

ICE: A Robust Framework for Learning Invariants

Pranav Garg¹, Christof Löding², P. Madhusudan¹, and Daniel Neider²

¹ University of Illinois at Urbana-Champaign

² RWTH Aachen University


Abstract. We introduce ICE, a robust learning paradigm for synthesizing invariants, that learns using examples, counter-examples, and *implications*, and show that it admits honest teachers and strongly convergent mechanisms for invariant synthesis. We observe that existing algorithms for black-box abstract interpretation can be interpreted as ICE-learning algorithms. We develop new strongly convergent ICE-learning algorithms for two domains, one for learning Boolean combinations of numerical invariants for scalar variables and one for *quantified* invariants for arrays and dynamic lists. We implement these ICE-learning algorithms in a verification tool and show they are robust, practical, and efficient.

1 Introduction

The problem of generating adequate inductive invariants to prove a program correct is at the heart of automated program verification. Synthesizing invariants is in fact the *hardest* aspect of program verification—once adequate inductive invariants are synthesized [1–5], program verification reduces to checking validity of verification conditions obtained from finite loop-free paths [6–8], which is a logic problem that has been highly automated over the years.

Invariant generation techniques can be broadly classified into two kinds: white-box techniques where the synthesizer of the invariant is acutely aware of the precise program and property that is being proved and black-box techniques where the synthesizer is largely agnostic to the structure of the program and property, but works with a partial view of the requirements of the invariant. Abstract interpretation [1], counter-example guided abstraction refinement, predicate abstraction [9, 10], the method of Craig interpolants [11, 12], IC3 [13], etc. all fall into the white-box category. In this paper, we are interested in the newly emerging black-box techniques for invariant generation.

Learning Invariants: One prominent black-box technique for invariant generation is the emerging paradigm of *learning*. Intuitively (see picture on the right), we have two components in the verification tool: a white-box *teacher* and a black-box *learner*.



The learner synthesizes suggestions for the invariants in each round. The teacher is completely aware of the program and the property being verified, and is responsible for two things: (a) to check if a purported invariant H (for hypothesis) supplied by the learner is indeed an invariant and is adequate in proving the property of the program (typically using a constraint solver), and (b) if the invariant is not adequate, to come up with concrete program configurations that need to be added or removed from the invariant

(denoted by + and – in the figure). The learner, who comes up with the invariant H is completely agnostic of the program and property being verified, and aims to build a simple formula that is consistent with the sample.

When learning an invariant, the teacher and learner talk to each other in rounds, where in each round the teacher comes up with additional constraints involving new data-points and the learner replies with some set satisfying the constraints, until the teacher finds the set to be an adequate inductive invariant. The above learning approach for invariants has been explored for quite some time in various contexts [14–16], and is gaining considerable excitement and traction in recent years [17–20].

Advantages of Learning: There are many advantages the learning approach has over white-box approaches. First, a synthesizer of invariants that works cognizant of the program and property is very hard to build, simply *due* to the fact that it has to deal with the complex logic of the program. When a program manipulates complex data-structures, pointers, objects, etc. with a complex memory model and semantics, building a set that is guaranteed to be an invariant gets extremely complex. However, the invariant for a loop in such a program may be much simpler, and hence a black-box technique that uses a “guess and check” approach guided by a finite set of configurations is much more light-weight and has better chances of finding the invariant. (See [4] where a similar argument is made for black-box generation of the abstract post in an abstract interpretation setting.) Second, learning, which typically concentrates on finding the *simplest* concept that satisfies the constraints, implicitly provides a tactic for generalization, while white-box techniques (like interpolation) need to build in tactics to generalize. Finally, the black-box approach allows us to seamlessly integrate highly scalable machine-learning techniques into the verification framework [21, 22].

ICE-learning: The problem with the learning approach described above is that it is *broken*, as we show in this paper! Approaches to learning invariants have been unduly influenced by algorithmic learning theory, automata learning, and machine learning techniques, which have traditionally offered learning from positive and negative examples. *As we show in this paper, learning using examples and counter-examples does not form a robust learning framework for synthesizing invariants.* To see why, consider the following simple program—

pre; S ; **while** (b) **do** L ; **od** S' ; *post*

with a single loop body for which we want to synthesize an invariant that proves that when the pre-condition to the program holds, the post-condition holds upon exit. Assume that the learner has just proposed a particular set H as a hypothesis invariant. In order to check if H is an adequate invariant, the teacher checks three things:

- (a) whether the strongest-post of the pre-condition across S implies H ; if not finds a concrete data-point p and passes this as a positive example to the learner.
- (b) whether the strongest-post of $(H \wedge \neg b)$ across S' implies the post-condition; if not, pass a data-point p in H that shouldn't belong to the invariant as a *negative* example.
- (c) whether H is inductive; i.e., whether the strongest post of $H \wedge b$ across loop body L implies H ; if not, finds two concrete configurations p and p' , with $p \in H$, $p' \notin H$.

In the last case above, the teacher is *stuck*. Since she does not *know* the precise invariant (there are after all many), she has no way of knowing whether p should be excluded from H or whether p' should be included. In many learning algorithms in the literature

[14–16, 20], the teacher cheats: she arbitrarily makes one choice and goes with that, hoping that it will result in an invariant. However, this makes the entire framework non-robust, causing divergence, blocking the learner from learning the simplest concepts, and introducing arbitrary bias that is very hard to control. If learning is to be seriously developed for synthesizing invariants, we need to fix this foundationally in the framework itself.

The main contribution of this paper is a new learning framework called ICE-learning, which stands for *learning using Examples, Counter-examples, and Implications*. We propose that we should build learning algorithms that do not take just examples and counter-examples, as most traditional learning algorithms do, but instead also handle *implications*. The teacher, when faced with non-inductiveness of the current conjecture H in terms of a pair (p, p') , simply communicates this implication pair to the learner, demanding that the learnt set satisfies the property that if p is included in H , then so is p' . The learner makes the choice, based on considerations of simplicity, generalization, etc., whether it would include both p and p' in its set or leave p out.

We show that ICE-learning is a *robust* learning model, in the sense that the teacher can always communicate to a learner precisely why a conjecture is not an invariant (even for programs with multiple loops, nested loops, etc.). This robustness then leads to new questions that we can formulate about learning, which we cannot ask in the setting of learning with only examples and counter-examples. In particular, we can ask whether the iterative learning process, for a particular learner, *strongly converges*— whether the learner will eventually learn the invariant, provided one exists expressible as a concept, no matter how the teacher gives examples, counter-examples, and implications to refute the learner’s conjectures.

We emphasize that earlier works in the literature have indeed seen *inductiveness* as an important aspect of synthesizing invariants, and several mechanisms for guiding the search towards an inductive property are known [13, 23–26]. Our work here is however the first that we know that develops a robust learning model that explicitly incorporates the search for inductive sets in black-box invariant generation.

Our main contributions are as follows:

- We propose the ICE-learning framework as a robust learning framework for synthesizing invariants. We study ICE-learning algorithms at two levels: ICE-learning for a particular sample as well as the *iterative* ICE-model in which the teacher and learner iteratively interact to find the invariant. The complexity of the ICE-learner for a sample, strong convergence of iterative learning, and the number of rounds of iteration required to learn are pertinent questions.
- We show that when the class of concepts forms a *lattice*, ICE learning can be often achieved, and in fact methods that already exist in the literature can be seen as ICE-learning algorithms. In particular, the abstract Houdini algorithm [3, 4] and the work reported in [27] for invariant synthesis over abstract numerical domain lattices are in fact ICE-learning algorithms. However, these algorithms are not typically *strongly convergent* and moreover, cannot learn from negative examples at all. We hence concentrate on strongly convergent ICE-learning algorithms for two different domains in this paper that do use negative examples and implications effectively.

- We develop a new ICE-learning algorithm for *Boolean combinations of numerical invariants*, which does not form a complete lattice. Given an ICE-sample, we show how to find the *simplest* expressible formula that satisfies the sample. Our algorithm *iterates* over all possible template formulas, growing in complexity, till it finds an appropriate formula, and adapts template-based synthesis techniques that use constraint solvers [28–31] to build a black-box ICE-learning algorithm. We prove that the resulting iterative ICE-algorithm is strongly convergent. Note that the user only specifies the logic for the invariants, and does not need to give templates. We build a tool over Boogie [8] for synthesizing invariants over scalar variables and show that it is extremely effective: it mostly outperforms other techniques, and furthermore gives guarantees of simplicity and strong convergence that other algorithms do not.
- As a second instantiation of the ICE-framework, we develop a new strongly convergent ICE-learning algorithm for *quantified invariants*. We develop a general technique of reducing ICE-learning of quantified properties to ICE-learning of quantifier-free properties, but where the latter is generalized to *sets of configurations* rather than single configurations. We instantiate this technique to build an ICE-learner for quantified properties of arrays and lists. This new learning algorithm (which is the most involved technical contribution of this paper) extends the classical RPNI learning algorithm for automata [32] to learning in the ICE-model and further learns *quantified data automata* [20], which can be converted to quantified logical formulas over arrays/lists. We build a prototype verifier by building this learner and the teacher as well, and show that this results in extremely efficient and robust learning of quantified invariants.

Related Work: Prominent white-box techniques for invariant synthesis include abstract interpretation [1], interpolation [11, 12] and IC3 [13]. Abstract interpretation has been used for generating invariants over mostly convex domains [2, 33], some non-convex domains [34, 35] and more recently even over non-lattice abstract domains [36]. Template based approaches to synthesizing invariants using constraint solvers have been explored in a white-box setting in [28–31], and we adapt these techniques in Section 4 for developing an ICE-learning algorithm for numerical invariants. Several white-box techniques for synthesizing quantified invariants are also known. Most of them are based on abstract interpretation or on interpolation theorems for array theories [37–45].

Turning to black-box learning-based techniques for synthesizing invariants, Daikon [46] was a prominent early technique proposed for conjunctive Boolean learning to find *likely* invariants from configurations recorded along test runs. Learning was introduced in the context of verification by Cobleigh et al. [14], which was followed by applications of Angluin’s L^* algorithm [47] to finding rely-guarantee contracts [15] and stateful interfaces for programs [16]. Houdini [3] uses essentially conjunctive Boolean learning (which can be achieved in polynomial time) to learn conjunctive invariants over templates of atomic formulas. In Section 3, we show that the Houdini algorithm along with its generalization by Thakur et al. [4] and [27] to arbitrary abstract domains like intervals, octagons, polyhedrons, linear equalities, etc. are in fact ICE-learning algorithms.

Recently, there is renewed interest in the application of learning to program verification, in particular to synthesize invariants [17–19] by using scalable machine learning techniques [21, 22] to find classifiers that can separate good states that the program can reach (positive examples) from the bad states the program is forbidden from reaching (counter-examples). Quantified *likely* invariants for linear data-structures and arrays are found from dynamic executions using learning in [20], but these aren't necessarily adequate. Boolean formula learning has also been applied recently for learning quantified invariants in [48]. In addition, learning has been applied towards inductive program synthesis [49, 50] and model extraction and testing of software [51, 52].

Counterexamples to inductiveness of an invariant have been handled in the past [24–26], but only in the context of lattice domains where the learned concepts grow monotonically and implications essentially yield positive examples. Recently, [23] tries to find inductive invariants by finding *common* interpolants for same program locations. Though [18] mentions a heuristic for handling implication samples in their algorithm for learning invariants their tool does not implement that heuristic. As far as we know, our work here is the first to explicitly incorporate the search for inductive sets in black-box invariant generation.

2 Illustrative Example

Consider the C program on the right. This program requires a scalar loop invariant ($i > p \Rightarrow j = 1$) for its verification using VCC [53]. Even in order to synthesize such a scalar invariant, white-box techniques would need to reason about the array `a[]` in the program, and in general have to deal with complex language features like objects, pointers, a complex memory model and its semantics, etc. A black-box approach can however learn such an invariant from a small set of program configurations restricted to scalars.

Consider a black-box engine that calls `foo` with the values for p — 25, 26, ... and that unrolls the loop a few times to find positive examples for (i, j, p) in the kind $(0, 0, 25), (1, 0, 25), (1, 1, 25), \dots$ for a small number of values of i , and counter-examples of the form $(100, 0, 25), (100, 2, 25), \dots (99, 0, 25), (99, 2, 25), \dots$ (values close to 100 for i and different from 1 for j). From these positive and negative examples, the learner could naturally come up with a conjecture such as $(i > 50 \Rightarrow j = 1)$ (machine learning algorithms tend to come up with such invariants).

Now notice that the teacher is *stuck* as all positive and negative examples are satisfied by the conjecture, though it is not inductive. Consequently, when using a learner from only positive and negative samples, the teacher cannot make progress. However, in ICE-learning, the teacher can give an implication pair, say of the form $((50, 0, 25), (51, 0, 25))$, and proceed with the learning. Hence we can make progress in learning, and a learner that produces the simplest conjectures satisfying the samples would eventually generalize a large enough sample to come up with a correct invariant. Our tool from Section 4 precisely learns the above mentioned invariant for this program.

```
#include <vcc.h>
int foo(int a[], int p)
_(requires (p>=25 && p<75))
_(requires a[p]==1)
_(requires \thread_local_array
    (a, 100))
{
    int i=0, j=0;
    while (i<100)
    _(invariant (i>p ==> j==1))
    {
        if (a[i]==1)
            j = 1;
        i = i+1;
    }
    _(assert j==1);
}
```

3 The ICE-Learning Framework

When defining a (machine) learning problem, one usually specifies a domain D (like points in the real plane or finite words over an alphabet), and a class of concepts C (like rectangles in the plane or regular languages), which is a class of subsets of the domain. In classical learning frameworks (see [22]), the teacher provides a set of positive examples in D that are part of the target concept, and a set of counter-examples (or negative examples) in D that are not part of the target concept. Based on these, the learner must construct a hypothesis that approximates the target concept the teacher has in mind.

ICE-learning: In our setting, the teacher does *not* have a precise target concept from C in mind, but is looking for an inductive set which meets certain additional constraints. Consequently, we extend this learning setting with a third type of information that can be provided by the teacher: implications. Formally, let D be some domain and $C \subseteq 2^D$ be a class of subsets of D , called the concepts. The teacher knows a triple (P, N, R) , where $P \subseteq D$ is an (infinite) set of positive examples, $N \subseteq D$ is an (infinite) set of counter-examples (or negative examples), and $R \subseteq D \times D$ is a relation interpreted as an (infinite) set of implications. We call (P, N, R) the *target description*, and these sets are typically *infinite* and are obtained from the program, but the teacher has the ability to query these sets effectively.

The learner is given a finite part of this information (E, C, I) with $E \subseteq P$, $C \subseteq N$, and $I \subseteq R$. We refer to (E, C, I) as an (ICE) *sample*. The task of the ICE-learner is to construct *some* hypothesis $H \in C$ such that $P \subseteq H$, $N \cap H = \emptyset$, and for each pair $(x, y) \in R$, if $x \in H$, then $y \in H$. A hypothesis with these properties is called a *correct hypothesis*. Note that a target description (P, N, R) may have several correct hypotheses (while H must include P , exclude N , and be R -closed, there can be several such sets).

Iterative ICE-learning: The above ICE-learning corresponds to a passive learning setting, in which the learner does not interact with the teacher. In general, the quality of the hypothesis will heavily depend on the amount of information contained in the sample. However, when the hypothesis is wrong, we would like the learner to gain information from the teacher using new samples. Since such a learning process proceeds in rounds, we refer to it as *iterative ICE-learning*.

The iterative ICE-learning happens in rounds, where in each round, the learner starts with some sample (E, C, I) (from previous rounds or an initialization) and constructs a hypothesis $H \in C$ from this information, and asks the teacher whether this is correct. If the hypothesis is correct (i.e., if $P \subseteq H$, $H \cap N = \emptyset$, and for every $(x, y) \in R$, if $x \in H$, then $y \in H$ as well), then the teacher answers “correct” and the learning process terminates. Otherwise, the teacher returns either some element $d \in D$ with $d \in P \setminus H$ or $d \in H \cap N$, or an implication $(x, y) \in R$ with $x \in H$ and $y \notin H$. This new information is added to the sample of the learner.

The learning proceeds in rounds and when the learning terminates, the learner has learnt *some* R -closed concept that includes P and excludes N .

Using ICE-Learning to Synthesize Invariants: Honesty and Progress

Given an ICE-learning algorithm for a concept class, we can build algorithms for synthesizing invariants by building the required (white-box) teacher. We can apply such learning for finding invariants in programs with multiple loops, nested loops, etc.

The learning will simultaneously learn all these invariant annotations. The teacher can check the hypotheses by generating verification conditions for the hypothesized invariants and by using automatic theorem provers to check their validity.

The two salient features of ICE-learning is that it facilitates *progress* and *honesty*. The teacher can always make progress by adding an example/counter-example/implication such that H (and any other previous hypothesis) does not satisfy it. Furthermore, while augmenting the sample, the teacher can answer honestly, not precluding *any* possible adequate inductive invariant of the program. Honesty and progress are impossible to achieve when learning just from positive and negative examples (when the hypothesis is not inductive, there is no way to make progress without making a dishonest choice).

Convergence: The setting of iterative ICE-learning naturally raises the question of convergence of the learner, that is, does the learner find a correct hypothesis in a finite number of rounds? We say that a learner *strongly converges*, if for every target description (P, N, R) it reaches a correct hypothesis (from the empty sample) after a finite number of rounds, no matter what information is provided by the teacher (of course, the teacher has to answer correctly according to the target description (P, N, R)).

Note that the definition above demands convergence for arbitrary triples (P, N, R) , and allows the teacher in each round to provide *any* information that contradicts the current hypothesis, and is hence a very strong property.

Observe now that for a *finite* class C of concepts, a learner strongly converges if it never constructs the same hypothesis twice. This assumption on the learner is satisfied if it only produces hypotheses H that are consistent with the sample (E, C, I) , that is, if $E \subseteq H$, $C \cap H = \emptyset$, and for each pair $(x, y) \in I$, if $x \in H$, then $y \in H$. Such a learner is called a *consistent learner*. Since the teacher always provides a witness for an incorrect hypothesis, the next hypothesis constructed by a consistent learner must be different from all the previous ones.

Lemma 1. *For a finite class C of concepts, every consistent learner strongly converges.*

For various iterative ICE-algorithm classes, where class of concepts may be infinite, we will study strong convergence.

ICE-Learning over Lattice Domains: It turns out that ICE-algorithms are especially easy to build when the class of concepts forms a lattice, as typical in an abstract interpretation setting.

Consider an abstract domain that is a lattice. Then given any sample (E, C, I) , we can compute the *best* (smallest) abstract element that satisfies the constraints (E, C, I) as follows. First, we take the least upper bound of the set of all $\alpha(e)$, for each $e \in E$. Then we see if these satisfy the implication constraints; if not, then for every pair $(p, p') \in I$ that is not satisfied, we know that p' *must* be added to the set (since p belongs to every set that includes E). Hence all these elements p' can be added by applying α to them, and we can take the lub with respect to the existing set. We continue in this fashion till we converge to an abstract element that is the smallest satisfying E and I . Now, we can check if C is excluded from it; if yes, we have computed the best set, else there is no set satisfying the constraints. The above is an ICE-algorithm for any abstract domain.

We can, using this argument, establish polynomial-time (non-iterative) ICE-learning algorithms for conjunctive formulas (in fact, this is what the classical Houdini algorithm

does [3, 22]), k -CNF formulas [22], and for abstract domains such as octagons, polyhedra, etc. as in [24, 25]

However, note that the iterative extension of the above ICE-algorithm may not halt (unless the domain has finite height). One can of course use a widening heuristically after some rounds to halt, but then clearly the iterative ICE algorithm will not be necessarily *strongly convergent*. The iterative ICE-algorithm with widening is, in fact, precisely the *abstract Houdini* algorithm proposed recently in [4], and is similar to another recent work in [27], and are not *strongly convergent*.

The iterative ICE-learning algorithms we develop in this paper *are strongly convergent*. While the above derived iterative ICE-algorithms essentially ignore counter-examples, and fail to use counter-examples and implications as a way to *come down* the lattice after a widening/over-generalization, the algorithms we propose in the next two sections are more general schemes that truly utilize examples, counter-examples, and implications to find succinct expressions.

4 An ICE-Learning Algorithm for Numerical Invariants

In this section, we describe a learning algorithm for synthesizing invariants that are arbitrary Boolean combinations of numerical atomic formulas. Since we want the learning algorithm to generalize the sample (and not capture precisely the finite set of implication-closed positive examples), we would like it to learn a formula with the *simplest* Boolean structure. In order to do so, we iterate over templates over the Boolean structure of the formulas, and learn a formula in the given template.

Note that the domain is a join-semilattice (every pair of elements has a least upper bound) since formulas are closed under disjunction. Hence we can employ the generic abstract Houdini algorithm [4] to obtain a passive ICE-learning algorithm. However, using the vanilla algorithm will learn only the precise set of positive and implication-closed set, and hence not generalize without a widening. Widening for disjunctive domains is not easy, as there are several ways to generalize disjunctive sets [54]. Furthermore, even with a widening, we will not get a *strongly convergent* iterative ICE-algorithm that we desire (see experiments in this section where abstract Houdini diverges even on conjunctive domains on some programs for this reason). The algorithm we build in this section will not only be strongly convergent but also will produce the *simplest* expressible invariant.

Let $Var = \{x_1, \dots, x_n\}$ be the set of (integer) variables in the scope of the program. For simplicity, let us restrict atomic formulas in our concept class to octagonal constraints, over program configurations, of the general form:

$$s_1 v_1 + s_2 v_2 \leq c, \quad s_1, s_2 \in \{0, +1, -1\}, \quad v_1, v_2 \in Var, \quad v_1 \neq v_2, \quad c \in \mathbb{Z}.$$

(We can handle more general atomic formulas as well; we stick to the above for simplicity and effectiveness.)

Our ICE-learning algorithm will work by iterating over more and more complex *templates* till it finds the simplest formula that satisfies the sample. A *template* fixes the Boolean structure of the desired invariants and also restricts the constants $c \in \mathbb{Z}$ appearing in the atomic formulas to lie within a finite range $[-M, +M]$, for some $M \in \mathbb{Z}^+$. Bounding the constants leads to strong convergence as we show below. For a

given template $\bigvee_i \bigwedge_j \alpha^{ij}$, the iterative ICE-learning algorithm we describe below learns an adequate invariant φ , of the form:

$$\varphi(x_1, \dots, x_n) = \bigvee_i \bigwedge_j (s_1^{ij} v_1^{ij} + s_2^{ij} v_2^{ij} \leq c^{ij}), \quad |c^{ij}| \leq M.$$

Given a sample (E, C, I) , the learner iterates through templates, and for each template, tries to find concrete values for s_k^{ij} , v_k^{ij} ($k \in \{1, 2\}$) and c^{ij} such that the formula φ is consistent with the sample; i.e., for every data-point $p \in E$, $\varphi(p)$ holds; for $p \in C$, $\varphi(p)$ does not hold; and for every implication pair $(p, p') \in I$, $\varphi(p')$ holds if $\varphi(p)$ holds. Unfortunately, finding these values in the presence of implications is hard; classifying each implication pair (p, p') as both positive or p as negative tends to create an exponential search space that is hard to search efficiently. Our ICE-learner uses a constraint solver to search this exponential space in a reasonably efficient manner. It does so by checking the satisfiability of the formula Ψ (below), over free variables s_k^{ij} , v_k^{ij} and c^{ij} , which precisely captures all the ICE-constraints. In this formula, b_p is a Boolean variable which tracks $\varphi(p)$; the Boolean variables b_p^{ij} represent the truth value of $(s_1^{ij} v_1^{ij} + s_2^{ij} v_2^{ij} \leq c^{ij})$ on point p , r_{kp}^{ij} encode the value of $s_k^{ij} \cdot v_k^{ij}$ (line 2); and d_{kp}^{ij} encode the value of v_k^{ij} (line 3).

$$\begin{aligned} \Psi(s_k^{ij}, v_k^{ij}, c^{ij}) \equiv & \left(\bigwedge_{p \in E} b_p \right) \wedge \left(\bigwedge_{p \in C} \neg b_p \right) \wedge \left(\bigwedge_{(p, p') \in I} b_p \Rightarrow b_{p'} \right) \wedge \left(\bigwedge_p (b_p \Leftrightarrow \bigvee_i \bigwedge_j b_p^{ij}) \right) \wedge \\ & \left(\bigwedge_{p, i, j} (b_p^{ij} \Leftrightarrow (\sum_{k \in \{1, 2\}} r_{kp}^{ij} \leq c^{ij})) \right) \wedge \left(\bigwedge_{i, j} (-M \leq c^{ij} \leq M) \right) \wedge \left(\bigwedge_{k \in \{1, 2\}} \left(\begin{array}{l} s_k^{ij} = 0 \Rightarrow r_{kp}^{ij} = 0 \\ s_k^{ij} = 1 \Rightarrow r_{kp}^{ij} = d_{kp}^{ij} \\ s_k^{ij} = -1 \Rightarrow r_{kp}^{ij} = -d_{kp}^{ij} \end{array} \right) \right) \wedge \\ & \left(\bigwedge_{k \in \{1, 2\}} \bigwedge_{l \in \{1, n\}} (v_k^{ij} = l \Rightarrow d_{kp}^{ij} = p(l)) \right) \wedge \left(\bigwedge_{k \in \{1, 2\}} (-1 \leq s_k^{ij} \leq 1) \right) \wedge \left(\bigwedge_{k \in \{1, 2\}} (1 \leq v_k^{ij} \leq n) \right) \wedge \left(\bigwedge_{i, j} (v_1^{ij} \neq v_2^{ij}) \right) \end{aligned}$$

Note that Ψ falls in the theory of quantifier-free linear integer arithmetic, the satisfiability of which is decidable. A satisfying assignment for Ψ gives a *consistent* formula that the learner conjectures as an invariant. If Ψ is unsatisfiable, then there is no invariant for the current template consistent with the given sample. In this case we iterate by increasing the complexity of the template. For a given template, the class of formulas conforming to the template is finite. Our enumeration of templates *dovetails* between the Boolean structure and the range of constants in the template, thereby progressively increasing the complexity of the template. Consequently, the ICE-learning algorithm always synthesizes a consistent hypothesis if there is one, and furthermore synthesizes a hypothesis of the simplest template.

A similar approach can be used for learning invariants over linear constraints, and even more general constraints if there is an effective solver for the resulting theory.

Convergence: Our iterative ICE-algorithm conjectures a consistent hypothesis in each round, and hence ensures that we do not repeat hypotheses. Furthermore, the enumeration of templates using dovetailing ensures that all templates are eventually considered, and together with the fact that there are a finite number of formulas conforming to any template ensures strong convergence.

Theorem 1. *The above ICE-learning algorithm always produces consistent conjectures and the corresponding iterative ICE-algorithm strongly converges.*

Our learning algorithm is quite different from earlier white-box constraint based approaches to invariant synthesis [28–31]. These approaches directly encode the adequacy

of the invariant (encoding the entire program’s body) into a constraint, and use Farkas’ lemma to reduce the problem to satisfiability of quantifier-free non-linear arithmetic formulas, which is harder and in general undecidable. We, on the other hand, split the task between a white-box teacher and a black-box learner, communicating only through ICE-constraints on concrete data-points. This greatly reduces the complexity of the problem, leading to a simple teacher and a much simpler learner. Our idea is more similar to [19] which use algebraic techniques to guess the coefficients.

Table 1. Results for ICE-learning numerical invariants. ICE is the total time taken by our tool. All times reported are in seconds. \times means an adequate invariant was not found.

Program	Inv-ariant	White Box		Black Box			Program	Invariant	White Box		Black Box		
		InvGen [31]	CPA [55]	absH [4]	ML [18]	ICE			InvGen [31]	CPA [55]	absH [4]	ML [18]	ICE
w1 [29]	$x \leq n$	0.1	\times	0.1	0.2	0.0	w2 [29]	$x \leq n - 1$	0.1	\times	0.2	0.1	0.0
fig6 [56]	<i>true</i>	0.1	1.3	0.1	0.1	0.0	fig1 [29]	$x \leq -1 \vee y \geq 1$	\times	4.5	\times	\times	0.1
fig8 [56]	<i>true</i>	0.0	1.4	0.0	0.0	0.0	fig3 [56]	$lock = 1 \vee x \leq y - 1$	0.1	1.4	\times	0.1	0.0
ex14 [57]	$x \geq 1$	\times	1.5	0.2	0.2	0.0	fig9 [56]	$x = 0 \wedge y \geq 0$	0.1	1.4	0.0	0.2	0.0
finf1	$x = 0$	0.1	1.5	0.1	0.4	0.0	ex23 [57]	$0 \leq y \leq z \wedge$ $z \leq c + 4572$	\times	90.5	0.2	\times	14.2
finf2	$x = 0$	0.1	1.4	0.0	0.1	0.0	ex7 [57]	$0 \leq i \wedge y \leq len$	\times	1.6	0.2	0.4	0.0
sum3	$sn = x$	0.1	1.5	0.1	0.1	0.0	sum1	$sn = i - 1 \wedge$ $(sn = 0 \vee sn \leq n)$	\times	\times	\times	\times	1.8
term2	<i>true</i>	0.0	1.6	0.0	0.0	0.0	sum4	$sn = i - 1 \wedge sn \leq 8$	0.1	2.8	\times	\times	2.6
term3	<i>true</i>	0.0	1.4	0.0	0.0	0.0	tcs [12]	$i \leq j - 1 \vee i \geq j + 1 \vee$ $x = y$	0.1	1.4	\times	0.5	1.4
trex1	$z \geq 1$	0.1	1.5	0.1	0.4	0.0	trex3	$0 \leq x1 \wedge 0 \leq x2 \wedge$ $0 \leq x3 \wedge d1 = 1 \wedge$ $d2 = 1 \wedge d3 = 1$	0.5	\times	\times	\times	2.2
trex2	<i>true</i>	0.0	1.4	0.0	0.0	0.0	matrix	$a[0][0] \leq m \vee j \leq 0;$ $a[0][0] \leq m \vee j + k \leq 0$	\times	\times	\times	\times	5.8
trex4	<i>true</i>	0.0	1.4	0.0	0.0	0.0	cgr2 [29]	$N \leq 0 \vee (x \geq 0 \wedge$ $0 \leq m \leq N - 1)$	\times	1.8	\times	\times	7.3
winf1	$x = 0$	0.0	1.4	0.0	0.0	0.0	array	$j \leq 0 \vee m \leq a[0]$	\times	\times	\times	0.2	0.3
winf2	$x = 0$	0.0	1.4	0.0	0.0	0.0	vbsd	$pathlim \leq tmp$	\times	1.6	0.5	\times	0.0
winf3	$x = 0$	\times	1.4	0.3	0.1	0.1							
vmail	$i \geq 0$	\times	1.4	0.1	0.3	0.0							
lucmp	$n = 5$	\times	77.0	0.0	0.1	0.0							
n.c11	$0 \leq len$ ≤ 4	0.1	2.2	\times	0.2	0.6							
cgr1 [29]	$x - y \leq 2$	0.1	1.5	0.1	0.2	0.0							
oct	$x + y \leq 2$	0.0	1.3	0.2	0.1	0.2							

Experimental Results: We have implemented our learning algorithm as an invariant synthesis tool¹ in Boogie [8]. In our tool we enumerate templates in an increasing order of their complexity. For a given Boolean structure of the template B_i , we fix the range of constants M in the template to be the greater value out of i and the maximum integer in the program. If an adequate invariant is not found, we increase i . If an adequate invariant is found, we use binary search on M to find an invariant that has the same Boolean structure but the smallest constants. This enumeration of templates is complete and it ensures that we learn the *simplest* invariant. In our tool, ICE-samples discovered while learning an invariant belonging to a simpler template are not wasted but used in subsequent rounds. As already mentioned, our learner uses an incremental Z3 [58] solver that adds a new constraint for every ICE-sample discovered by the Boogie based teacher.

¹ Available at <http://www.cs.uiuc.edu/~madhu/cav14/>

We evaluate our tool on SV-COMP benchmarks² and several other programs from the literature (see Table 1). We use SMACK [59] to convert C programs to Boogie and use our tool to learn loop invariants for the resulting Boogie programs. We use inlining to infer invariants for programs with multiple procedures. In Table 1 we compare our tool to invariant synthesis using abstract Houdini [4] (called absH in Table 1), [18] (called ML), Invgen [31] and interpolation based Impact algorithm [60] implemented in CPAchecker (called CPA) [55]. We implemented the octagonal domain in abstract Houdini for a comparison with our tool. As mentioned in Section 3, abstract Houdini is an ICE-learning algorithm but is not strongly convergent. Unlike our tool, abstract Houdini is not able to learn disjunctive octagonal invariants. In addition, it is unable to prove programs like `trex3` and `n.c11` where it loses precision due to widening. InvGen [31] uses a white-box constraint-based approach to invariant synthesis. Unlike our tool that enumerates all templates, InvGen requires the user to specify a template for the invariants. Being white-box, it cannot handle programs with arrays and pointers, even if the required invariants are numerical constraints over scalar variables. Being incomplete, it is also unable to prove several scalar programs like `fig1` and `cegar2`. Finally, [18] is a machine learning algorithm for inferring numerical invariants. From our experience, the inference procedure in [18] is very sensitive to the test harness used to obtain the set of safe/unsafe program configurations. For several programs, we could not learn an adequate invariant using [18] despite many attempts with different test harnesses.

The experiments show that our tool outperforms [4, 18, 31, 55] on most programs, and learns an adequate invariant for all programs in reasonable time. Though we use the more complex but more robust framework of ICE-learning that promises to learn the simplest invariants and is strongly convergent, it is generally faster than other learning algorithms like [17, 18] that learn invariants from just positive and negative examples, and lack any such promises.

5 Learning Universally Quantified Properties

In this section we describe a setting of ICE-learning for *universally quantified* concepts over linear data-structures like arrays and lists. A configuration of a program can be described by the heap structure (locations, the various field-pointers, etc.), and a finite set of pointer variables pointing into the heap. Since the heap is unbounded, typical invariants for programs manipulating heaps require universally quantified formulas. For example, a list is sorted if the data at all pairs y_1, y_2 of successive positions is sorted correctly. We consider synthesis of universal properties of the form $\psi = \forall y_1, \dots, y_k \varphi(y_1, \dots, y_k)$, where φ is a quantifier-free formula. We now describe how to modify the ICE-learning framework so that we can use a learner for the quantifier-free property described by $\varphi(y_1, \dots, y_k)$.

We consider for each concrete program configuration c the set S_c of *valuation configurations* of the form (c, val) , where val is a valuation of the variables y_1, \dots, y_k . For example, if the configurations are heaps, then the valuation maps each quantified variable y_i to a cell in the heap, akin to a scalar pointer variable. Then $c \models \psi$ if $(c, val) \models \varphi$ for all valuations val , and $c \not\models \psi$ if $(c, val) \not\models \varphi$ for some valuation val .

² <https://svn.sosy-lab.org/software/sv-benchmarks/tags/svcomp13/loops/>

This leads to the setting of *data-set based ICE-learning*. In this setting, the target description is of the form $(\hat{P}, \hat{N}, \hat{R})$ where $\hat{P}, \hat{N} \subseteq 2^D$ and $\hat{R} \subseteq 2^D \times 2^D$. A hypothesis $H \subseteq D$ is correct if $P \subseteq H$ for each $P \in \hat{P}$, $N \not\subseteq H$ for each $N \in \hat{N}$, and for each pair $(X, Y) \in \hat{R}$, if $X \subseteq H$, then also $Y \subseteq H$. The sample is a finite part of the target description, that is, it is of the form $(\hat{E}, \hat{C}, \hat{I})$, where $\hat{E}, \hat{C} \subseteq 2^D$, and $\hat{I} \subseteq 2^D \times 2^D$.

An ICE-learner for the data-set based setting corresponds to an ICE-learner for universally quantified concepts in the original data-point based setting using the following connection. Given a standard target description (P, N, R) over D , we now consider the domain D_{val} extended with valuations of the quantified variables y_1, \dots, y_k as described above. Replacing each element c of the domain by the set $S_c \subseteq D_{val}$ transforms (P, N, R) into a set-based target description. Then a hypothesis H (described by a quantifier-free formula $\varphi(y_1, \dots, y_k)$) is correct w.r.t. the set-based target description iff the hypothesis described by $\forall y_1, \dots, y_k \varphi(y_1, \dots, y_k)$ is correct w.r.t. the original target description. Unlike [40] that uses ‘‘Skolem constants’’, learning over data-sets allows us to learn from not only examples, but also from counter-examples and implications (where simple Skolem constants will not work).

Recap of Quantified Data Automata and Related Results [20]:

We will develop ICE-learning algorithms for universally quantified invariants over arrays and lists that can be expressed by an automaton model called quantified data automata (QDA) introduced by Garg et al. in [20]. We here briefly recall the main ideas concerning this model and refer the reader to [20] for more detailed definitions.

For example, consider a typical invariant in a sorting program over an array A : $\forall y_1, y_2. ((0 \leq y_1 \wedge y_2 = y_1 + 1 \wedge y_2 \leq i) \Rightarrow A[y_1] \leq A[y_2])$. This says that for all successive cells y_1, y_2 that occur somewhere in the array A before the cell pointed to by a scalar pointer variable i , the data stored at y_1 is no larger than the data stored at y_2 .

We model arrays (or other linear data structures) by *data words*, in which each position corresponds to an element or cell in the data structure. Each position in such a word is labeled with a tuple of a set of pointer variables of the program that indicates their position in the data structure and a data value from some data domain (e.g., integers) that indicates the value contained in the cell of the data structure. A QDA defines a set of data words. However, to capture the idea of expressing universally quantified properties, a QDA reads *valuation words*, which are additionally annotated with universally quantified variables. The alphabet of a QDA is a pair in which the first component corresponds to the pointer variables, and the second component contains the universally quantified variable at that position (if any).




Fig. 1. An example QDA representing an invariant of a sorting routine

The sortedness invariant above is captured by the QDA in Figure 1. The QDA accepts a valuation word if the data values at the positions of y_1 and y_2 satisfy the formula at the final state it reaches. Moreover, the automaton accepts a data word w if for *all* possible valuations of y_1 and y_2 , the automaton accepts the corresponding valuation word.

We assume that the set of formulas used to label the final states of a QDA forms a finite lattice in which the order \sqsubseteq is compatible with implication of formulas, that is, if $\varphi_1 \sqsubseteq \varphi_2$, then $\varphi_1 \Rightarrow \varphi_2$.

In [20] the subclass of elastic QDAs (EQDAs) is considered because they have a decidable emptiness problem and can be translated into decidable logics, like the array property fragment (APF) [61] for arrays, or a decidable fragment of the logic STRAND [62] for lists. The key property of these logics is their inability to express that quantified variables are only a bounded distance away from each other. This is captured at the automaton level by only allowing self loops on \underline{b} in EQDAs. The example QDA in Figure 1 is *not* an elastic QDA because there is a \underline{b} -transition from q_2 to q_5 . However, there is an EQDA for an equivalent invariant in which the sortedness property is checked for every pair of cells y_1, y_2 such that $y_1 \leq y_2$. Note that since each variable can occur only once, the blank symbol is the only one that can appear arbitrarily often in an input word. Therefore, there are only finitely many EQDAs for a fixed alphabet (set of variables). We refer the reader to [20] for more details on EQDAs.

ICE-Learning Algorithms for EQDAs. The goal of this section is to develop an iterative ICE-learner for concepts represented by EQDAs. The first relevant question is whether there is a *polynomial time* iterative ICE-learner. We show that this is *impossible* when the set of pointers and quantified variables is unbounded (see the technical report [63] for a proof sketch).

Theorem 2. *There is no polynomial time iterative ICE-learner for EQDAs, when the alphabet size is unbounded.*

The theorem is proved by adapting a result from [64], namely that there is no polynomial time learning algorithm for DFAs that only uses equivalence queries. This shows that there is no hope of obtaining an iterative ICE-learner for EQDAs (or even QDAs) in the style of the well-known L^* algorithm of Angluin, which learns DFAs in polynomial time using equivalence and membership queries.

Though we cannot hope for a polynomial time iterative ICE-learner, we develop a (non-iterative) ICE-learner that constructs an EQDA from a given sample in polynomial time. In the iterative setting this yields a learner for which each round is polynomial, while the number of rounds is not polynomial, in general. Our ICE-learning algorithm is adapted from the classical RPNI passive learning algorithm [32], which takes as input a sample (E, C) of positive example words E and counter-example words C , and constructs a DFA that DFA accepts all words in E and rejects all words in C .

RPNI can be viewed as an instance of an abstract state merging algorithm that is sketched as Algorithm 1. In this general setting, the algorithm takes a finite collection \mathcal{S} of data, called a *sample*, as input and produces a *Moore machine* (i.e., a DFA with output) that is consistent with the sample (we define this formally later). In the case of classical RPNI, $\mathcal{S} = (E, C)$ consists of two finite sets of example and counter-example words, the resulting Moore machine is interpreted as a DFA, and we require that all words in E be accepted whereas all words in C be rejected by the DFA.

Algorithm 1 proceeds in two consecutive phases. In Phase 1 (Lines 1 and 2), it calls $\text{init}(\mathcal{S})$ to construct an initial Moore machine $\mathcal{A}_{\text{init}}$ from \mathcal{S} that satisfies the sample (assuming that this is possible). Then, it picks a total order $q_0 < \dots < q_n$ on the states of

$\mathcal{A}_{\text{init}}$, which determines the order in which the states are to be merged in the subsequent phase. The actual state merging then takes place in Phase 2 (Lines 3 to 14). According to the given order, Algorithm 1 tries to merge each state q_i with a “smaller” state q_j (i.e., $j < i$) and calls `test` on the resulting Moore machine to check whether this machine

Algorithm 1: Generic State Merging algorithm.

Input: A sample \mathcal{S}

Output: A Moore machine \mathcal{A} that passes `test`(\mathcal{A})

```

1  $\mathcal{A}_{\text{init}} = (Q, \Sigma, \Gamma, q_0, \delta, f) \leftarrow \text{init}(\mathcal{S});$ 
2  $(q_0, \dots, q_n) \leftarrow \text{order}(Q);$ 
3  $\sim_0 \leftarrow \{(q, q) \mid q \in Q\};$ 
4 for  $i = 1, \dots, n$  do
5   if  $q_i \not\sim_{i-1} q_j$  for all  $j \in \{0, \dots, i-1\}$  then
6      $j \leftarrow 0;$ 
7     repeat
8       Let  $\sim$  be the smallest congruence that
9       contains  $\sim_{i-1}$  and the pair  $(q_i, q_j);$ 
10       $j \leftarrow j + 1;$ 
11     until test( $\mathcal{A}_{\text{init}}/\sim$ );
12      $\sim_i \leftarrow \sim;$ 
13   else
14      $\sim_i \leftarrow \sim_{i-1};$ 
15 end
16 return  $\mathcal{A}_{\text{init}}/\sim_n;$ 

```

still satisfies the sample; since a merge can cause nondeterminism, it might be necessary to merge further states in order to restore determinism. A merge is kept if the Moore machine passes `test`; otherwise the merge is discarded, guaranteeing that the final Moore machine still satisfies sample. Note that we represent merging of states by means of a congruence relation $\sim \subseteq Q \times Q$ over the states (i.e., \sim is an equivalence relation that is compatible with the transitions) and the actual merging operation as constructing the quotient Moore machine $\mathcal{A}_{\text{init}}/\sim$ in the usual way. Note that in the case of DFAs, each merge increases the language and thus can be seen as a generalization step in the learning algorithm.³

We are now ready to describe our new ICE-learning algorithm for EQDAs that extends the above Algorithm 1, handling both EQDAs and implication samples. In our setting, a sample is of the form $(\hat{E}, \hat{C}, \hat{I})$ where \hat{E}, \hat{C} are sets of sets of valuation words and \hat{I} contains pairs of sets of valuation words. From [20] we know that EQDAs can be viewed as Moore machines that read valuation words and output data formulas. Hence we can adapt the RPNI algorithm to learn EQDAs, as explained below.

For the initialization `init`(\mathcal{S}) we construct an EQDA whose language is the smallest (w.r.t. inclusion) EQDA-definable language that is consistent with the sample \mathcal{S} . To do this, we consider the set of all positive examples, i.e., the set $E := \bigcup \hat{E}$. This is a set of valuation words, from which we strip off the data part, obtaining a set E' of symbolic words only made up of pointers and universally quantified variables. We start with the prefix tree of E' using the prefixes of words in E' as states (as the original RPNI does). The final states are the words in E' . Each such word $w \in E'$ originates from a set of valuation words in E (all the extensions of w by data that result in a valuation word in E). If we denote this set by E_w , then we label the state corresponding to w with the least formula that is satisfied in all valuation words in E_w (recall that the formulas form a lattice). This defines the smallest QDA-definable set that contains all words in E . If this QDA is not consistent with the sample, then either there is no such QDA, or the QDA is not consistent with an implication, that is, for some $(X, Y) \in \hat{I}$ it accepts everything in

³ We refer the reader to the technical report [63] more details.

X but not everything in Y . In this case, we add X and Y to \hat{E} and restart the construction (since every QDA consistent with the sample must accept all of X and all of Y).

To make this QDA \mathcal{A} elastic, all states that are connected by a \underline{b} -transition are merged. This defines the smallest EQDA-definable set that contains all words accepted by \mathcal{A} (see [20]). Hence, if this EQDA is not consistent with the sample, then either there is no such EQDA, or an implication $(X, Y) \in \hat{I}$ is violated, and we proceed as above by adding X and Y to \hat{E} and restarting the computation. This adapted initialization results in an EQDA whose language is the *smallest EQDA-definable language that is consistent with the sample*.

Once Phase 1 is finished, our algorithm proceeds to Phase 2, in which it successively merges states of $\mathcal{A}_{\text{init}}$, to obtain an EQDA that remains consistent with the sample but has less states. When merging accepting states, the new formula at the combined state is obtained as the lub of the formulas of the original states. Note that merging states of an EQDA preserves the self-loop condition for \underline{b} -transitions. Finally, the `test` routine simply checks whether the merged EQDA is consistent with the sample.

It follows that the hypothesis constructed by this adapted version of RPNI is an EQDA that is consistent with the sample. Hence we have described a consistent learner. For a fixed set of pointer variables and universally quantified variables there are only a finite number of EQDAs. Therefore by Lemma 1 we conclude that the above learning is strongly convergent (though the number of rounds need not be polynomial).

Theorem 3. *The adaption of the RPNI algorithm for iterative set-based ICE-learning of EQDAs strongly converges.*

Experiments: We built a prototype tool implementing the set-based ICE-learning algorithm for EQDAs, consisting of both a learner and a teacher. The ICE-learner is implemented by extending the classical RPNI algorithm from the LIBALF library [65]. Given an EQDA conjectured by the learner, the teacher we build converts it to a quantified formula in the APF [61] or decidable STRAND for lists [62], and uses a constraint solver to check adequacy of invariants. Since there is no tool implementing the decision procedure for STRAND, we evaluate our prototype on array programs only.

Table 2 presents the results of our prototype on typical programs manipulating arrays⁴. We compare our results to SAFARI [44], a verification tool based on interpolation in array theories. SAFARI, in general, cannot handle list programs, and also array

Table 2. RPNI-based ICE-learning for quantified array invariants. R : # rounds of iterative-ICE; $|Q|$: # states in final EQDA. \times means timeout of 5 min.

Program	White-Box	Black-Box		
	SAFARI (s)	R	Q	ICE(s)
<i>copy</i>	0.0	4	8	0.7
<i>copy-lt-key</i>	\times	5	13	1.2
<i>init</i>	0.7	4	8	0.6
<i>init-partial</i>	\times	8	12	1.5
<i>compare</i>	0.1	9	8	1.3
<i>find</i>	0.2	9	8	1.2
<i>max</i>	0.1	3	8	0.4
<i>increment</i>	\times	5	8	0.7
<i>sorted-find</i>	\times	8	17	5.1
<i>sorted-insert</i>	\times	6	21	2.0
<i>sorted-reverse</i>	\times	18	17	9.4
<i>devres</i> [48]	0.1	3	8	0.7
<i>rm_pkey</i> [48]	0.3	3	8	0.7

⁴ Available at <http://www.cs.uiuc.edu/~madhu/cav14/>

programs like *sorted-find* that have quantified pre-conditions. On the others, SAFARI diverges for some programs, and probably needs manually provided term abstractions to achieve convergence. The results show that our ICE-learning algorithm for quantified invariants is effective, in addition to promising polynomial-per-round efficiency, promising invariants that fall in decidable theories, and promising strong convergence.

Acknowledgements. This work was partially funded by NSF CAREER award #0747041 and NSF Expeditions in Computing ExCAPE Award #1138994.

References

1. Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: POPL, ACM, pp. 238–252. ACM (1977)
2. Miné, A.: The octagon abstract domain. In: WCRE, pp. 310–319 (2001)
3. Flanagan, C., Leino, K.R.M.: Houdini, an annotation assistant for eSC/Java. In: Oliveira, J.N., Zave, P. (eds.) FME 2001. LNCS, vol. 2021, pp. 500–517. Springer, Heidelberg (2001)
4. Thakur, A., Lal, A., Lim, J., Reps, T.: PostHat and all that: Attaining most-precise inductive invariants. Technical Report TR1790, University of Wisconsin, Madison, WI (April 2013)
5. Fähndrich, M., Logozzo, F.: Static contract checking with abstract interpretation. In: Beckert, B., Marché, C. (eds.) FoVeOOS 2010. LNCS, vol. 6528, pp. 10–30. Springer, Heidelberg (2011)
6. Floyd, R.: Assigning meaning to programs. In: Schwartz, J.T. (ed.) Mathematical Aspects of Computer Science. Proceedings of Symposia in Applied Mathematics, vol. 19, pp. 19–32. AMS (1967)
7. Hoare, C.A.R.: An axiomatic basis for computer programming. Commun. ACM 12(10), 576–580 (1969)
8. Barnett, M., Chang, B.-Y.E., DeLine, R., Jacobs, B., M. Leino, K.R.: Boogie: A modular reusable verifier for object-oriented programs. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) FMCO 2005. LNCS, vol. 4111, pp. 364–387. Springer, Heidelberg (2006)
9. Ball, T., Rajamani, S.K.: The SLAM project: debugging system software via static analysis. In: POPL, pp. 1–3. ACM (2002)
10. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: POPL, pp. 58–70. ACM (2002)
11. McMillan, K.L.: Interpolation and SAT-based model checking. In: Hunt Jr., W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 1–13. Springer, Heidelberg (2003)
12. Jhala, R., McMillan, K.L.: A practical and complete approach to predicate refinement. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 459–473. Springer, Heidelberg (2006)
13. Bradley, A.R.: SAT-based model checking without unrolling. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 70–87. Springer, Heidelberg (2011)
14. Cobleigh, J.M., Giannakopoulou, D., Păsăreanu, C.S.: Learning assumptions for compositional verification. In: Garavel, H., Hatcliff, J. (eds.) TACAS 2003. LNCS, vol. 2619, pp. 331–346. Springer, Heidelberg (2003)
15. Alur, R., Madhusudan, P., Nam, W.: Symbolic compositional verification by learning assumptions. In: Etessami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 548–562. Springer, Heidelberg (2005)

16. Alur, R., Cerný, P., Madhusudan, P., Nam, W.: Synthesis of interface specifications for java classes. In: POPL, pp. 98–109. ACM (2005)
17. Sharma, R., Nori, A.V., Aiken, A.: Interpolants as classifiers. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 71–87. Springer, Heidelberg (2012)
18. Sharma, R., Gupta, S., Hariharan, B., Aiken, A., Nori, A.V.: Verification as learning geometric concepts. In: Logozzo, F., Fähndrich, M. (eds.) Static Analysis. LNCS, vol. 7935, pp. 388–411. Springer, Heidelberg (2013)
19. Sharma, R., Gupta, S., Hariharan, B., Aiken, A., Liang, P., Nori, A.V.: A data driven approach for algebraic loop invariants. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 574–592. Springer, Heidelberg (2013)
20. Garg, P., Löding, C., Madhusudan, P., Neider, D.: Learning universally quantified invariants of linear data structures. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 813–829. Springer, Heidelberg (2013)
21. Mitchell, T.M.: Machine learning. McGraw-Hill (1997)
22. Kearns, M.J., Vazirani, U.V.: An introduction to computational learning theory. MIT Press, Cambridge (1994)
23. Albarghouthi, A., McMillan, K.L.: Beautiful interpolants. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 313–329. Springer, Heidelberg (2013)
24. Reps, T., Sagiv, M., Yorsh, G.: Symbolic implementation of the best transformer. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 252–266. Springer, Heidelberg (2004)
25. Yorsh, G., Ball, T., Sagiv, M.: Testing, abstraction, theorem proving: better together! In: ISSTA, pp. 145–156 (2006)
26. van Eijk, C.A.J.: Sequential equivalence checking without state space traversal. In: DATE, pp. 618–623 (1998)
27. Garoche, P.-L., Kahsai, T., Tinelli, C.: Incremental invariant generation using logic-based automatic abstract transformers. In: Brat, G., Rungta, N., Venet, A. (eds.) NFM 2013. LNCS, vol. 7871, pp. 139–154. Springer, Heidelberg (2013)
28. Colón, M., Sankaranarayanan, S., Sipma, H.: Linear invariant generation using non-linear constraint solving. In: Hunt Jr., W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 420–432. Springer, Heidelberg (2003)
29. Gulwani, S., Srivastava, S., Venkatesan, R.: Program analysis as constraint solving. In: PLDI, pp. 281–292. ACM (2008)
30. Gupta, A., Majumdar, R., Rybalchenko, A.: From tests to proofs. In: Kowalewski, S., Philippou, A. (eds.) TACAS 2009. LNCS, vol. 5505, pp. 262–276. Springer, Heidelberg (2009)
31. Gupta, A., Rybalchenko, A.: Invgen: An efficient invariant generator. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 634–640. Springer, Heidelberg (2009)
32. Oncina, J., Garcia, P.: Inferring regular languages in polynomial update time. Pattern Recognition and Image Analysis, 49–61 (1992)
33. Cousot, P., Halbwachs, N.: Automatic discovery of linear restraints among variables of a program. In: POPL, pp. 84–96. ACM (1978)
34. Filé, G., Ranzato, F.: The powerset operator on abstract interpretations. Theor. Comput. Sci. 222(1-2), 77–111 (1999)
35. Sankaranarayanan, S., Ivančić, F., Shlyakhter, I., Gupta, A.: Static analysis in disjunctive numerical domains. In: Yi, K. (ed.) SAS 2006. LNCS, vol. 4134, pp. 3–17. Springer, Heidelberg (2006)
36. Gange, G., Navas, J.A., Schachte, P., Søndergaard, H., Stuckey, P.J.: Abstract interpretation over non-lattice abstract domains. In: Logozzo, F., Fähndrich, M. (eds.) SAS 2013. LNCS, vol. 7935, pp. 6–24. Springer, Heidelberg (2013)
37. Cousot, P., Cousot, R., Logozzo, F.: A parametric segmentation functor for fully automatic and scalable array content analysis. In: POPL, pp. 105–118. ACM (2011)

38. Gulwani, S., McCloskey, B., Tiwari, A.: Lifting abstract interpreters to quantified logical domains. In: POPL, pp. 235–246. ACM (2008)
39. Bouajjani, A., Dragoi, C., Enea, C., Sighireanu, M.: On inter-procedural analysis of programs with lists and data. In: PLDI, pp. 578–589 (2011)
40. Flanagan, C., Qadeer, S.: Predicate abstraction for software verification. In: POPL, pp. 191–202 (2002)
41. Lahiri, S.K., Bryant, R.E.: Predicate abstraction with indexed predicates. *ACM Trans. Comput. Log.* 9(1) (2007)
42. Jhala, R., McMillan, K.L.: Array abstractions from proofs. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 193–206. Springer, Heidelberg (2007)
43. McMillan, K.L.: Quantified invariant generation using an interpolating saturation prover. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 413–427. Springer, Heidelberg (2008)
44. Alberti, F., Bruttomesso, R., Ghilardi, S., Ranise, S., Sharygina, N.: SAFARI: SMT-based abstraction for arrays with interpolants. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 679–685. Springer, Heidelberg (2012)
45. Seghir, M.N., Podolski, A., Wies, T.: Abstraction refinement for quantified array assertions. In: Palsberg, J., Su, Z. (eds.) SAS 2009. LNCS, vol. 5673, pp. 3–18. Springer, Heidelberg (2009)
46. Ernst, M.D., Czeisler, A., Griswold, W.G., Notkin, D.: Quickly detecting relevant program invariants. In: ICSE, pp. 449–458. ACM (2000)
47. Angluin, D.: Learning regular sets from queries and counterexamples. *Inf. Comput.* 75(2), 87–106 (1987)
48. Kong, S., Jung, Y., David, C., Wang, B.-Y., Yi, K.: Automatically inferring quantified loop invariants by algorithmic learning from simple templates. In: Ueda, K. (ed.) APLAS 2010. LNCS, vol. 6461, pp. 328–343. Springer, Heidelberg (2010)
49. Alur, R., Bodík, R., Juniwal, G., Martin, M.M.K., Raghothaman, M., Seshia, S.A., Singh, R., Solar-Lezama, A., Torlak, E., Udupa, A.: Syntax-guided synthesis. In: FMCAD, pp. 1–17 (2013)
50. Solar-Lezama, A., Tancau, L., Bodík, R., Seshia, S.A., Saraswat, V.A.: Combinatorial sketching for finite programs. In: ASPLOS, pp. 404–415 (2006)
51. Ackermann, C., Cleaveland, R., Huang, S., Ray, A., Shelton, C., Latronico, E.: Automatic requirement extraction from test cases. In: Barringer, H., et al. (eds.) RV 2010. LNCS, vol. 6418, pp. 1–15. Springer, Heidelberg (2010)
52. Choi, W., Necula, G.C., Sen, K.: Guided gui testing of android apps with minimal restart and approximate learning. In: OOPSLA, pp. 623–640 (2013)
53. Cohen, E., Dahlweid, M., Hillebrand, M., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.: VCC: A practical system for verifying concurrent C. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 23–42. Springer, Heidelberg (2009)
54. Bagnara, R., Hill, P.M., Zaffanella, E.: Widening operators for powerset domains. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 135–148. Springer, Heidelberg (2004)
55. Beyer, D., Keremoglu, M.E.: CPAchecker: A tool for configurable software verification. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 184–190. Springer, Heidelberg (2011)
56. Gulavani, B.S., Henzinger, T.A., Kannan, Y., Nori, A.V., Rajamani, S.K.: Synergy: a new algorithm for property checking. In: SIGSOFT FSE, pp. 117–127. ACM (2006)
57. Ivancic, F., Sankaranarayanan, S.: NECLA Benchmarks, http://www.nec-labs.com/research/system/systems_SAV-website/small_static_bench-v1.1.tar.gz
58. de Moura, L., Bjørner, N.S.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)

59. Rakamaric, Z., Emmi, M.: SMACK: Static Modular Assertion Checker, <https://github.com/smackers/smack>
60. McMillan, K.L.: Lazy abstraction with interpolants. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 123–136. Springer, Heidelberg (2006)
61. Bradley, A.R., Manna, Z., Sipma, H.B.: What’s decidable about arrays? In: Emerson, E.A., Namjoshi, K.S. (eds.) VMCAI 2006. LNCS, vol. 3855, pp. 427–442. Springer, Heidelberg (2006)
62. Madhusudan, P., Parlato, G., Qiu, X.: Decidable logics combining heap structures and data. In: POPL, pp. 611–622. ACM (2011)
63. Garg, P., Löding, C., Madhusudan, P., Neider, D.: Ice: A robust framework for learning invariants. Technical report, University of Illinois (October 2013), <http://hdl.handle.net/2142/45973>
64. Angluin, D.: Negative results for equivalence queries. *Machine Learning* 5, 121–150 (1990)
65. Bollig, B., Katoen, J.-P., Kern, C., Leucker, M., Neider, D., Piegdon, D.R.: libalf: The Automata Learning Framework. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 360–364. Springer, Heidelberg (2010)