

# Polynomial-time Implicit Learnability in SMT

Ionela G. Mocanu<sup>1</sup> and Vaishak Belle<sup>2</sup> and Brendan Juba<sup>3</sup>

**Abstract.** To deploy knowledge-based systems in the real world, the challenge of knowledge acquisition must be addressed. Knowledge engineering by hand is a daunting task, so machine learning has been widely proposed as an alternative. However, machine learning has difficulty acquiring rules that feature the kind of exceptions that are prevalent in real-world knowledge. Moreover, it is conjectured to be impossible to reliably learn representations featuring a desirable level of expressiveness. Works by *Khaldon and Roth* and by *Juba* proposed solutions to such problems by learning to reason directly, bypassing the intractable step of producing an explicit representation of the learned knowledge. These works focused on Boolean, propositional logics. In this work, we consider such implicit learning to reason for arithmetic theories, including logics considered with satisfiability modulo theory (SMT) solvers. We show that for standard fragments of linear arithmetic, we can learn to reason efficiently. These results are consequences of a more general finding: we show that there is an efficient reduction from the learning to reason problem for a logic to any sound and complete solver for that logic.

## 1 Introduction

To deploy knowledge-based systems in the real world, the challenge of knowledge acquisition must be addressed. Knowledge engineering by hand is a daunting task, so machine learning has been widely proposed as an alternative. In that regard, standard techniques such as inductive logic programming (ILP) [19] and concept learning [23] have been very influential in the area. Moreover, in a logical context, Valiant [24] recognized that the challenge of learning should be integrated with deduction. In particular, he proposed a semantics to capture the quality possessed by the output of (*probably approximately correct*) PAC-learning algorithms when formulated in a logic. Although weaker than classical entailment, it allows for a powerful model-theoretic framework for answering queries.

What concerns us in this work is the extension of these frameworks to fragments of first-order logic commonly used for representing continuous domains, such as the logics considered with *satisfiability modulo theory* (SMT) solvers [1]. SMT has been widely used for applications such as model checkers, verification, unit test generators, interactive theorem provers for higher-order logic, as well as probabilistic inference [4]. Specifying the appropriate domain constraints can be very challenging, owing to the numeric nature of the underlying language. Somewhat surprisingly, there is very little work that addresses this gap. To the best of our knowledge, Kolb et al. [16] were the first to tackle this issue and propose an approach to find a

set of constraints that is consistent with a given data set. However, this may be an over-approximation, and there is no guarantee for how close these constraints are to characterizing the actual set of solutions. Moreover, this approach requires a noiseless data set. So we ask the question: *can we provide any guarantees for robust learning of SMT formulas?*

From the standpoint of learning an expressive logical knowledge base and reasoning with it, most PAC results are somewhat discouraging. For example, in agnostic learning [14] where one does not require examples (drawn from an arbitrary distribution) to be fully consistent with learned sentences, efficient algorithms for learning conjunctions would yield an efficient algorithm for PAC-learning disjunctive normal form (DNF) formulas (also over arbitrary distributions), which current evidence suggests to be intractable [9]. Works by Khaldon and Roth [15] and by Juba [13] proposed solutions to such problems by learning to reason directly, bypassing the intractable step of producing an explicit representation of the learned knowledge. Thus, there is no “discovery” of the representation, which also means that no syntactic biases are necessary, beyond the assumption that the hypothesis is drawn from the same language as the examples and the background knowledge. It is perhaps also interesting to contrast this line of work with (standard) ILP: the latter searches for a hypothesis  $H$  (a set of formulas) that is consistent with the examples by appealing to entailment. (See table 1.) It does not, however, seek to analyze the degree to which the resulting formulas capture an unknown, ground-truth process that produced the examples. Our learning task is similar to an

Learning model	PAC Learning	ILP
Language	Propositional [13], FOL + equality [3] SMT (current work)	Propositional and FOL
Examples	Partial interpretations $\{\rho^{(1)}, \dots, \rho^{(m)}\}$	Positive examples $P$ (and sometimes Negative examples $N$ )
Hypothesis	Explicit [15], Implicit ([13, 3] and current work)	Explicit
Framework	Given $KB$ , $\rho^{(i)}$ and $\alpha$ , is it true that $KB \cup$ implicit $H \models \alpha$	Given $KB$ and $P$ find $H$ such that $KB \cup H \models P$

**Table 1.** Contrast between the two learning models: the PAC learning model and the Inductive Logic Programming model

unsupervised learning model, however our end-task is deciding query entailment with respect to background knowledge and partial interpretations. This clearly defined objective function gives a principled way for deciding validity from exponentially many valid constraints. Prior relevant work proposed to produce learning of highly constrained families (in the case of [20]) or alternative ad-hoc heuristic ways of

<sup>1</sup> School of Informatics, University of Edinburgh, UK, email: i.g.mocanu@ed.ac.uk

<sup>2</sup> School of Informatics, University of Edinburgh & Alan Turing Institute, UK, email: vaishak@ed.ac.uk

<sup>3</sup> Washington University in St. Louis, USA, email: bjuba@wustl.edu

deciding what is returned.

In this work, we show how to extend the implicit learning approach to SMT formulas, yielding agnostic (implicit) learning of SMT formulas for the purposes of deciding entailment queries. Specifically, we first describe extensions of the definitions and operations underlying the implicit learning technique to SMT formulas and continuous-valued data. Second, we establish in general that implicit learning can be done for fragments of logical languages that are closed under substitutions of values for variables, and which possess sound and complete decision procedures. We do this by showing that completeness plus the closure under substitutions of the language implies the key “restriction-closure” property underpinning Juba’s framework [13]. This turns out to be convenient: we finally note that since many popular fragments of SMT are known to have such sound and complete decision procedures, we immediately obtain such constraint learning for these fragments. Previously, using Juba’s approach, this would have required a separate analysis of the corresponding logic to establish that restriction closure holds. Indeed, in some cases, the reorientation is essential: previous work suggested that logics capturing the CDCL approach do not satisfy the restriction closure property [2], and so Juba’s analysis does not provide guarantees that implicit learning succeeds with such solvers. But, our new analysis shows that we can nevertheless add implicit learning to these solvers.

In summary, our results contribute to a theoretical understanding of the PAC learnability of logical theories. The PAC learning of Boolean functions (i.e., formulas), which corresponds to the “explicit” discovery of a hypothesis, is only feasible for rather simple representations. As we mentioned, learning of CNF/DNF itself is likely intractable. Implicit learning, in contrast, simply focuses on answering queries without trying to explicitly identify the hypothesis (set of KB rules), which is the computational bottleneck. The algorithm gets as input the background knowledge and a finite set of partial assignments. The goal is to then decide entailment of the query using both the KB and the partial assignments. In other words, in the explicit approach, a formula is learned from the assignments, and entailment judgements are then computed from that formula, but in the implicit approach, the partial assignments are used to determine which queries are entailed.

## 2 Formal Framework

### 2.1 Logical background

*Satisfiability* (SAT) is the problem of deciding whether there exists an assignment of truth values (i.e., model) to variables (propositional symbols) such that a propositional logical formula  $\alpha$  is true. *Satisfiability modulo theories* (SMT) is a generalization to SAT for deciding satisfiability for fragments and extensions of first-order logics with equality, for which specialized decision procedures have been developed. Deciding satisfiability for these logics is done with respect to some decidable background theory which fixes the interpretations of functions and predicates [1]. In this work, we are especially interested in the background theories of quantifier-free arithmetic over the integers and over the reals. The following formal exposition of the logical language is adapted from [1].

**Syntax:** We assume a logical signature consisting of the set of predicates denoted as  $\mathcal{P}$ , and a set of functions symbols  $\mathcal{F}$ , including 0-ary functions, logical variables, and standard connectives. An atomic formula is one of the form:  $a$  (a propositional symbol),  $pred(t_1, \dots, t_k)$ ,  $t_1 = t_2$ ,  $\perp$  (false),  $\top$  (true). A literal  $l$  is an atomic formula or its negation  $\neg l$ . A clause is a disjunction  $l_1 \vee \dots \vee l_k$  of literals. We denote clauses as  $c$  (with superscripts and subscripts) and identify the empty

clause with the formula  $\perp$ . A ground expression is one where all the variables are replaced by the domain of discourse (e.g., integers, reals, finite set of named objects).

**Semantics:** In terms of meaning, formulas are given a truth value from the set  $\{\perp, \top\}$  by means of first order models. A model  $\rho$  is a pair consisting of a non-empty set  $\Sigma$ , the universe of the model and a mapping assigning to each constant symbol  $a$  an element  $a \in \Sigma$  (the domain), to each function symbol  $f \in \mathcal{F}$  of arity  $n > 0$  a total function  $f : \Sigma^n \rightarrow \Sigma$ , to each propositional symbol  $b$  an element  $b \in \{\perp, \top\}$  and to each predicate  $p \in \mathcal{P}$  of arity  $n > 0$  a total function  $p : \Sigma^n \rightarrow \{\perp, \top\}$ . (A partial model will be written using the regular font as  $\rho$  vs the bold font for a full model as  $\boldsymbol{\rho}$ .) For simplicity, throughout the paper we assume that the language consists of  $n$  0-arity function symbols, and so we are dealing with expressions of the form:  $x < y$ ,  $x + y > 10$ ,  $2 \cdot x > z$ , etc. In this case, a model  $\rho$  can be seen simply as an element of  $\Sigma^n$ .

Terms are interpreted as usual, as is the satisfaction relation that is defined inductively. As discussed above, we assume satisfaction and entailment with respect to a suitable background theory (e.g., theory of reals with the understanding that inequalities and other mathematical operators are interpreted as usual). See [1] for details.

### 2.2 PAC semantics and the learning model

Inductive generalization (as opposed to deduction) inherently has to cope with mistakes. Thus, the kind of knowledge produced by learning algorithms cannot hope to be valid in the traditional (Tarskian) sense, except in extreme cases, such as assuming we see every data point in a noise-free manner. The PAC semantics was introduced by Valiant [24] to capture the quality possessed by the output of PAC-learning algorithms when formulated in a logic. Essentially, it is a relaxed semantics to capture the quality possessed by knowledge learned from independently drawn examples. The relaxed notion of *validity* that formulas may satisfy is then also defined in terms of the same distribution  $D$  used to produce the examples.

**Definition 1:** [ $1 - \epsilon$ -validity [13]] Given a distribution  $D$  over  $\Sigma^n$ , we say that a Boolean function  $b$  is  $(1 - \epsilon)$ -valid if  $\Pr_{\rho \in \Sigma^n} [b(\rho) = 1] \geq 1 - \epsilon$ .

The notation and formulation introduced in [13] applies immediately to our setting<sup>4</sup>, despite the fact that over the domain of reals, we are dealing with infinitely many models and so we will need a continuous distribution. We will clarify such subtleties as and when they are introduced.

Previously, [13] would define a discrete distribution over  $\{0, 1\}^n$ , but we lift this to  $\Sigma^n$ , where  $\Sigma$  may be the set of reals  $\mathbb{R}$ , in which case any single model  $\rho$  would be assigned a density by  $D$ , and so we are saying the probability of the region satisfying  $b$  would be  $\geq 1 - \epsilon$ . For example, suppose that all points in the region  $0 \leq x_1 \leq 4$  are accorded a density of 0.25, and  $b$  denotes the formula  $x_1 \leq 2$ , then  $\Pr_{\rho \in \mathbb{R}} [b(\rho) = 1] = 0.5$ , so we say  $[x_1 \leq 2]$  is 0.5-valid.

We consider the reasoning problem of deciding whether a query formula  $\alpha$  is  $(1 - \epsilon)$ -valid with respect to a data distribution  $D$ . We suppose we have an explicit knowledge base  $\Delta$ , where we generally presume  $\Delta$  is 1-valid, i.e., satisfied except on a set of measure zero under  $D$ . If examples are drawn from  $D$ , Hoeffding’s inequality guarantees that with high probability, the proportion of times that the query

<sup>4</sup> We remark that this work is in the context of PAC-semantics as opposed to the traditional PAC learning model which focuses on learning the classifier robustly. PAC semantics refers to  $1 - \epsilon$  validity of formulas; the connection between the two is that if we use a PAC-learning algorithm to produce a rule  $f(x)$  that predicts the value  $y$ , then the formula  $f(x) = y$  is  $(1 - \epsilon)$ -valid (with probability  $1 - \delta$ ).

formula/input evaluates to *true* is a good estimate of the degree of validity of that formula.

**Theorem 1 (Hoeffding’s inequality)** *Let  $X_1, \dots, X_m$  be independent random variables taking values in  $[0, 1]$ . Then for any  $\epsilon > 0$ ,*

$$\Pr \left[ \frac{1}{m} \sum_{i=1}^m X_i \geq \mathbb{E} \left[ \frac{1}{m} \sum_{i=1}^m X_i \right] + \epsilon \right] \leq e^{-2m\epsilon^2}.$$

For complete observations, that is when the algorithm is provided with full assignments of the variables used, we could approximately decide  $1 - \epsilon$ -validity directly, distinguishing formulas that are  $1 - \epsilon$ -valid from those that are not  $1 - \epsilon - \gamma$ -valid for any desired  $\gamma > 0$  given enough data. In practice, we are often interested in queries that refer to values or properties that are not explicitly represented in the data. In other words, the algorithm only gets to see partial models which is the case we focus on. So, instead of drawing our samples directly from the distribution  $D$ , the algorithm will receive information about  $D$  in the form of partial assignments drawn from a masking process introduced below.

**Definition 2:** [Masking process [18]] A *mask* is a function  $M : \Sigma \rightarrow \Sigma \cup \{*\}$ , with the property that for any  $\rho \in \Sigma^n$ ,  $M(\rho)$  is *consistent* with  $\rho$ , i.e., whenever  $M(\rho)_i \neq *$  then  $M(\rho)_i = \rho_i$ . We refer to elements of the set  $(\Sigma \cup \{*\})^n$  as *partial assignments*. A *masking process*  $\mathbf{M}$  is a mask-valued random variable. As in [13], we denote the distribution over partial examples obtained by applying the masking process as  $\mathbf{M}(D)$ .

The masking function takes an element from  $\Sigma$  and returns either the same element or the masked value, represented as the symbol  $*$ . When applied to a full assignment, it returns the partial assignment with some of the variables being masked.

In this way, a full model, describing the state of the world becomes an observation or a partial assignment by applying the masking process to it. Once we have the partial assignments we can attempt to evaluate a formula  $\alpha$  on the partial assignment obtained from the masking process. If evaluation produces a Boolean value *true* or *false*, then we will say that this formula is *witnessed* in the partial assignment. Otherwise we will call the result of our partial evaluation a *restricted formula*:

**Definition 3:** [Restriction and witnessed formulas] Given a formula  $\alpha$  and a partial assignment  $\rho$ , the restricted formula, denoted by  $\alpha|_\rho$  is inductively defined as follows:

- If  $\alpha$  is an atomic formula and none of the terms are given value  $*$  in  $\rho$ , then  $\alpha|_\rho$  is the formula representing the value that  $\alpha$  evaluates to under the assignment given by  $\rho$ , and we say that  $\alpha$  is *witnessed*. Otherwise,  $\alpha|_\rho$  is given by substituting the assignments  $\rho_i$  for variables not given value  $*$  by  $\rho$ .
- If  $\psi = \neg\alpha$  and  $\alpha$  is not witnessed in  $\rho$ , then  $\psi|_\rho = \neg(\alpha|_\rho)$ ; otherwise,  $\psi$  is witnessed and takes the negation of the value of  $\alpha|_\rho$ .
- If  $\psi = l_1 \vee \dots \vee l_n$  is a clause, if any  $l_i|_\rho$  is witnessed true,  $\psi|_\rho$  is also witnessed true; if every  $l_i|_\rho$  is witnessed false,  $\psi|_\rho$  is also witnessed false. And finally, otherwise  $\psi|_\rho = (l_1|_\rho) \vee \dots \vee (l_n|_\rho)$
- For a restriction  $\rho$  and a set of formulas  $F$ , we let  $F|_\rho$  denote the set  $\{\alpha|_\rho : \alpha \in F\}$ .

We have modified the definition slightly from [13]; there, the partial examples consisted of sets of values from atomic formulas directly, as opposed to values for the free variables, so our treatment of atomic formulas is different. Witnessed formulas correspond to the implicit

knowledge base. We use partial models to simplify complex formulas in the  $\Delta$  and query in order to capture the inferences in the knowledge base. As shown in [13], this is sufficient whenever the reasoning algorithm is “restriction closed”:

**Definition 4:** [Restriction closure] We say that a procedure  $A$  is restriction closed if and only if  $\Delta \models \alpha$  implies that  $A$  proves  $\alpha|_\rho$  from  $\Delta|_\rho$ , for any partial assignment  $\rho$ . Essentially, from any derivation of  $\alpha$  from  $\Delta$  by  $A$ , there is a proof of  $\alpha|_\rho$  from  $\Delta|_\rho$  by  $A$ .

Precisely then, we have the following:

**Theorem 2 (Implicit learning [13])** *Let  $\Delta$  be a conjunction of constraints representing the knowledge base and an input query  $\alpha$ . We draw at random  $m = \frac{1}{2\gamma^2} \ln \frac{1}{\delta}$  partial assignments  $\{\rho^{(1)}, \rho^{(2)}, \dots, \rho^{(m)}\}$  from  $\mathbf{M}(D)$  for the distribution  $D$  and a masking process  $\mathbf{M}$ . Suppose that we have a sound, restriction-closed decision procedure  $A$ . Then with probability  $1 - \delta$ :*

- *If  $(\Delta \Rightarrow \alpha)$  is not  $(1 - \epsilon - \gamma)$ -valid with respect to the distribution  $D$ , Algorithm 1 returns *Reject*; and*
- *If there exists some KB  $I$  such that  $\Delta \wedge I \models \alpha$  and  $I$  is witnessed true with probability at least  $(1 - \epsilon + \gamma)$  on  $\mathbf{M}(D)$ , then Algorithm 1 returns *Accept*.*

*Moreover, if  $A$  runs in polynomial-time (on the number of variables, size of query and size of knowledge base), so does Algorithm 1.*

Although we have modified the definitions of  $(1 - \epsilon)$ -validity and witnessing slightly, the proof remains the same.

We define the reasoning problem as follows: an agent has some background knowledge encoded as knowledge base (KB  $\Delta$ ) and receives information about the environment as partial observations ( $\rho$ ). We then ask the agent a query  $\alpha$  and the agent reasons about it using the KB  $\Delta$  and returns an answer with some degree of validity, confidence and error, after looking at all  $m$  partial observations. The parameter  $\gamma$  in the algorithm is the accuracy of the examples used, while the parameter  $\delta$  represents the confidence of the sample received. Both parameters are bound by the interval  $[0, 1]$ . In the main algorithm, the set of derivation steps is represented by the symbol  $S$ . For example, given a constraint  $(2x + 3) = 7$  a derivation proof  $S$  would be:  $S : \{2x = 7 - 3, 2x = 4, x = 4/2, x = 2\}$ .

---

**Algorithm 1:** Implicit learning reduction

---

**Input:** Procedure  $A$ , formula  $\alpha$ , variables  $\epsilon, \delta, \gamma \in (0, 1)$ , list of partial assignments  $\{\rho^{(1)}, \rho^{(2)}, \dots, \rho^{(m)}\}$ , list of hypothesis formulas (Knowledge base)  $\Delta$

**Output:** *Accept* if there exists a derivation proof  $S$  of  $\alpha$  from  $\Delta$  and formulas  $\phi_1, \phi_2, \dots$  that are simultaneously witnessed *true* with probability at least  $(1 - \epsilon + \gamma)$ -valid on  $\mathbf{M}(D)$   
*Reject* if  $\Delta \Rightarrow \alpha$  is not  $(1 - \epsilon - \gamma)$ -valid under  $D$

```

begin
   $B \leftarrow \lfloor \epsilon \times m \rfloor$ ,  $FAILED \leftarrow 0$ .
  foreach  $i$  in  $m$  do
    if  $A(\alpha|_{\rho^{(i)}}, \Delta|_{\rho^{(i)}})$  returns UNSAT then
      Increment  $FAILED$ . if  $FAILED > B$  then
        return Reject
  return Accept

```

---

### 3 Implicit Learnability from Completeness

We now turn to considering implicit learning of arithmetic. In other words, the knowledge base  $\Delta$  contains conjunctions of linear (or non-linear) inequalities. We also suppose we are given a set of partial assignments  $\{\rho^{(1)}, \rho^{(2)}, \dots, \rho^{(m)}\}$  from which we wish to implicitly learn knowledge in service of deciding whether or not a query,  $\alpha$ , is (approximately)  $(1 - \epsilon)$ -valid. Theorem 2 establishes that it's enough to possess a polynomial-time solver that is restriction-closed. Here, we observe that as long as the solver is sound and complete for  $\mathcal{L}$ , it will be restriction-closed. Thus, for many languages of interest, since we already know that we possess sound and complete solvers, we can immediately obtain implicit learning. Recall, formally:

**Definition 5:** [Sound and complete procedure] We say that  $A$  is a sound and complete decision procedure for language  $\mathcal{L}$  if and only if, for any  $\Delta, \alpha \in \mathcal{L}$ ,  $\Delta \models \alpha$  if and only if  $A(\Delta \wedge \neg\alpha) = \text{UNSAT}$ .

Towards establishing that sound and complete solvers are restriction-closed, it will first be convenient to observe that the effect of restrictions is captured by adding conjunctive constraints. For a partial assignment  $\rho$ ,  $\rho \downarrow$  denotes a conjunction of formulas equating the variables assigned by  $\rho$  to the values  $\rho$  assigns them. For example, by an assignment  $\rho = \{x_1 = 1, x_2 = 0, x_3 = 1\}$ , we mean  $\rho \downarrow = (x_1 = 1 \wedge x_2 = 0 \wedge x_3 = 1)$ .

**Lemma 6:** Suppose  $\alpha \in \mathcal{L}$  and  $\rho$  is a partial assignment. Then  $\alpha|_\rho$  is satisfiable if and only if  $\alpha \wedge (\rho \downarrow)$  is satisfiable.

We noted that correctness suffices to ensure restriction-closure. Sometimes it is not straightforward to show that the logic is restriction-closed;<sup>5</sup> however, the guarantee that the method is sound and complete becomes a sufficient condition for the restriction closure property:

**Theorem 3** Let  $\mathcal{L}$  be a logical language such that for any  $\alpha \in \mathcal{L}$  and any partial assignment  $\rho$  for  $\mathcal{L}$ ,  $\alpha|_\rho \in \mathcal{L}$  also. Let  $A$  be a sound and complete procedure for deciding entailment for  $\mathcal{L}$ . Suppose  $\Delta, \alpha \in \mathcal{L}$ , and  $\rho$  is a partial assignment for  $\mathcal{L}$ . If  $\Delta \models \alpha$ , then  $\Delta|_\rho \models \alpha|_\rho$ .

**Proof** Suppose  $\Delta \models \alpha$ . By definition of soundness and completeness,  $\Delta \models \alpha$  if and only if  $A(\Delta \wedge \neg\alpha) = \text{UNSAT}$ . For any partial assignment  $\rho$ , we have that  $A(\Delta \wedge \neg\alpha \wedge \rho \downarrow) = \text{UNSAT}$ . By the definition of a sound and complete decision procedure,  $\Delta|_\rho \models \alpha|_\rho$  if and only if  $A(\Delta|_\rho \wedge \neg\alpha|_\rho) = \text{UNSAT}$ , if and only if  $A((\Delta \wedge \neg\alpha)|_\rho) = \text{UNSAT}$ , if and only if  $A(\Delta \wedge \neg\alpha \wedge \rho \downarrow) = \text{UNSAT}$ . ■

Thus, as a corollary of Theorem 2 together with Theorem 3, languages  $\mathcal{L}$  with sound and complete solvers have solvers with implicit learning, that moreover are polynomial-time whenever the original solver was. Theorem 3 establishes that whenever the solver is complete for a restriction closed language, Algorithm 1 will be correct.

Algorithm 1 therefore represents the reduction from a learning to reason problem to a sound and complete solver by combining Theorem 2 and Theorem 3. The rest of the results in the paper (Theorems 10, 12, 13) illustrate the breadth of applicability of this style of analysis. In short, since we often have established the completeness of our solvers, Theorem 3 immediately establishes that we can add implicit learning.

**Example 7:** Consider a smart house system which maintains the rooms temperature and ventilation at the optimal conditions. The system is capable of answering queries such as whether a particular room

is sufficiently ventilated and in which case it will activate/turn on the fans for that room. The system is characterized by the following variables:  $\{t$  (temperature in the room),  $CO$  (amount of carbon monoxide detected in the room),  $occupants$  (the number of people present in the room),  $vol$  (volume of the room, depending on which current room is observed),  $vent$  (ventilation score which can be an integer between 1 and 5)  $\}$ .

All these variables are constrained within some ranges and can be represented as a conjunction of constraints, i.e., an SMT formula. Now consider a knowledge base of the form:  $KB = (15 < t < 32) \wedge (CO < 180) \wedge (0 \leq occupants < 10) \wedge (0 < vent < 5)$ . The system now observes values for some of the variables, say, the number of people in a room or the  $CO$  density. Because the observations are not full assignments of the state of the room at that specific moment, some of the variables are not seen. Consider that the system receives some partial interpretations of the form:

$$\rho^{(1)} = \{occupants = 3, vent = 3, t = 20, CO = *\}$$

$$\rho^{(2)} = \{occupants = 4, vent = 3, t = 23, CO = *\}$$

$$\rho^{(3)} = \{occupants = 5, vent = 3, t = 24, CO = *\}$$

which does not contain any information about the carbon monoxide in that room, but we know from the knowledge base that the amount of  $CO$  is within some range ( $CO < 180$ ). Consider answering whether the following query:  $\alpha = (CO/occupants) \leq 60?$  is entailed. Using the knowledge base alone is not sufficient to decide its entailment, but after receiving observations about how many people are in the room, we may assume an implicit knowledge base  $I : (occupants \geq 3)$ . (That is, this formula is witnessed by the partial interpretations.) We are now able to decide that  $KB \cup I \models \alpha$ .

We reiterate that the  $KB$  alone is not sufficient to decide whether the query holds. Moreover, the partial assignments only give us information about some of the variables and so it is not sufficient to infer the truth of the query. We require both the  $KB$  and the partial assignments to establish that the query is entailed.

### 4 Difference Logic

Difference logic is a fragment of linear arithmetic where predicates are restricted to be of the form  $x - y \bowtie k$ , where  $x$  and  $y$  are variables from  $\mathbb{Z}$  (or  $\mathbb{R}$ ),  $k$  is a numeric constant, and  $\bowtie \in \{<, >, \leq, \geq, =\}$ . Nonetheless, the set of constraints will be assumed in the normal form  $x - y \leq k$  by applying transformations for the other types of inequalities, as discussed in [25]. The fragment is useful for various verification problems involving timed automata, as well as for representing a class of probabilistic densities [4]. One can determine the satisfiability of difference constraints by viewing the constraints as a computational problem over a graph [17]. We discuss the essential ideas as they will help us understand why the restriction closure property holds, and obtain the learnability result. First, we discuss the construction of the *inequality graph* corresponding to a set of constraints:

**Definition 8:** [17] Let  $\Delta$  be a set of difference predicates and let the inequality graph  $G(V, E)$  be the graph comprising one edge  $(x_i, x_{i+1})$  with weight  $k_i$  for every constraint of the form  $x_i - x_{i+1} \leq k_i$  in  $\Delta$ .

Given a difference logic formula  $\Delta$  with non-strict inequalities only, the inequality graph corresponding to the set of difference predicates in  $\Delta$  can be used for deciding entailment of a query  $\alpha$  by refutation. A rephrasing of the needed result is:

**Theorem 4 (Adapted from [17])** Let  $\Delta' = \Delta \wedge \neg\alpha$  be a conjunction of difference constraints and let  $G$  be the corresponding inequality graph. Then  $\Delta' = \Delta \wedge \neg\alpha$  is UNSAT iff there is a negative cycle in  $G$ .

<sup>5</sup> In particular, for more sophisticated SMT solvers, see, for example, Beame et al. [2] for issues with CDCL solvers.

The computation of the negative cycle proceeds by means of the Bellman-Ford algorithm [5, 11], which solves single-source shortest-paths problem and is sound and complete. However note that, as mentioned earlier, we also need to be able to express restricted formulas in the language. For example, given  $x - y \leq 4$  and on observing that  $y$  gets a value 2, we get that  $x \leq 6$  which is *not* in the standard form for difference logic, so *prima facie* difference logic is *not* restriction-closed. But, we can consider an extension of the language by introducing a common integer variable ZERO that can be added to encode the predicate as  $x - \text{ZERO} \leq 6$  [25]; when partial evaluation would introduce an inequality with a single variable, we encode it with the ZERO variable in this way. Indeed, in any feasible solution we can subtract the value assigned to ZERO from the values assigned to all of the variables, and it is easy to see that the constraints will remain feasible, and ZERO will obtain the value 0 as desired. Using this extension of the language, we get the following result on restriction closure:

**Proposition 9:** *The constraint graph decision procedure is restriction closed.*

The restriction closure property then allows us to state the learnability result:

**Corollary 10:** *Let  $\delta, \epsilon, \gamma, m, \Delta, \alpha, \rho^{(i)}$  and  $\mathbf{M}$  be as introduced above. By utilizing the Bellman-Ford algorithm, Algorithm 1 returns Accept or Reject such that with probability at least  $(1 - \delta)$ :*

- *If  $\Delta \wedge \neg\alpha$  is  $\epsilon + \gamma$ -valid with respect to the distribution  $D$ , it returns Reject; and*
- *If there is some implicit KB  $I$  such that  $\Delta \wedge I \models \alpha$  and every formula in  $I$  is witnessed true with probability  $1 - \epsilon + \gamma$  under the partial models  $\rho^{(i)}$ , then it returns Accept.*

*The algorithm runs in time  $O(n(|\Delta| + |\alpha|)\frac{1}{2\gamma^2} \log \frac{1}{\delta})$ , where  $n$  is the number of variables in  $\Delta$ .*

**Proof** The Bellman-Ford algorithm runs in time  $O(|V||E|)$  [5, 11], where  $V$  corresponds to the number of variables in  $\Delta$  (so,  $n$ ), and the set of edges is equivalent to the set of constraints plus the query; that is,  $|E| = |\Delta| + |\alpha|$ . Every iteration costs the time for checking feasibility<sup>6</sup> which is bounded by  $O(n \times (|\Delta| + |\alpha|))$ . The total number of iterations is  $\frac{1}{2\gamma^2} \log \frac{1}{\delta}$  corresponding to the number of samples drawn, hence the total time bound is  $O(n(|\Delta| + |\alpha|)\frac{1}{2\gamma^2} \log \frac{1}{\delta})$ . ■

We discuss the full fragment of linear arithmetic below but, of course, the benefit of the above result is that the Bellman-Ford algorithm is faster than any known algorithm for deciding the full fragment of linear arithmetic. Thus, if it suffices to use difference logic for the domain of interest, we should naturally restrict the reasoner to the constraint graph procedure. Note that the sample size is computed in the same manner for all the solvers used, so we focus on the change of size complexity of the solver in the context of the main algorithm.

## 5 Linear Arithmetic

Linear arithmetic is arguably one of the most important languages considered with SMT solvers. We assume formulas of the form  $a_1x_1 + \dots + a_nx_n \bowtie b$ , where  $x_1, \dots, x_n$  are real (or integer) variables,  $a_1, \dots, a_n, b$  are rationals, and  $\bowtie$  can be any of the inequalities,

<sup>6</sup> The arithmetic operations are assumed to be performed in unit cost operation time,  $O(1)$ . This is in contrast to, for example, a model where integers are represented as strings.

as introduced for difference logic. When  $\bowtie$  ranges over the relations  $\{=, \leq, \geq\}$ , the satisfiability problem for conjunctions of such formulas is the standard *Linear Programming Feasibility* problem. Numerous sound and complete polynomial-time algorithms exist for this problem. In particular, Cohen et al. [8] present a relatively efficient sound and complete algorithm. Assuming the program is given with rational number coefficients, the strict inequalities  $\{<, >\}$  may be represented by adding a sufficiently small  $\delta$  to non-strict inequalities, using bounds on the size of the denominators obtained from Cramer's rule (a standard technique; see e.g., [10] for an effective implementation in practice). We note that partial evaluations continue to keep the program in the expected normal form. Thus, we get:

**Proposition 11:** *The linear programming decision procedure is restriction closed.*

For learnability, we therefore get:

**Corollary 12:** *Let  $\delta, \gamma, m, \epsilon, \Delta, \alpha, \rho^{(i)}$  and  $\mathbf{M}$  be as introduced above. Suppose we solve a linear program using the procedure from [8] which runs in time  $O(n^{\omega+\alpha(1)} \log(\frac{n}{\delta}))$ , where  $n$  is the number of variables in the program and  $\omega$  is the current matrix multiplication exponent ( $\omega \approx 2.373$ ). Then Algorithm 1 using this decision procedure returns Accept or Reject such that with probability greater than  $(1 - \delta)$ :*

- *If  $\Delta \wedge \neg\alpha$  is  $\epsilon + \gamma$ -valid with respect to the distribution  $D$ , it returns Reject; and*
- *If there is some implicit KB  $I$  such that  $\Delta \wedge I \models \alpha$  and every formula in  $I$  is witnessed true with probability  $1 - \epsilon + \gamma$  under the partial models  $\rho^{(i)}$ , then it returns Accept.*

*The algorithm runs in time  $O(n^{\omega+\alpha(1)} \log(\frac{n}{\delta})\frac{1}{2\gamma^2} \log \frac{1}{\delta})$ .*

**Proof** We use the algorithm of Cohen et al. for solving a system of inequalities, which runs in time  $O(n^{\omega+\alpha(1)} \log(\frac{n}{\delta}))$ . For every iteration of the main algorithm, the system of inequalities is evaluated for some partial interpretation  $\rho$ , which is logically equivalent to evaluating whether  $KB \wedge \rho \models \alpha \wedge \rho$ . To decide entailment, we need determine whether the set of inequalities obtained from that statement all together offer a feasible solution. Every iteration costs the time for checking feasibility, bounded by  $O(n^{\omega+\alpha(1)} \log(\frac{n}{\delta}))$ . The total number of iterations is  $\frac{1}{2\gamma^2} \log \frac{1}{\delta}$ , corresponding to the number of samples drawn, hence the total time bound is  $O(n^{\omega+\alpha(1)} \log(\frac{n}{\delta})\frac{1}{2\gamma^2} \log \frac{1}{\delta})$ . ■

## 6 Non-linear arithmetic

We now consider the more general case of “nonlinear” real arithmetic (NRA), i.e., of general semi-algebraic sets, that is, systems of polynomial equations and inequalities. The canonical form of polynomial constraints is  $p \bowtie 0$ , where  $\bowtie \in \{<, >, \leq, \geq\}$  and  $p$  is a sum of terms. Tarski [21] showed, using the method of quantifier elimination, that the first-order theory of the real numbers under addition and multiplication is decidable. As the result of plugging in some values for variables in to a polynomial is indeed another polynomial, and Tarski's algorithm is complete, we could apply our method to yield a solver based on Tarski's algorithm that implicitly learns such polynomial constraints:

**Corollary 13:** *(Implicit learnability over non-linear arithmetic) For a system of polynomial constraints  $\Delta$  and a polynomial constraint  $\alpha$ , Let  $\delta, \gamma, \mathbf{M}, m$ , and  $\rho^{(i)}$  be as before. Then using Tarski's decision procedure [21] and  $m = \frac{1}{2\gamma^2} \ln \frac{1}{\delta}$  partial assignments, with probability at least  $(1 - \delta)$ :*

- If  $\Delta \wedge \neg\alpha$  is  $\epsilon + \gamma$ -valid with respect to the distribution  $D$ , it returns *Reject*; and
- If there is some implicit KB  $I$  such that  $\Delta \wedge I \models \alpha$  and every formula in  $I$  is witnessed true with probability  $1 - \epsilon + \gamma$  under the partial models  $\rho^{(i)}$ , then it returns *Accept*.

However, the worst case time complexity of solving such first-order polynomial systems is doubly exponential in the number of variables [26], and in general Tarski’s method would not be effective in practice. More recent algorithms have been proposed that appear to be much more effective in practice, e.g., Jovanovic and de Moura [12], in spite of the impossibility of a strong worst-case time complexity guarantee. The same argument holds for other fragments; for example, the fragment of polynomial equalities. More effective methods are known that are sound and complete for such fragments, for example those based on Gröbner bases [7, 6]: we can decide entailment by checking if the ideal generated by  $\Delta$  remains the same when  $\alpha$  is added by computing the Gröbner basis using Buchberger’s algorithm. Although there is again no polynomial-time guarantee for such solvers – the problem is easily seen to be NP-hard – we can still apply our reduction to obtain an algorithm that implicitly learns polynomial equality constraints. Similarly, we note that for example, Tiwari and Lincoln [22] presented a solver that is effective for the instances arising in verification and synthesis, but is only complete for the fragment of polynomial equalities that have at most finitely many solutions. Since the number of solutions cannot increase when we substitute values for variables, we can apply our method to this fragment as well.

## 7 Discussion

We have motivated and proved learnability results for a number of fragments of SMT, by considering some of the main algorithmic schemes for the appropriate fragment. In this section, we briefly reflect the impact on learnability when considering other types of algorithmic schemes. A very popular and generic approach for solving SMT fragments is the *eager encoding* paradigm: an input formula in a non-propositional fragment can be encoded as a Boolean formula. This Boolean formula is equisatisfiable, but its effectiveness rests on a technique referred to as *bit-blasting* [1], which could result in a significant increase in the size of the resulting formula. In this case, we can resort to lifting the propositional learnability result from [13], but there is a catch. A polynomial learnability result is only possible if the entailment question of Algorithm 1 can be solved in polynomial time, which we do not get for the full propositional fragment owing to the NP-completeness of propositional satisfiability. The approach taken in [13] is to “promise” that the query is provable in some low-complexity fragment; for example, it is provable by a small treelike resolution proof (where “small” refers to the number of lines of the proof). Equivalently, we give up on completeness, and only seek completeness with respect to conclusions provable in low complexity in a given fragment. In general, then, one obtains a running time guarantee that is parameterized by the size of the proof of the query. We can take a similar approach here, by using an algorithm for deciding entailment that is efficient when parameterized in such terms. In general, what is needed is a fragment for which we can decide the existence of proofs efficiently, and that is “restriction-closed,” meaning that for any partial model, if we consider the restriction of each line of the proof, we obtain a proof in the same fragment. Most fragments we might consider, including specifically treelike or bounded-width resolution, are restriction-closed. So, while possible in principle if we consider the eager encoding paradigm, the caveat

is that we have to rely on the low-complexity fragment capturing the reasoning problem in question.

Analogously, one could, in principle, appeal to the *lazy encoding* paradigm too [1]. Here, a Boolean *abstraction* of the input formula is considered, which consists of substituting all predicates over respective theories with fresh propositions. If a satisfying assignment is found, specialized theory solvers determine the validity of a proposed solution with respect to the underlying theories. If such a procedure is sound and complete, then it would follow that it is restriction closed, but once again, the caveat would be that to obtain polynomial learnability, the Boolean reasoning has to be made tractable in the way discussed above. This might make the query answering process somewhat opaque, but it nonetheless would allow us to provide a learnability algorithm for a large class of SMT fragments [1]. We briefly note that Beame et al. [2] considered fragments of logics associated with CDCL solvers, and conjectured that these were the rare examples of logics that are *not* restriction-closed. Thus, using the techniques of Juba [13], it would appear that implicit learning cannot be added to CDCL solvers. But, since the solvers are complete, our approach here shows that the reduction is still correct for these solvers. The difference is that the running time, which corresponds to the size of the associated proof (extracted from the trace of the solver), could increase significantly.

## 8 Conclusion

We present the first results on learning to reason with SMT. Our results show that for common fragments, such as difference logic and LRA, where sound and complete solvers are known, we can efficiently and robustly learn constraints when deciding entailment queries. Indeed, we showed more generally that for languages closed under substitutions of values for variables, including nonlinear arithmetic, implicit learning can always be added to sound and complete decision procedures. The main contributions are the extension of the framework to handle numeric data and Theorem 3 which establishes that we can add implicit learning to any complete solver.

We also considered alternative strategies such as eager encoding. One interesting direction for future work is to consider other kinds of partial information in our examples: currently, these take the form of partial assignments. However, sometimes we might have domains in which sensors provide partial information of the form that some signal exceeds a detection threshold, or some value can be resolved to some interval (but no more precisely). We believe that these kinds of partial information can be expressed naturally as additional LRA constraints, and thus we expect we should be able to use such partial examples in implicit learning. We would like a nice characterization of what implicit knowledge bases can be learned from such partial information and when.

## Acknowledgements

Ionela G. Mocanu was supported by the Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training in Pervasive Parallelism (grant EP/L01503X/1) at the University of Edinburgh, School of Informatics. Vaishak Belle was supported by a Royal Society University Research Fellowship. Brendan Juba was supported by NSF Award CCF-1718380. We would also like to thank our reviewers for their helpful suggestions.

## REFERENCES

- [1] Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli, *Satisfiability modulo theories*, 825–885, Frontiers in Artificial Intelligence and Applications, 1 edn., (2009).
- [2] Paul Beame, Henry A. Kautz, and Ashish Sabharwal, ‘Towards understanding and harnessing the potential of clause learning’, *J. Artif. Intell. Res.*, **22**, 319–351, (2004).
- [3] Vaishak Belle and Brendan Juba, ‘Implicitly learning to reason in first-order logic’, in *NeurIPS*, (2019).
- [4] Vaishak Belle, Andrea Passerini, and Guy Van den Broeck, ‘Component caching in hybrid domains with piecewise polynomial densities’, in *AAAI*, (2016).
- [5] Richard Bellman, ‘On a routing problem’, *Quarterly of applied mathematics*, **16**(1), 87–90, (1958).
- [6] Bruno Buchberger, *Ein Algorithmus zum Auffinden der Basis-elemente des Restklassenrings nach einem nulldimensionalen Polynomideal*, Ph.D. dissertation, (1965).
- [7] Bruno Buchberger, ‘Groebner bases: An algorithmic method in polynomial ideal theory’, in *Multidimensional Systems Theory*, ed., N.K. Bose, 184–232, D. Reidel Publ. Comp., (1985).
- [8] Michael B. Cohen, Yin Tat Lee, and Zhao Song, ‘Solving linear programs in the current matrix multiplication time’, in *ACM SIGACT*, pp. 938–942, (2019).
- [9] Amit Daniely and Shai Shalev-Shwartz, ‘Complexity theoretic limitations on learning dnf’s’, in *COLT*, pp. 815–830, (2016).
- [10] Bruno Dutertre and Leonardo de Moura, ‘A fast linear-arithmetic solver for dpll(t)’, in *Computer Aided Verification*, pp. 81–94, Berlin, Heidelberg, (2006).
- [11] Lester R. Ford Jr., ‘Network flow theory’, Technical Report P-923, Rand Corporation, (1956).
- [12] Dejan Jovanović and Leonardo De Moura, ‘Solving non-linear arithmetic’, in *IJCAR*, pp. 339–354. Springer, (2012).
- [13] Brendan Juba, ‘Implicit learning of common sense for reasoning’, in *IJCAI*, pp. 939–946, (2013).
- [14] Michael J Kearns, Robert E Schapire, and Linda M Sellie, ‘Toward efficient agnostic learning’, *Machine Learning*, **17**(2-3), 115–141, (1994).
- [15] Roni Khardon and Dan Roth, ‘Learning to reason’, *J. ACM*, **44**(5), 697–725, (1997).
- [16] Samuel Kolb, Stefano Teso, Andrea Passerini, and Luc De Raedt, ‘Learning SMT(LRA) constraints using SMT solvers’, *IJCAI’18*, pp. 2333–2340. AAAI Press, (2018).
- [17] Daniel Kroening and Ofer Strichman, *Linear Arithmetic*, Berlin, Heidelberg, (2008).
- [18] Loizos Michael, ‘Partial observability and learnability’, *Artificial Intelligence*, **174**(11), 639–669, (2010).
- [19] Stephen Muggleton and Luc de Raedt, ‘Inductive logic programming: Theory and methods’, *The Journal of Logic Programming*, **19-20**, 629–679, (1994).
- [20] Luc De Raedt and Sašo Džeroski, ‘First-order jk-clausal theories are pac-learnable’, *Artificial Intelligence*, **70**(1), 375 – 392, (1994).
- [21] Alfred Tarski, ‘A decision method for elementary algebra and geometry’, *Journal of Symbolic Logic*, **17**(3), 207–207, (1952).
- [22] Ashish Tiwari and Patrick Lincoln, ‘A nonlinear real arithmetic fragment’, in *Computer Aided Verification*, eds., Armin Biere and Roderick Bloem, p. 729–736, Cham, (2014).
- [23] Leslie G. Valiant, ‘A theory of the learnable’, *Communications of the ACM*, **27**(11), 1134–1142, (1984).
- [24] Leslie G. Valiant, ‘Robust logics’, *Artificial Intelligence*, **117**(2), 231–253, (2000).
- [25] Chao Wang, Franjo Ivančić, Malay Ganai, and Aarti Gupta, ‘Deciding separation logic formulae by sat and incremental negative cycle elimination’, in *Logic for Programming, Artificial Intelligence, and Reasoning*, pp. 322–336, (2005).
- [26] Volker Weispfenning, ‘The complexity of linear problems in fields’, *Journal of Symbolic Computation*, **5**(1), 3–27, (1988).