

Learning Implicitly with Noisy Data in Linear Arithmetic

Alexander P Rader¹, Ionela G Mocanu², Vaishak Belle² and Brendan Juba³

¹Imperial College London

²University of Edinburgh

³Washington University in St. Louis

alexander.rader20@imperial.ac.uk, {i.g.mocanu, vaishak}@ed.ac.uk, bjuba@wustl.edu

Abstract

Robust learning in expressive languages with real-world data continues to be a challenging task. Numerous conventional methods appeal to heuristics without any assurances of robustness. While probably approximately correct (PAC) Semantics offers strong guarantees, learning explicit representations is not tractable, even in propositional logic. However, recent work on so-called “implicit” learning has shown tremendous promise in terms of obtaining polynomial-time results for fragments of first-order logic. In this work, we extend implicit learning in PAC-Semantics to handle noisy data in the form of intervals and threshold uncertainty in the language of linear arithmetic. We prove that our extended framework keeps the existing polynomial-time complexity guarantees. Furthermore, we provide the first empirical investigation of this hitherto purely theoretical framework. Using benchmark problems, we show that our implicit approach to learning optimal linear programming objective constraints significantly outperforms an explicit approach in practice.

1 Introduction

Data in the real world can be incomplete, noisy and imprecise. Approaches from the knowledge representation communities take great care to represent expert knowledge; however, this knowledge can be hard to come by, challenging to formalize for non-experts, and brittle, especially in the presence of noisy measurements. In contrast, connectionist approaches, such as neural networks, have been particularly successful in learning from real-world data. However, they represent knowledge as distributed networks of nodes, rendering it difficult to incorporate knowledge from other sources. Moreover, such methods usually do not come with any guarantees of soundness, and are also quite brittle to small shifts in the domain [Recht *et al.*, 2019; Koh *et al.*, 2020].

In this work, we are concerned with learning in expressive languages, such as fragments of first-order logic, where knowledge and queries are represented as logical formulas. In that regard, Valiant [2000] 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 PAC-learning algorithms when formulated in a logic.

Unfortunately, many results on learning logical languages have been discouraging. For example, in agnostic learning, [Kearns *et al.*, 1994], where one does not require examples to be fully consistent with learned sentences, efficient algorithms for learning conjunctions of formulas would yield an efficient algorithm for PAC-learning disjunctive normal forms (DNF), which current evidence suggests to be intractable [Daniely and Shalev-Shwartz, 2016].

Interestingly, Kharden and Roth [1997] and Juba [2013] observed that, by circumventing the need to produce an explicit representation, learning to reason can be effectively reduced to classical reasoning, leading to a notion of “implicit learning”. Since reasoning in propositional logic is NP-complete, implicit learning in full propositional logic also cannot be done in polynomial time (unless P=NP), but when limited to tractable fragments of propositional logic, learning inherits the tractability guarantee. Very recently, the learnability results have been extended to first-order clauses in [Belle and Juba, 2019], and then to certain fragments of *satisfiability modulo theories* (SMT) in [Mocanu *et al.*, 2020].

In this work, we extend the implicit learning approach to handle partial information given by bounds (e.g., intervals), thus capturing an even larger class of problem settings, and develop an implementation strategy for a hitherto purely theoretical framework.

1.1 Theoretical Contributions

1. Extending the PAC-Semantics framework to be able to handle noisy data.
2. Proving polynomial running time guarantees.

To model noisy and imprecise data, such as sensor readings, we represent data as sets of intervals. Previous results [Mocanu *et al.*, 2020] only allowed for sets of assignments. We prove polynomial-time guarantees for implicitly learning any fragments of SMT with polynomial-time decision procedures, such as linear arithmetic. Since we are extending the PAC-Semantics framework, we are inheriting all of its useful features: e.g., these examples could be drawn an arbitrary distribution and are not required to be fully consistent with the learned example; nonetheless, we are able to give approx-

imately correct answers with a tight guarantee.

1.2 Empirical Contributions

1. Realising the first implementation of the PAC-Semantics framework.
2. Showing that implicit reasoning can be much faster and more noise-resistant than explicit methods for linear programs (LPs).

We test our implementation on benchmark LP problems, where the task is to find the optimum objective value subject to the constraints of the problem. We compare the running times with an alternative approach of first creating an explicit model using IncaLP [Schede *et al.*, 2019] and then finding the optimum objective value based on the model. Our results show significantly lower running times using the implicit approach, and this advantage becomes more pronounced as the models grow. We also demonstrate that, in the presence of noisy data, implicit learning returns a feasible, approximate solution to the true linear program, whereas IncaLP often fails to find any model at all.

2 Preliminaries

We will now briefly review concepts from logic, satisfiability, PAC-Semantics and the formal setup used in our work.

2.1 Logic

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 is true. SMT is a generalization to SAT for deciding satisfiability for fragments and extensions of first-order logic with equality, for which specialized decision procedures have been developed. Deciding satisfiability for these logics is done with respect to some decidable background theory [Barrett *et al.*, 2009]. In this work, we are especially interested in the background theories of quantifier-free linear real arithmetic (QFLRA). The following formal exposition is adapted from [Barrett *et al.*, 2009]:

Syntax. We assume a logical signature consisting of a set of predicate symbols, and a set of functions, logical variables, and standard connectives (\wedge, \vee, \neg). An atomic formula is one of the form: b (a propositional symbol), $\text{pred}(t_1, \dots, t_k)$, $t_1 = t_2$, \perp (false), \top (true), where t_1, \dots, t_k are the terms of the signature, in the standard logical interpretation. A literal l is an atomic formula or its negation $\neg l$. A ground expression is one where all the variables are replaced by constants and in general terms of the domain of discourse.

Semantics. Formulas are given a truth value from the set $\{\perp, \top\}$ by means of first-order models. A model ρ can be seen simply as an element of Σ^n , which is the universe of the model in n dimensions. Throughout the paper we work with the signature of linear real arithmetic, with function symbols of the form $\{0, 1, +, -, \leq, <, \geq, >, =, \neq\}$, interpreted in the usual way over the reals.

2.2 PAC-Semantics

PAC-Semantics was introduced by Valiant [2000] to capture the quality possessed by knowledge learned from indepen-

dently drawn examples from some unknown distribution D . The output produced using this approach does not express validity in the traditional (Tarskian) sense. Instead, the notion of *validity* is then defined as follows:

Definition 1. [($1 - \epsilon$)-validity [Valiant, 2000]] Given a joint distribution D over Σ^n , we say that a Boolean function f is $(1 - \epsilon)$ -valid if $\Pr_{\rho \in D}[f(\rho) = 1] \geq 1 - \epsilon$. If $\epsilon = 0$, we say that f is perfectly valid.

The reasoning problem of interest is deciding whether a query formula α is $(1 - \epsilon)$ -valid. Knowledge about the distribution D comes from the set of examples ρ , independently drawn from this distribution. Additional knowledge can come from a collection of axioms Δ , known as the knowledge base (KB).

In implicit learning, the query α is answered from examples directly, without creating an explicit model. This is done by means of entailment: we repeatedly ask whether $\Delta \wedge \rho^{(k)} \models \alpha$ for examples $\rho^{(k)} \in D$. If at least $(1 - \epsilon)$ of the examples entail α , we accept. The more examples we use, the more accurate our estimate is and the more confident we can be in it. The concepts of accuracy and confidence are captured by the hyper-parameters $\gamma, \delta \in (0, 1)$, where γ represents the accuracy of the examples used and δ captures the confidence of the example received. The number of examples needed will be determined from δ and γ ; cf. Section 4 for a precise formulation.

Reasoning from complete examples is trivial: Hoeffding's inequality [Hoeffding, 1963] guarantees that with high probability, the proportion of times that the query formula evaluates to 'true' is a good estimate of the degree of validity of that formula. To capture the more interesting problem of learning from incomplete examples, previous works used the notion of a masking process developed by Michael [2010]. A masking process randomly replaces some of the values of variables in a formula with * before passing the examples to the learning algorithm.

Definition 2. [Masking process [Michael, 2010]] A *mask* is a function $M : \Sigma^n \rightarrow \{\Sigma \cup \{\ast\}\}^n$, with the property that for any $\rho \in \Sigma^n$, $M(\rho)$ is *consistent* with ρ , i.e., whenever $M(\rho)_i \neq *$ then $M(\rho)_i = \rho_i$. We refer to elements of the set $(\Sigma \cup \{\ast\})^n$ as *partial assignments*. A *masking process* M is a mask-valued random variable. We denote the distribution over partial examples obtained by applying the masking process as $M(D)$.

Example 3. Assume a language with two real-valued variables x and y . Given a full assignment $\rho : \{x = 4, y = 5\}$, after applying a masking process on this assignment, we may obtain the partial assignment $\rho : \{x = 4, y = *\}$.

In the above example, x is not masked, and thus remains consistent with ρ , while y gets masked. A partial model will be written using the regular font as ρ vs the bold font for a full model as ρ .

Work by [Mocanu *et al.*, 2020] showed that implicitly learning to reason can be done efficiently with partial assignments for standard fragments of arithmetic theories, which integrates a decision procedure to solve SMT formulas. They proposed a reduction from the learning to reason problem for




Figure 1: High-level example of the framework

a logic to any sound and complete solver for that fragment of logic. They showed that, as long as there is a sound and complete solver for a fragment of arithmetic, then reasoning under partial evaluation (denoted as $|\rho|$) can be obtained.

Theorem 4. [[Mocanu et al., 2020]] Let \mathcal{L} be a logical language such that for any $\alpha \in \mathcal{L}$ and any partial assignment ρ 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 ρ is a partial assignment for \mathcal{L} . If $\Delta \models \alpha$, then $\Delta|_{\rho} \models \alpha|_{\rho}$.

They provide an algorithm (a less general version of DecidePAC as we discuss later) that enables implicit learning of fragments of SMT under partial assignments. We continue this line of work and present an extension of the assignments used in DecidePAC to sets of intervals.

3 Use Case Example

Let us demonstrate the functionality of the extended PAC framework with a small example: Consider a smart fitness watch, which monitors the heart rate and blood oxygen levels of the wearer. It calculates the wearer's stress level using a formula: $stress = hr - 5 \times (ox - 90)$, where hr is the heart rate in beats per minute and ox is the percentage of oxygen saturation in the blood.

As illustrated in Figure 1, this formula is encoded in the knowledge base of the system, Δ , along with two chosen bounds for hr and ox . The bound for ox is set between 90% and 100%, as anything lower is unnatural and a sign of hypoxemia [Mayo Clinic, 2020]. Similarly, hr can be bounded by the values 40 and 200 [American Heart Association, 2020].

The watch alerts the user if the stress level exceeds 50, which is encoded with the query $\alpha : (stress > 50)$. Depending on the values of the heart rate and oxygen within the boundaries set in Δ , the stress value might be below or above 50. Therefore, the explicit knowledge base alone cannot answer the query. However, the watch gets regular sensor readings in the form of intervals $\phi^{(k)}$, which it can use to answer the query. Since the watch might not be tight enough

Algorithm 1: DecidePAC

Input: Procedure A , query α , validity $\epsilon \in (0, 1)$, list of partial intervals $\{\phi^{(1)}, \dots, \phi^{(m)}\}$, knowledge base Δ

Output: Accept/Reject

begin

$$B \leftarrow \lfloor \epsilon \times m \rfloor, FAILED \leftarrow 0.$$

foreach k in $1..m$ **do**

- if** $A(\alpha, \phi^{(k)}, \Delta)$ returns *UNSAT* **then**
- Increment *FAILED*.
- if** *FAILED* $> B$ **then return** Reject;

return Accept

on the wrist, no guarantees can be made about all the sensor readings being complete. In this example, some values for oxygen are missing.

Using DecidePAC (cf. Algorithm 1), the watch answers the query by means of entailment for each of the readings received and the knowledge base: $\Delta \wedge \phi^{(m)} \models \alpha$. It receives the readings $\phi^{(1)}, \phi^{(2)}, \phi^{(3)}$, out of which the first and third yield entailment, whereas the second does not (because $\{hr = 92, ox = 99\}$ leads to $stress = 47 < 50$). The third set of intervals, $\phi^{(3)}$, triggers validity despite missing data, as the stress level is above 50 for any possible values of ox . Our validity parameter $\epsilon = 0.4$ allows the algorithm to return acceptance with $(1 - \epsilon)$ validity of the query after seeing two out of three positive readings.

4 Extending the Learning Model

In this section, we extend the framework to deal with intervals as opposed to assignments and prove that polynomial-time guarantees are preserved. We refer to a masking process in the context of intervals ϕ as a blurring process, as introduced below¹:

Definition 5. [Blurring process] Given a full assignment $\rho^{(k)} = \{\rho_1, \dots, \rho_n\}$, a *blurring* function is defined as $B : \Sigma^n \rightarrow \{\Sigma \cup \{-\infty, +\infty\}\}^{2n}$, which produces a set of intervals $\phi^{(k)}$ consistent with the assignment $\rho^{(k)}$, i.e., with the property that for each ρ_i , $B(\rho)_{2i-1} \leq \rho_i \leq B(\rho)_{2i}$, where $B(\rho)_{2i-1}$ is a random value from $(-\infty, \rho_i]$ marking the lower bound of the interval and $B(\rho)_{2i}$ is a random value from $[\rho_i, \infty)$ marking the upper bound. We refer to elements of the resulting set as *partial intervals*, where a full assignment is bound by the lower and upper bound. A *blurring process* \mathbf{B} is a blur-valued random variable (i.e. a random function). We denote the distribution over intervals obtained by applying the blurring process as $\mathbf{B}(D)$.

The blurring function essentially offers a generalization of the masking process in order to model uncertainty about complete assignments. For each interval, we draw one lower and one upper bound, hence the endpoints of the n^{th} interval occupy indices $2n - 1$ and $2n$. Using the blurring process, we

¹We use superscript to represent the index of the assignment, and subscript to represent the index of the variable in each assignment.

allow observations to be not only (i) unknown (masked) but also (ii) uncertain, and this degree of uncertainty is given by the width of the interval in which the real value lies.

Example 6. Given a full assignment $\rho : \{x = 4, y = 5\}$, after applying a blurring process on this assignment, we may obtain the intervals set $\phi : \{(1 \leq x), (x \leq 6), (-\infty \leq y), (y \leq \infty)\}$.

In the above example, the value of x has increased in uncertainty, in the sense that we have a small range of possible values. The value of y has been extended on both sides to ∞ and denotes complete uncertainty (equivalent to the effect of a masking process). In other words, the relaxation of a fixed value/assignment to the interval directly corresponds to the uncertainty about the value.

The output of a blurring process is a list of partial intervals. In order to integrate these observations in a decision procedure, i.e., entailment, they require an additional transformation into logical formulas:

Definition 7. [Grounding \downarrow] Given a partial interval $\phi^{(k)} = \{\phi_1, \dots, \phi_{2n}\}$, we define $\phi^{(k)} \downarrow$ as the ground formula corresponding to the set of intervals and represented as a set of constraints, i.e., $\phi^{(k)} \downarrow = \phi_1 \wedge \dots \wedge \phi_{2n}$.

In the evaluation procedure, we define a formula to be witnessed to evaluate to *true* given partial intervals as follows:

Definition 8. [Witnessed formulas] A formula φ is witnessed true under partial intervals ϕ if $(\phi \downarrow) \models \varphi$, i.e. φ is true under every assignment possible under ϕ .

In [Mocanu *et al.*, 2020], it was shown that, if $\Delta \models \alpha$, then $\Delta|_\rho \models \alpha|_\rho$, where ρ is a partial assignment. The key difference in our work is that instead of partial assignments, we use partial intervals. However, since partial intervals are well defined formulas of QFLRA, this property directly extends to our case and the proof is analogous to [Mocanu *et al.*, 2020].

Proposition 9. For the language of QFLRA, let A be a sound and complete decision procedure. If $\Delta \models \alpha$, then $\Delta|_\phi \models \alpha|_\phi$. Equivalently, if $A(\Delta, \alpha)$, then $A(\Delta|_\phi, \alpha|_\phi)$ iff $A(\Delta \wedge \phi \downarrow, \alpha)$.

We can now prove that implicit learning remains possible with our extended framework and that existing polynomial-time guarantees from the assignment setting are preserved:

Theorem 10. [Implicit learning] Let Δ be a conjunction of constraints representing the knowledge base and α an input query. We draw at random $m = \frac{1}{2\gamma^2} \ln \frac{1}{\delta}$ sets of intervals $\{\phi^{(1)}, \phi^{(2)}, \dots, \phi^{(m)}\}$ from $\mathbf{B}(D)$ for the distribution D and a blurring process \mathbf{B} . Suppose that we have a decision procedure A . Then with probability $1 - \delta$:

- If $(\Delta \Rightarrow \alpha)$ is not $(1 - \epsilon - \gamma)$ -valid with respect to the distribution D , the DecidePAC algorithm 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 $\mathbf{B}(D)$, then DecidePAC returns Accept.

Moreover, if A runs in polynomial-time (in the number of variables, size of the query, and size of the knowledge base), so does DecidePAC.

Proof. Consider a sound and complete decision procedure A for the language domain aforementioned and the reasoning problem of deciding $\Delta \models \alpha$. By definition of soundness and completeness, $\Delta \models \alpha$ if and only if $A(\Delta \wedge \neg\alpha) = \text{UNSAT}$. Suppose we receive observations about the world as sets of blurred intervals ϕ and we wish to decide entailment of the aforementioned problem with respect to these blurred observations, hence calculate $A(\Delta|_\phi \wedge \neg\alpha|_\phi) = \text{UNSAT}$. This holds iff $A(\Delta \wedge \phi \downarrow \wedge \neg\alpha) = \text{UNSAT}$, iff $A(\Delta \wedge I \wedge \phi \downarrow \wedge \neg\alpha) = \text{UNSAT}$ for any KB I that is witnessed true under ϕ . Now the argument becomes analogous to Juba [2013]: if $\Delta \Rightarrow \alpha$ is unsatisfied on a point drawn from D , it is not entailed by the blurred example from $\mathbf{B}(D)$ either, so FAILED increments when such points are drawn, and does not increment when a suitable I is witnessed true. By Hoeffding's inequality, the returned value satisfies the given conditions with probability $1 - \delta$ for m examples. The decision procedure will return UNSAT in polynomial time $T(n)$ depending on the size of the knowledge base and query. Every iteration costs the time for checking feasibility which is bounded by the time complexity of the decision procedure used for deciding satisfiability. The total number of iterations is $m = \frac{1}{2\gamma^2} \log \frac{1}{\delta}$, corresponding to the number of samples drawn, which gives us the total time bound of $O(T(n) \cdot \frac{1}{\gamma^2} \log \frac{1}{\delta})$. ■

With this theorem, we have extended the previous result to deal with the more complex setting of threshold uncertainty, which is a much more realistic learning model. In QFLRA, the decision procedure is polynomial-time [Barrett *et al.*, 2009] and blurred examples have the natural interpretation of being imprecise measurements. However, Theorem 10 is in principle applicable to any decidable fragment of SMT. Interpreting blurred examples for other domains (e.g., as noisy observations in the form of disjunctions of partial models) is certainly interesting, but will be left as a direction for the future.

Like the previous PAC-Semantics results, implicit learning is geared for answering queries, and has been primarily a theoretical framework. In contrast, conventional methods, although not always offering learnability guarantees, provide an explicit model, making it harder to compare with the above framework and study the trade-offs. The next section shows how this is now possible.

5 Applying PAC-Semantics to Optimisation

The goal in optimisation problems is to find an optimal objective value given constraints. We propose Algorithm 2, OptimisePAC, which solves optimisation problems using DecidePAC. This allows us to compare our framework to other more traditional methods, since optimisation problems are widely studied. In particular, we will use OptimisePAC in our empirical analysis for the case of LPs. The correctness and running time are as follows:

Theorem 11. Let Δ be a conjunction of constraints representing the knowledge base and as input preference function f . We draw at random $m = O(\frac{1}{\gamma^2} \log \frac{1}{\delta})$ partial intervals $\phi^{(1)}, \dots, \phi^{(m)}$ from $\mathbf{B}(D)$ for a distribution D and a blurring

Algorithm 2: OptimisePAC

Input: Procedure A, preference function f , validity $\epsilon \in (0, 1)$, accuracy $a \in \mathbb{Z}^+$, list of intervals $\phi = \{\phi^{(1)}, \phi^{(2)}, \dots, \phi^{(m)}\}$, goal $\in \{\text{"max"}, \text{"min"}\}$

Output: estimated optimal value w.r.t. f

```

begin
  if goal = "min" then  $f \leftarrow -f$ ;
  if DecidePAC(A,  $0 \geq f, \epsilon, \phi$ ) accepts then
    if DecidePAC(A,  $-1 \geq f, \epsilon, \phi$ ) rejects then
      |  $l \leftarrow -1, u \leftarrow 0$ 
    else
      |  $b \leftarrow -2$ 
      | while DecidePAC(A,  $b \geq f, \epsilon, \phi$ ) accepts do
        | |  $b \leftarrow b \times 2$ 
        | |  $l \leftarrow b, u \leftarrow b/2$ 
    else
      | | if DecidePAC(A,  $1 \geq f, \epsilon, \phi$ ) accepts then
        | | |  $l \leftarrow 0, u \leftarrow 1$ 
      else
        | | |  $b \leftarrow 2$ 
        | | | while DecidePAC(A,  $b \geq f, \epsilon, \phi$ ) rejects do
          | | | |  $b \leftarrow b \times 2$ 
          | | | |  $l \leftarrow b/2, u \leftarrow b$ 
    for a iterations do
      | | if DecidePAC(A,  $(l + u)/2 \geq f, \epsilon, \phi$ ) accepts then
        | | |  $u \leftarrow (l + u)/2$ 
      else
        | | |  $l \leftarrow (l + u)/2$ 
  if goal = "min" then return  $-l$  else return  $l$  ;

```

process B. Suppose that we have a decision procedure A running in time $T(n)$. Then, the OptimisePAC algorithm will return a significant bits of a value v^* that is attainable on $I \wedge \Delta$ for all KBs I that are witnessed with probability $1 - \epsilon + \gamma$, and such that for the value u^* obtained by incrementing the a^{th} bit, $\Delta \Rightarrow (f \leq u^*)$ is $(1 - \epsilon - \gamma)$ -valid (resp., $f \geq u^*$ with the a^{th} bit decreased if minimising) in time $O(T(n) \cdot m \cdot a)$.

Proof. OptimisePAC first finds the approximate bounds of the optimal value by doubling the bound each iteration and running DecidePAC. If a lower bound $b/2$ and an upper bound b are found, it runs binary search to find the optimal value to the desired precision. If the decision procedure runs in time $T(n)$, the total running time stated in the theorem follows.

To prove correctness, we will use a theorem due to Talagrand [1994] (Theorem 4.9 of [Anthony and Bartlett, 1999]):

Theorem 12. [[Anthony and Bartlett, 1999, Theorem 4.9]] There are positive constants c_1, c_2, c_3 such that the following holds. Suppose that F is a set of functions defined on the domain X and that F has a finite VC dimension d . Let $\gamma \in (0, 1)$ and $m \in \mathbb{Z}^+$. Then the probability that the empirical mean of any $f \in F$ on m examples differs from its expectation by more than γ is at most $c_1 c_2^d e^{-c_3 \gamma^2 m}$.

Thus, for $m \geq \frac{c_3}{\gamma^2} (d \ln c_2 + \ln \frac{c_1}{\delta})$, the bound is at most δ .

Recall, the VC dimension is the size of the largest set of points that can be given all labels by a class ("shattered"). Consider a fixed class of Boolean functions on the blurred samples, which is parameterized by the objective value bounds b . This function outputs the value 1 whenever $\Delta \wedge \phi \wedge (f(x) \leq b)$ returns UNSAT, and 0 otherwise. We will show that this class has VC-dimension at most 1.

We will show that for any two blurred examples ϕ_1 and ϕ_2 , it is not possible to get all the labellings $\{(1, 1), (1, 0), (0, 1), (0, 0)\}$ by varying b . Suppose there is some b^* for which ϕ_1 gives the label 1 and ϕ_2 gives 0, meaning that for ϕ_1 the bound $f \leq b^*$ does not hold and for ϕ_2 it does. Since $f \leq b^*$ holds for ϕ_2 , then for any $b > b^*$, the decision procedure will return 0 for ϕ_2 . On the other hand, the bound $f \leq b^*$ will not hold for ϕ_1 for all values $b < b^*$. Thus, in either direction, one of the labels for one of the two remains the same. So, it is not possible to get all possible labellings of ϕ_1 and ϕ_2 and so the VC-dimension is ≤ 1 .

Therefore, by Talagrand's bound, given m examples, with probability $1 - \delta$, DecidePAC answers correctly for all queries made by OptimisePAC. In particular, we note that the algorithm maintains the invariant that l is the largest value for which $l \geq f$ was rejected by DecidePAC. Since it was not accepted, we see that for any I that is witnessed with probability $\geq 1 - \epsilon + \gamma$, there must exist some x satisfying $I \wedge \Delta$ with $f(x) > l$ (resp., $f(x) < l$ if minimising). Since DecidePAC does not reject with $u, \Delta \Rightarrow f(x) \leq u$ is $(1 - \epsilon - \gamma)$ -valid, where u and l only differ by the value of the a^{th} most significant bit. Thus l is as needed. ■

6 Empirical Analysis

Implicit learning in PAC-Semantics has been described rigorously on a theoretical level. However, to the best of our knowledge, the framework has never been analysed in practice. In this section, we present the first implementation of the PAC-Semantics framework.² We will use it to empirically study the differences between an implicit and explicit approach for LP problems. LP is an optimisation technique for a linear objective function. The problem consists of an optimisation function $f(\vec{x})$, where $\vec{x} \in \mathbb{R}^n$, and feasibility region expressed as the set of constraints $A \cdot \vec{x} \leq b$. We chose LPs as the domain, as they can be represented in QFLRA, and are thus polynomial-time. It also allows us to directly evaluate our framework using existing benchmarks. We will analyse the following four points: running time, accuracy of objective value estimates, noise and outlier resistance.

6.1 Implementation Setup

The implementation is written in Python and uses the Z3 Theorem Solver [de Moura and Bjørner, 2008] for SMT queries. The explicit solver we compare it to is IncaLP(SMT), an algorithm that learns linear arithmetic SMT formulas from data [Schede et al., 2019]. In this paper, we refer to IncaLP(SMT) simply as IncaLP.

²The code can be found at <https://github.com/APRader/pac-smt-arithmetic>




Figure 2: Comparing running times and objective value estimates for *simplexn*. Here, n is the number of dimensions, *distance from true f* is the absolute difference between the true and the estimated objective value.

We use the following benchmark LP problems for our analysis: **simplexn**, **cuben**, **pollution** and **police** [Hillier and Lieberman, 1995]. All these problems consist of a set of hard constraints, which define the feasible region or the boundaries of the shape and a set of soft constraints in the form of a linear objective function, which can be either maximised or minimised. The goal is to find the optimum objective value within the hard constraints. For example, the feasible region in *simplexn* is an intersection of $\frac{1}{2}n(n-1)$ uniform triangular prisms and the objective function for *pollution* represents the amount of pollution and has to be minimised.

The implicit PAC approach and the explicit IncaLP approach differ in how they reach this goal: in the PAC model, we directly compute the optimal objective value from positive examples using OptimisePAC. In contrast, the IncaLP approach first creates an SMT model from positive and negative examples. Using standard MaxSMT techniques, we can then find the optimal objective value within the model.

We ran tests on all four problems. For each test, we had 20 independent runs on increasing sample sizes from 50 to 500 and increasing dimensions from 2 to 4, if applicable. We used 50% positive and 50% negative samples. A timeout of 30 minutes was set for each run. To ensure reproducibility, we ran each test with the seed 1613 (a numerical encoding of “PAC”). The hardware we used was an Intel Core i7-6700 3.40GHz, with 16 GB of RAM running Ubuntu 20.04.

As for parameter settings, we chose the SelectDT heuristic for the IncaLP approach, as it is the most appropriate according to their paper. We also ran our experiments using the same

initial configurations as their released code. We set the accuracy for OptimisePAC to 60, which is the number of times our intervals are divided by 2. The reason being that we can match the accuracy of double precision floating-point values, which is about 2^{-53} .

6.2 Results

On the theoretical front, one of the advantages of implicit learning is argued to be efficiency, since you can skip the step of creating an explicit model. As shown in Figure 2, this effect is significant in practice for linear programs. PAC is able to get similarly good objective value estimates at significantly lower running times for *simplexn*. The larger the sample size and the higher the dimensionality, the bigger the gap between running times. (For all of our graphs, the sample standard deviation is shown.)

With the extension we introduced in this paper, PAC can now handle noisy data using intervals. If we add Gaussian noise with a standard deviation σ to each data point, we can represent it using intervals of width $4 \log d \cdot \sigma$, where d is the dimensionality. This interval covers about 95% of the density, which is why we set a validity of 95% for DecidePAC. The noise and intervals were capped at the domain boundaries. We adjusted the noise to the dimensionality, meaning that for a noise value of n , the std of the Gaussian $\sigma = \frac{n}{\sqrt{d}}$. We also tested robustness against outliers. For our purposes, an outlier is a point within the domain with a wrong label.

Figure 3 shows how PAC compares to IncaLP. In cases where IncaLP does find a model, it gives an estimate that is




Figure 3: Objective values for *pollution*. True minimum: 32.15. IncaLP failed to find a model 30, 71 and 82 out of 120 times respectively.

closer to the true objective value in noisy cases. However, it failed to find a model 25%, 59% and 68% of the time for a noise of 0.1, 0.25 and outliers of 0.01, respectively. Moreover, the PAC estimate is pessimistic on purpose. This results in values farther away from the optimum but ensures that they are reachable. PAC always finds feasible objective values, while IncaLP undershoots the minimum for *pollution* almost always.³

It is also worth noting these results do not mean that IncaLP, on the whole, is inferior to our PAC implementation. IncaLP creates a model, while PAC answers queries implicitly. For some problems, such as objective value optimisation, it is possible to skip the model making process. As we have shown, doing that comes with advantages in speed and noise resistance. However, in some contexts, having an explicit model is desirable, in which case the implicit learning paradigm might be harder to work with. Put simply, OptimisePAC is not a replacement for IncaLP.

7 Related Work

We identify two threads of related research. On the theoretical side, the problem of dealing with missing information in learning settings was recognized early on in the literature. For example, to analyze how incomplete data affects predictions, [Schuurmans and Greiner, 1994] look at various ways of altering the input data, according to a product distribution or via masking. In [Kearns and Schapire, 2001], a framework is proposed that uses incomplete observations to predict a hypothesis function and never output a “*don’t know*” answer. These and related approaches, however, focus on a discrete model and do not cover continuous-valued domains. There is, of course, also a large body of work on techniques like imputation, but these are far afield from our goals.

Conceptually, it is perhaps also interesting to contrast this line of work with *inductive logic programming* (ILP) [Muggleton and de Raedt, 1994]: the latter searches for a hypothesis H that is consistent with the examples by appealing to entailment. Like IncaLP, it does not, however, seek to analyze the degree to which the resulting formulas capture an unknown, ground-truth process that produced the examples.

³The empirical behavior is similar across the four problems. All graphs are included in the extended report [Rader *et al.*, 2020].

For more discussions on how implicit learning via the PAC-Semantics differs from such hypothesis generation methodologies, including probabilistic structure learning schemes, see [Juba, 2013; Belle and Juba, 2019].

On the empirical side, many methods for inducing linear constraints from data use greedy or approximate learning schemes, e.g., [Kantchelian *et al.*, 2014]. To overcome the problems of overfitting, local optima or noise-robustness, approaches like syntax-guided learning [Alur *et al.*, 2013] have been proposed, in which the learning problem is reduced to an optimisation problem. See [Schede *et al.*, 2019] for a comprehensive discussion. Similar to IncaLP, work by [Pawlak and Krawiec, 2017] proposes an approach of approximating the set of constraints from feasible and infeasible examples, which focuses on mixed integer linear programs, but without any theoretical guarantees.

8 Conclusion

In this work, we proposed a general framework for learning implicitly in linear arithmetic from noisy examples. By then considering a novel optimisation variant, we were able to empirically compare and outperform an explicit approach in terms of running time and resistance to noise and outliers for LP problems. A natural direction for the future is to consider whether this implementation strategy extends to other classes of formulas in first-order logic and/or SMT. In particular, it is worth conducting empirical investigations into polynomial-time fragments of SMT other than LPs.

Acknowledgements

We would like to thank our reviewers for their helpful suggestions. Alexander Rader conducted this work while at the University of Edinburgh. Ionela Georgiana 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/Amazon award IIS-1939677.

References

- [Alur *et al.*, 2013] Rajeev Alur, Rastislav Bodík, Garvit Jiniwal, Milo Martin, Mukund Raghothaman, Sanjit Se-

- shia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. Syntax-guided synthesis. In *2013 Formal Methods in Computer-Aided Design, FMCAD 2013*, pages 1–17, 10 2013.
- [American Heart Association, 2020] American Heart Association. Know your target heart rates for exercise, losing weight and health. <https://www.heart.org/en/healthy-living/fitness/fitness-basics/target-heart-rates>, 2020. Accessed: 2021-05-25.
- [Anthony and Bartlett, 1999] Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, USA, 1st edition, 1999.
- [Barrett *et al.*, 2009] Clark Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. *Satisfiability modulo theories*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 825–885. IOS Press, 1 edition, 2009.
- [Belle and Juba, 2019] Vaishak Belle and Brendan Juba. Implicitly learning to reason in first-order logic. In *NeurIPS*, 2019.
- [Daniely and Shalev-Shwartz, 2016] Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning dnf’s. In *COLT*, pages 815–830, 2016.
- [de Moura and Bjørner, 2008] Leonardo de Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In C. R. Ramakrishnan and Jakob Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [Hillier and Lieberman, 1995] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Mathematical Programming*. McGraw-Hill, 1995.
- [Hoeffding, 1963] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [Juba, 2013] Brendan Juba. Implicit learning of common sense for reasoning. In *IJCAI*, pages 939–946, 2013.
- [Kantchelian *et al.*, 2014] Alex Kantchelian, Michael C Tschantz, Ling Huang, Peter L Bartlett, Anthony D Joseph, and J Doug Tygar. Large-margin convex polytope machine. In *Advances in Neural Information Processing Systems*, pages 3248–3256, 2014.
- [Kearns and Schapire, 2001] Michael Kearns and Robert Schapire. Efficient distribution-free learning of probabilistic concepts. *Journal of Computer and System Sciences*, 48:464–497, 03 2001.
- [Kearns *et al.*, 1994] Michael J Kearns, Robert E Schapire, and Linda M Sellie. Toward efficient agnostic learning. *Machine Learning*, 17(2-3):115–141, 1994.
- [Kharden and Roth, 1997] Roni Kharden and Dan Roth. Learning to reason. *J. ACM*, 44(5):697–725, 1997.
- [Koh *et al.*, 2020] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akashay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. *arXiv preprint arXiv:2012.07421*, 2020.
- [Mayo Clinic, 2020] Mayo Clinic. Hypoxemia (low blood oxygen). <https://www.mayoclinic.org/symptoms/hypoxemia/basics/definition/sym-20050930>, 2020. Accessed: 2021-05-25.
- [Michael, 2010] Loizos Michael. Partial observability and learnability. *Artificial Intelligence*, 174(11):639–669, 2010.
- [Mocanu *et al.*, 2020] Ionela Mocanu, Vaishak Belle, and Brendan Juba. Polynomial-time implicit learnability in smt. In *Proceedings to the 24th European Conference on Artificial Intelligence (ECAI 2020)*, 1 2020.
- [Muggleton and de Raedt, 1994] Stephen Muggleton and Luc de Raedt. Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 19-20:629–679, 1994.
- [Pawlak and Krawiec, 2017] Tomasz Pawlak and Krzysztof Krawiec. Automatic synthesis of constraints from examples using mixed integer linear programming. *European Journal of Operational Research*, 261, 02 2017.
- [Rader *et al.*, 2020] Alexander P. Rader, Ionela G. Mocanu, Vaishak Belle, and Brendan Juba. Learning implicitly with noisy data in linear arithmetic. *arXiv preprint arXiv:2010.12619*, 2020.
- [Recht *et al.*, 2019] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet classifiers generalize to ImageNet? In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5389–5400. PMLR, 09–15 Jun 2019.
- [Schede *et al.*, 2019] Elias A. Schede, Samuel Kolb, and Stefano Teso. Learning linear programs from data. In *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 1019–1026, 2019.
- [Schuurmans and Greiner, 1994] Dale Schuurmans and Russell Greiner. Learning default concepts. In *In Proceedings of the Tenth Canadian Conference on Artificial Intelligence (CSCSI-94*, pages 519–523, 1994.
- [Talagrand, 1994] Michel Talagrand. Sharper bounds for gaussian and empirical processes. *The Annals of Probability*, 22(1):28–76, 1994.
- [Valiant, 2000] Leslie G. Valiant. Robust logics. *Artificial Intelligence*, 117(2):231–253, 2000.