way to address these limitations is to augment the learning algorithms with white-box information extraction methods. Some directions will be discussed in the next section.

6 Exploiting White-Box Techniques

We have discussed the potentially positive impact of exploiting white-box techniques in automata learning in previous sections and have sketched the current limitations and open questions when learning register automata models. We conclude the paper by discussing directions for future research with a particular focus on using white-box techniques to improve learning of expressive register automata models while retaining the benefits of the black-box approach, i.e., concise and abstract models.

6.1 Improving the Equivalence Oracle

Black-box learning approaches, although effective in constructing hypothesis models for finite state machines, typically have difficulties to find counterexamples for hypotheses with a large number of states and events [65]. If we have direct access to the code or binary of the SUL, several additional techniques become available to discover counterexamples for hypothesis models. There is a range of white-box symbolic execution techniques, such as veritesting [7], concolic testing [35], and white-box fuzz testing [34] that can be adapted to find counterexamples for hypothesis models. We survey some works that exploit this idea.

Smetsers et al. [66] used American fuzzy lop (AFL)³ to efficiently obtain counterexamples for hypothesis models. AFL is a fuzzer that uses compile-time instrumentation and genetic algorithms to automatically discover test cases that trigger new internal states in the targeted binary. By combining model learning with AFL, Smetsers et al. were successful in the RERS 2016 challenge⁴, which aims to compare verification techniques and tools.

The PSYCO tool integrates active automata learning and dynamic symbolic execution for generating component interfaces [32]. Recent work explores the potential of using symbolic search on a component's state space [58] for estimating an upper bound on the length of counterexamples that can be found during learning by PSYCO. Fully symbolic and synchronized exploration of a SUL and a conjectured model, i.e., checking the conjecture against the SUL, would allow it to decide equivalence (assuming decidability of the corresponding SMT encoding). The white-box learning algorithm Σ^* [13] uses predicate abstraction to construct models that overapproximate stream processing code in order to answer equivalence queries.

³ <http://lcamtuf.coredump.cx/afl/>.

⁴ <http://www.rers-challenge.org/2016>.

6.2 Introducing New Queries

In a black-box setting, one important reason why many membership queries are needed is that the learner cannot see directly whether a value is stored by the SUL, or whether it is compared to other values in the guard of a transition. The task of the learner is to figure this out using black-box experiments only. In [15], learning algorithms are presented that accomplish this task for some simple theories with equality and/or inequality. These algorithms have a high query complexity, however, and it is not clear how they can be generalized to richer theories. Consider, for instance, a trace *offer*(2) *offer*(3) *offer*(4) *offer*(4) in a setting with equality, successor and addition. What guard was used on the last transition? Is the last value required to be equal to the previous one? Or is it the successor of the second value ($3 + 1 = 4$)? Or has it been obtained by addition from the first value ($2 + 2 = 4$)? A number of works use forms of symbolic execution for making symbolic constraints on execution paths visible when performing tests on a component during membership queries. In this way, a learner can observe directly which values are stored and which predicates are tested in a trace. The idea is to replace membership queries (is word w in the language?) by a symbolic version in which the reply consists not only of a yes/no answer but also includes the complete symbolic run induced by input word w . The white-box learning algorithm Σ^* [13] is an example of an approach in which the learner may pose “symbolic” membership queries.

A lightweight alternative for the use of symbolic executions is provided by taint analysis. Dynamic taint analysis (also referred to as dynamic information flow tracking) [8, 20, 21, 41] is a technique in which code is instrumented in order to mark and track data in a program at run time. Inputs to a program are “tainted”, i.e. labeled with a tag. When the program executes, these tagged values are propagated such that any new values derived from a tagged value also carry a tag derived from these source input tags. Dynamic tainting, as implemented for instance in Autogram [41], allows to precisely identify which program inputs are stored, tested, or processed at any point in time.

In order to see how a learner may use tainting information, consider the priority queue example of Fig. 3. Suppose the learner is using the SL^* algorithm of Sect. 5.2 and needs to construct a symbolic decision tree for the prefix *offer*(3) and the set of symbolic suffixes $V = \{\textit{offer}(p), \textit{poll}(p)\}$. Suppose that taint analysis reveals that the parameter of a call to *offer* is stored in some variable, say x_1 , and that the return value of a subsequent call to *poll* is equal to the value of this variable. Based on this information, the learner may deduce the right branch of Fig. 5. After calls *offer*(3) *offer*(4), taint analysis tells that the first parameter is stored in x_1 , the second parameter is compared with x_1 in a test $p > x_1$, and then stored in x_2 . Based on this, the learner deduces the left branch of Fig. 5, and infers that the tree is still incomplete. In order to complete the tree, the learner performs calls *offer*(3) *offer*(2), where 2 is an arbitrary value less than 3. Using taint analysis, the learner can now infer the third and final branch in the tree.

Fig. 5. Symbolic decision tree constructed using dynamic taint analysis

It is interesting to compare the SDTs of Figs. 4 and 5. Whereas a black-box tree query initially returns the incorrect decision tree of Fig. 4 for prefix $offer(3)$, and several additional tree queries (and even an equivalence query) are required before a correct decision tree is found, a white-box tree query produces the correct decision tree of Fig. 5 right away. This benefit comes at a price, as application of taint analysis requires instrumentation of the code. According to [41], their tainting framework for Java programs (which is not yet optimized) runs approximately 100 times slower than the original, uninstrumented code.

A key benefit of taint analysis is that it enables the application of the SL^* algorithm in settings with richer theories. The SL^* algorithm crucially depends on an oracle that answers tree queries. In [15] it is shown how a tree oracle can be implemented via (black-box) membership queries for some commonly occurring theories: the theory of equalities, the theory of equality and inequality over rational (or real) numbers, and the theory of equality and inequality over integers. Every single tree query is mapped to an exponential number of membership queries (in the number of data values in a queried word w) that identify relevant relations between data values in w . Moreover, these tree oracle implementations are nontrivial and their correctness proofs are involved. Computation of decision trees and counterexample analysis becomes harder (or even impossible) when more relations are added to the theory.

6.3 Computing Domain-Specific Information

The state of a component is usually maintained as valuations of internal variables. A static code analysis that produces, e.g., for a Java class, which methods write internal variables, can help identifying methods that cannot alter the state of a component before actually learning a model of the component. A comparative analysis (reads and writes) can identify pairs of independent methods for which the order of execution is irrelevant. Information of this kind can be used to reduce the number of membership queries as has been shown in [45, 54].

When inferring register automata, static code analysis can reveal important information about a component.

Live Variables. Knowing about live variables can help deciding if values have to become memorable or not in certain states without performing extensive testing or expensive symbolic analysis.

Tests and Computations on Variables. Knowing that only certain tests are performed on a variable can help when determining the theory that is used for learning the behavior that is possibly associated with this variable. This can help reducing the number of tests that have to be performed when learning models. On the other hand, knowledge about the set of computations that can be used on parameters or register contents may enable learning more expressive models that, e.g., describe application of cryptographic primitives on parameters or internal fields.

Independent Parameters/Fields. If there are subsets of parameters and fields that are independent, a learning algorithm does not have to perform tests for inferring potential relations between these sets. The typing mechanism that is presented in [14] allows to exploit this information during learning. Static code analysis could be used to compute it in a fully automated approach.

The above list is not exhaustive but rather a starting point for future research.

In all three directions, initial positive results exist and indicate the potential that lies in the application of white-box techniques for the implementation of more performant learning algorithms, inferring more expressive classes of models.

7 Conclusion

We have outlined current techniques and applications for model learning, a.k.a. active automata learning, and pointed at challenges for improving its scalability and applicability to richer models. We then outlined proposals for exploiting white-box techniques in order to overcome these limitations. We indicated some approaches that have started in these directions, and we expect a significant body of techniques to be developed over the next years.

References

1. 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
2. Aarts, F., Howar, F., Kuppens, H., Vaandrager, F.: Algorithms for inferring register automata. In: Margaria, T., Steffen, B. (eds.) ISoLA 2014. LNCS, vol. 8802, pp. 202–219. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45234-9_15
3. Aarts, F., Jonsson, B., Uijen, J., Vaandrager, F.W.: Generating models of infinite-state communication protocols using regular inference with abstraction. *Formal Methods Syst. Des.* **46**(1), 1–41 (2015)
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. Aarts, F., Heidarian, F., Vaandrager, F.: A theory of history dependent abstractions for learning interface automata. In: Koutny, M., Ulidowski, I. (eds.) CONCUR 2012. LNCS, vol. 7454, pp. 240–255. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32940-1_18
6. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* **75**(2), 87–106 (1987)
7. Avgerinos, T., Rebert, A., Cha, S.K., Brumley, D.: Enhancing symbolic execution with veritesting. In: Proceedings of the 36th International Conference on Software Engineering, ICSE 2014, pp. 1083–1094. ACM, New York (2014)
8. Bell, J., Kaiser, G.: Phosphor: illuminating dynamic data flow in commodity JVMs. In: Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, pp. 83–101. ACM, New York (2014)
9. Berg, T., Grinchtein, O., Jonsson, B., Leucker, M., Raffelt, H., Steffen, B.: On the correspondence between conformance testing and regular inference. In: Cerioli, M. (ed.) FASE 2005. LNCS, vol. 3442, pp. 175–189. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-31984-9_14
10. Berg, T., Jonsson, B., Raffelt, H.: Regular inference for state machines using domains with equality tests. In: Fiadeiro, J.L., Inverardi, P. (eds.) FASE 2008. LNCS, vol. 4961, pp. 317–331. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78743-3_24
11. Bernot, G., Gaudel, M.C., Marre, B.: Software testing based on formal specifications: a theory and a tool. *Softw. Eng. J.* **6**(6), 387–405 (1991)
12. Bossert, G., Hiet, G., Henin, T.: Modelling to simulate botnet command and control protocols for the evaluation of network intrusion detection systems. In: Proceedings of the 6th Conference on Network and Information Systems Security, SAR-SSI 2011, pp. 1–8. IEEE Computer Society (2011)
13. Botinčan, M., Babić, D.: Sigma*: symbolic learning of input-output specifications. In: Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2013, pp. 443–456. ACM, New York (2013)
14. Cassel, S., Howar, F., Jonsson, B.: RALib: a LearnLib extension for inferring EFSMs. In: International Workshop on Design and Implementation of Formal Tools and Systems, DIFTS 2015, Austin, Texas (2015). http://www.faculty.ece.vt.edu/chaowang/difts2015/papers/paper_5.pdf
15. Cassel, S., Howar, F., Jonsson, B., Steffen, B.: Active learning for extended finite state machines. *Formal Asp. Comput.* **28**(2), 233–263 (2016)
16. Chalupar, G., Peherstorfer, S., Poll, E., de Ruiter, J.: Automated reverse engineering using Lego. In: Proceedings of 8th USENIX Workshop on Offensive Technologies (WOOT 2014), San Diego, California, Los Alamitos, CA, USA, August 2014. IEEE Computer Society (2014)
17. Cho, C.Y., Babić, D., Poosankam, P., Chen, K.Z., Wu, E.X., Song, D.: Mace: model-inference-assisted concolic exploration for protocol and vulnerability discovery. In: Proceedings of the 20th USENIX Conference on Security, SEC 2011, Berkeley, CA, USA, p. 10. USENIX Association (2011)
18. Choi, W.: Automated testing of graphical user interfaces: a new algorithm and challenges. In: Proceedings of the 2013 ACM Workshop on Mobile Development Lifecycle, MobileDeLi 2013, pp. 27–28. ACM, New York (2013)
19. Choi, W., Necula, G., Sen, K.: Guided GUI testing of android apps with minimal restart and approximate learning. *SIGPLAN Not.* **48**(10), 623–640 (2013)

20. Clause, J., Li, W., Orso, A.: Dytan: a generic dynamic taint analysis framework. In: Proceedings of the 2007 International Symposium on Software Testing and Analysis, ISSTA 2007, pp. 196–206. ACM, New York (2007)
21. Clause, J., Orso, A.: Penumbra: automatically identifying failure-relevant inputs using dynamic tainting. In: Proceedings of the Eighteenth International Symposium on Software Testing and Analysis, ISSTA 2009, pp. 249–260. ACM, New York (2009)
22. Cobleigh, J.M., Giannakopoulou, D., Pasareanu, C.S.: Learning assumptions for compositional verification. In: Garavel, H., Hatcliff, J. (eds.) TACAS 2003. LNCS, vol. 2619, pp. 331–346. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-36577-X_24
23. Dinca, I., Ipate, F., Mierla, L., Stefanescu, A.: Learn and test for event-B – a rodin plugin. In: Derrick, J., Fitzgerald, J., Gnesi, S., Khurshid, S., Leuschel, M., Reeves, S., Riccobene, E. (eds.) ABZ 2012. LNCS, vol. 7316, pp. 361–364. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30885-7_32
24. Dinca, I., Ipate, F., Stefanescu, A.: Model learning and test generation for event-B decomposition. In: Margaria, T., Steffen, B. (eds.) ISO/LA 2012. LNCS, vol. 7609, pp. 539–553. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34026-0_40
25. Drews, S., D’Antoni, L.: Learning symbolic automata. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10205, pp. 173–189. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54577-5_10
26. Feng, L., Han, T., Kwiatkowska, M., Parker, D.: Learning-based compositional verification for synchronous probabilistic systems. In: Bultan, T., Hsiung, P.-A. (eds.) ATVA 2011. LNCS, vol. 6996, pp. 511–521. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24372-1_40
27. Feng, L., Kwiatkowska, M., Parker, D.: Automated learning of probabilistic assumptions for compositional reasoning. In: Giannakopoulou, D., Orejas, F. (eds.) FASE 2011. LNCS, vol. 6603, pp. 2–17. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19811-3_2
28. Fiterău-Broștean, P., Howar, F.: Learning-based testing the sliding window behavior of TCP implementations. In: Petrucci, L., Seceleanu, C., Cavalcanti, A. (eds.) FMICS/AVoCS -2017. LNCS, vol. 10471, pp. 185–200. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67113-0_12
29. 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
30. Fiterău-Broștean, P., Lenaerts, T., Poll, E., de Ruiter, J., Vaandrager, F., Verleg, P.: Model learning and model checking of SSH implementations. In: Proceedings of the 24th ACM SIGSOFT International SPIN Symposium on Model Checking of Software, SPIN 2017, pp. 142–151. ACM, New York (2017)
31. Gaudel, M.-C.: Testing can be formal, too. In: Mosses, P.D., Nielsen, M., Schwartzbach, M.I. (eds.) CAAP 1995. LNCS, vol. 915, pp. 82–96. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-59293-8_188
32. Giannakopoulou, D., Rakamarić, Z., Raman, V.: Symbolic learning of component interfaces. In: Miné, A., Schmidt, D. (eds.) SAS 2012. LNCS, vol. 7460, pp. 248–264. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33125-1_18
33. van Glabbeek, R.J.: The linear time – branching time spectrum I. The semantics of concrete, sequential processes. In: Bergstra, J.A., Ponse, A., Smolka, S.A. (eds.) Handbook of Process Algebra, pp. 3–99. North-Holland, Amsterdam (2001)

34. Godefroid, P., Levin, M.Y., Molnar, D.A.: Automated whitebox fuzz testing. In: Proceedings of the Network and Distributed System Security Symposium, NDSS 2008, 10th February - 13th February 2008, San Diego, California, USA. The Internet Society (2008)
35. Godefroid, P., Klarlund, N., Sen, K.: DART: directed automated random testing. *SIGPLAN Not.* **40**(6), 213–223 (2005)
36. Grinchtein, O., Jonsson, B., Leucker, M.: Learning of event-recording automata. *Theoret. Comput. Sci.* **411**(47), 4029–4054 (2010)
37. Groce, A., Peled, D., Yannakakis, M.: Adaptive model checking. *Logic J. IGPL* **14**(5), 729–744 (2006)
38. Hagerer, A., Hungar, H., Niese, O., Steffen, B.: Model generation by moderated regular extrapolation. In: Kutsche, R.-D., Weber, H. (eds.) FASE 2002. LNCS, vol. 2306, pp. 80–95. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45923-5_6
39. Hagerer, A., Margaria, T., Niese, O., Steffen, B., Brune, G., Ide, H.-D.: Efficient regression testing of CTI-systems: testing a complex call-center solution. *Ann. Rev. Commun. Int. Eng. Consortium (IEC)* **55**, 1033–1040 (2001)
40. Hooimeijer, P., Veanes, M.: An evaluation of automata algorithms for string analysis. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 248–262. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-18275-4_18
41. Hörschele, M., Zeller, A.: Mining input grammars from dynamic taints. In: Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE 2016, pp. 720–725. ACM, New York (2016)
42. Howar, F., Isberner, M., Steffen, B., Bauer, O., Jonsson, B.: Inferring semantic interfaces of data structures. In: Margaria, T., Steffen, B. (eds.) ISO/LA 2012. LNCS, vol. 7609, pp. 554–571. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34026-0_41
43. Howar, F., Steffen, B., Jonsson, B., Cassel, S.: Inferring canonical register automata. In: Kunčak, 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
44. Howar, F., Steffen, B., Merten, M.: Automata learning with automated alphabet abstraction refinement. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 263–277. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-18275-4_19
45. Howar, F., Giannakopoulou, D., Rakamarić, Z.: Hybrid learning: interface generation through static, dynamic, and symbolic analysis. In: Proceedings of the 2013 International Symposium on Software Testing and Analysis, ISSA 2013, pp. 268–279. ACM, New York (2013)
46. Hungar, H., Niese, O., Steffen, B.: Domain-specific optimization in automata learning. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 315–327. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45069-6_31
47. Isberner, M., Howar, F., Steffen, B.: Inferring automata with state-local alphabet abstractions. In: Brat, G., Rungta, N., Venet, A. (eds.) NFM 2013. LNCS, vol. 7871, pp. 124–138. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38088-4_9
48. 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
49. Isberner, M., Howar, F., Steffen, B.: Learning register automata: from languages to program structures. *Mach. Learn.* **96**(1–2), 65–98 (2014)

50. 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
51. Khalili, A., Natale, L., Tacchella, A.: Reverse engineering of middleware for verification of robot control architectures. In: Brugali, D., Broenink, J.F., Kroeger, T., MacDonald, B.A. (eds.) SIMPAR 2014. LNCS (LNAI), vol. 8810, pp. 315–326. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11900-7_27
52. Lee, D., Yannakakis, M.: Principles and methods of testing finite state machines – a survey. *Proc. IEEE* **84**(8), 1090–1123 (1996)
53. Margaria, T., Niese, O., Raffelt, H., Steffen, B.: Efficient test-based model generation for legacy reactive systems. In: Proceedings of the 2004 Ninth IEEE International High-Level Design Validation and Test Workshop, HLDVT 2004, Washington, DC, USA, pp. 95–100. IEEE Computer Society (2004)
54. Margaria, T., Raffelt, H., Steffen, B.: Knowledge-based relevance filtering for efficient system-level test-based model generation. *Innov. Syst. Softw. Eng.* **1**(2), 147–156 (2005)
55. Meinke, K., Sindhu, M.A.: LBTest: a learning-based testing tool for reactive systems. In: Sixth IEEE International Conference on Software Testing, Verification and Validation, ICST 2013, Luxembourg, Luxembourg, 18–22 March 2013, pp. 447–454 (2013)
56. Mens, I.-E., Maler, O.: Learning regular languages over large ordered alphabets. *Log. Methods Comput. Sci.* **11**(3), 1–22 (2015)
57. Moerman, J., Sammartino, M., Silva, A., Klin, B., Szyrwelski, M.: Learning nominal automata. In: Castagna, G., Gordon, A.D. (eds.) Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, 18–20 January 2017, Paris, France, pp. 613–625. ACM (2017)
58. Mues, M., Howar, F., Luckow, K., Kahsai, T., Rakamaric, Z.: Releasing the PSYCO: using symbolic search in interface generation for Java. *ACM SIGSOFT Softw. Eng. Notes* **41**(6), 1–5 (2016)
59. Neven, F., Schwentick, T., Vianu, V.: Finite state machines for strings over infinite alphabets. *ACM Trans. Comput. Logic* **5**(3), 403–435 (2004)
60. Peled, D., Vardi, M.Y., Yannakakis, M.: Black box checking. *J. Automata Lang. Comb.* **7**(2), 225–246 (2002)
61. Raffelt, H., Merten, M., Steffen, B., Margaria, T.: Dynamic testing via automata learning. *STTT* **11**(4), 307–324 (2009)
62. de Ruiter, J., Poll, E.: Protocol state fuzzing of TLS implementations. In: 24th USENIX Security Symposium (USENIX Security 2015), Washington, D.C., pp. 193–206. USENIX Association, August 2015
63. 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
64. Shahbaz, M., Groz, R.: Analysis and testing of black-box component-based systems by inferring partial models. *Softw. Test. Verif. Reliab.* **24**(4), 253–288 (2014)
65. Smeenk, W., Moerman, J., Vaandrager, F., Jansen, D.N.: Applying automata learning to embedded control software. In: Butler, M., Conchon, S., Zaidi, F. (eds.) ICFEM 2015. LNCS, vol. 9407, pp. 67–83. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25423-4_5
66. Smetsters, R., Moerman, J., Janssen, M., Verwer, S.: Complementing model learning with mutation-based fuzzing. *CoRR*, abs/1611.02429 (2016)

67. Sun, J., Xiao, H., Liu, Y., Lin, S.W., Qin, S.: TLV: abstraction through testing, learning, and validation. In: Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015, pp. 698–709. ACM, New York (2015)
68. Tretmans, J.: A formal approach to conformance testing. Ph.D. thesis, University of Twente, December 1992
69. Tretmans, J.: Model-based testing and some steps towards test-based modelling. In: Bernardo, M., Issarny, V. (eds.) SFM 2011. LNCS, vol. 6659, pp. 297–326. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21455-4_9
70. Vaandrager, F.W.: Model learning. Commun. ACM **60**(2), 86–95 (2017)
71. Veanes, M., Hooimeijer, P., Livshits, B., Molnar, D., Bjørner, N.: Symbolic finite state transducers: algorithms and applications. In: POPL, pp. 137–150. ACM (2012)
72. Volpato, M., Tretmans, J.: Active learning of nondeterministic systems from an ioco perspective. In: Margaria, T., Steffen, B. (eds.) ISOoLA 2014. LNCS, vol. 8802, pp. 220–235. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45234-9_16
73. Windmüller, S., Neubauer, J., Steffen, B., Howar, F., Bauer, O.: Active continuous quality control. In: Proceedings of the 16th International ACM Sigsoft Symposium on Component-based Software Engineering, CBSE 2013, pp. 111–120. ACM, New York (2013)
74. Xue, Y., Wang, J., Liu, Y., Xiao, H., Sun, J., Chandramohan, M.: Detection and classification of malicious JavaScript via attack behavior modelling. In: Proceedings of the 2015 International Symposium on Software Testing and Analysis, ISSTA 2015, pp. 48–59. ACM, New York (2015)