

Solving Recurrence Relations using Machine Learning, with Application to Cost Analysis *

MAXIMILIANO KLEMEN

IMDEA Software Institute & Technical University of Madrid (UPM)

MIGUEL Á. CARREIRA-PERPIÑÁN

University of California, Merced

PEDRO LOPEZ-GARCIA

Spanish Council for Scientific Research & IMDEA Software Institute

Abstract

Automatic static cost analysis infers information about the resources used by programs without actually running them with concrete data, and presents such information as functions of input data sizes. Most of the analysis tools for logic programs (and other languages) are based on setting up recurrence relations representing (bounds on) the computational cost of predicates, and solving them to find closed-form functions that are equivalent to (or a bound on) them. Such recurrence solving is a bottleneck in current tools: many of the recurrences that arise during the analysis cannot be solved with current solvers, such as Computer Algebra Systems (CASs), so that specific methods for different classes of recurrences need to be developed. We address such a challenge by developing a novel, general approach for solving arbitrary, constrained recurrence relations, that uses machine-learning sparse regression techniques to *guess* a candidate closed-form function, and a combination of an SMT-solver and a CAS to *check* whether such function is actually a solution of the recurrence. We have implemented a prototype and evaluated it with recurrences generated by a cost analysis system (the one in CiaoPP). The experimental results are quite promising, showing that our approach can find closed-form solutions, in a reasonable time, for classes of recurrences that cannot be solved by such a system, nor by current CASs.

KEYWORDS: Cost Analysis, Recurrence Relations, Static Analysis, Machine Learning, Sparse Linear Regression, Resource Usage Analysis.

1 Introduction and Motivation

The motivation of the work presented in this paper stems from automatic static cost analysis and verification of programs represented as Horn clauses (Debray et al. 1990; Debray and Lin 1993; Debray et al. 1997; Navas et al. 2007; Serrano et al. 2014; Lopez-Garcia et al. 2016; 2018). The goal of such analysis is to infer information about the resources used by programs without actually running them with concrete data, and present such information as functions of input data sizes and possibly other (environmental) parameters. We assume a broad concept of resource as a numerical property of the execution of a program, such as

* This paper is submitted for presentation only. It is under review for ICLP 2023. We thank John Gallagher, Manuel Hermenegildo and José F. Morales for useful discussions and feedback, and their work as developers of the CiaoPP System. Research partially funded by MINECO MICINN PID2019-108528RB-C21 *ProCode* project, the Madrid M141047003 *N-GREENS* and P2018/TCS-4339 *BLOQUES-CM* programs, and the Tezos foundation.




Fig. 1: Control flow diagram of our novel solver based on machine learning.

number of *resolution steps*, *execution time*, *energy consumption*, *memory*, number of *calls* to a predicate, number of *transactions* in a database, etc. Estimating in advance the resource usage of computations is useful for a number of applications, such as automatic program optimization, verification of resource-related specifications, detection of performance bugs, helping developers make resource-related design decisions, security applications (e.g., detection of side channels attacks), or blockchain platforms (e.g., smart-contract gas analysis and verification).

The challenge we address originates from the established approach of setting up recurrence relations representing the cost of predicates, parameterized by input data sizes (Wegbreit 1975; Rosendahl 1989; Debray et al. 1990; Debray and Lin 1993; Debray et al. 1997; Navas et al. 2007; Albert et al. 2011; Serrano et al. 2014; Lopez-Garcia et al. 2016), which are then solved to obtain *closed forms* of such recurrences (i.e., functions that provide either exact, or upper/lower bounds on resource usage in general). Such approach can infer different classes of functions (e.g., polynomial, factorial, exponential, summation, or logarithmic).

The applicability of these resource analysis techniques strongly depends on the capabilities of the component in charge of solving (or safely approximating) the recurrence relations generated during the analysis, which has become a bottleneck in some systems.

A common approach to automatically solving such recurrence relations consists of using a Computer Algebra System (CAS) or a specialized solver to find a closed form. However, this approach poses several difficulties and limitations. For example, some recurrence relations contain complex expressions or recursive structures that most of the well-known CASs cannot solve, making it necessary to develop ad-hoc techniques to handle such cases. Moreover, some recurrences may not have the form required by such systems because an input data size variable does not decrease, but increases instead. Note that a decreasing-size variable could be implicit in the program, i.e., it could be a function of a subset input data sizes (a ranking function), which could be inferred by applying established techniques used in termination analysis (Podelski and Rybalchenko 2004). However, such techniques are usually restricted to linear arithmetics.

In order to address this challenge we have developed a novel, general method for solving arbitrary, constrained recurrence relations. It is a *guess and check* approach that uses machine learning techniques for the *guess* stage, and a combination of an SMT-solver and a Computer Algebra System for the *check* stage (see Figure 1). To the best of our knowledge, there is no other approach that does this. The resulting closed-form function solutions can be of different kinds, such as polynomial, factorial, exponential, summation, or logarithmic.

The rest of this paper is organized as follows. Section 2 gives an overview of our novel *guess and check* approach. Then Section 3 provides some background information and preliminary notation. Section 4 presents a more detailed, formal and algorithmic description of our approach. Section 5 describes the use of our approach in the context of static cost analysis. Section 6 comments on our prototype implementation and its experimental evaluation. Finally, Section 7 summarizes some conclusions and lines for future work.

2 Overview of our Approach

We now give an overview of the two stages of our approach already mentioned: *guess* a candidate closed-form function, and *check* whether such function is actually a solution of the recurrence relation.

Given a recurrence relation for a function $f(\vec{x})$, solving it means to find a closed-form function $\hat{f}(\vec{x})$ that has the same domain as $f(\vec{x})$, and for all \vec{x} in such domain, $\hat{f}(\vec{x}) = f(\vec{x})$. By a closed-form function \hat{f} we mean an expression that is built by using only elementary arithmetic functions, e.g., constants, addition, subtraction, multiplication, division, exponential, or even factorial functions. In particular, this means that \hat{f} does not contain any subexpressions built by using the same function \hat{f} (i.e., \hat{f} is not recursively defined).

We will use the following recurrence as an example to illustrate our approach:

$$\begin{aligned} f(x) &= 0 && \text{if } x = 0 \\ f(x) &= f(f(x-1)) + 1 && \text{if } x > 0 \end{aligned} \quad (1)$$

2.1 The “guess” stage (sparse linear regression via Lasso)

We use a sparse linear regression mechanism (see Sec. 3 for more details), so that any possible model we can obtain (which constitutes a candidate solution) must be a linear combination of a predefined set of terms, but using a usually small subset of terms. That is, a function $\hat{f}(\vec{x})$ of the form:

$$\hat{f}(\vec{x}) = \beta_0 + \beta_1 t_1(\vec{x}) + \beta_2 t_2(\vec{x}) + \cdots + \beta_n t_n(\vec{x})$$

where the t_i ’s are arbitrary functions on \vec{x} from a set T of candidate terms that we call *base functions*, and the β_i ’s are the coefficients (real numbers) that are estimated by regression, but so that only a few coefficients are nonzero. Currently, the set T is fixed, and contains the base functions that are representative of the common complexity orders (in Sec. 7 we comment on future plans to obtain it). For illustration purposes, assume that we use the following set T of base functions:

$$T = \{\lambda x.x, \lambda x.x^2, \lambda x.x^3, \lambda x.\lceil \log_2(x) \rceil, \lambda x.2^x, \lambda x.x \cdot \lceil \log_2(x) \rceil\}$$

where each base function is represented as a lambda expression. Then, the sparse linear regression is performed as follows:

1. Generate a training set S . First, a set $X_{\text{train}} = \{\vec{x}_1, \dots, \vec{x}_k\}$ of input values to the recurrence function is randomly generated. Then, starting with an initial $S = \emptyset$, for each input value $\vec{x}_i \in X_{\text{train}}$, a training case s_i is generated and added to S . For any input value $\vec{x} \in X_{\text{train}}$ the corresponding training case s is a tuple of the form:

$$s = \langle b, c_1, \dots, c_n \rangle$$

where $c_i = \llbracket t_i \rrbracket_{\vec{x}}$ for $1 \leq i \leq n$, and $\llbracket t_i \rrbracket_{\vec{x}}$ represents the result (a scalar) of evaluating the base function $t_i \in T$ for input value \vec{x} , where T is a set of n base functions, as

already explained. The (dependent) value b (also a constant number) is the result of evaluating the recurrence $f(\vec{x})$ that we want to solve or approximate, in our example, the one defined in Eq. 1. Assuming that there is an $\vec{x} \in X_{\text{train}}$ such that $\vec{x} = \langle 5 \rangle$, its corresponding training case s in our example will be:

$$\begin{aligned} s &= \langle \mathbf{f}(\mathbf{5}), \llbracket x \rrbracket_5, \llbracket x^2 \rrbracket_5, \llbracket x^3 \rrbracket_5, \llbracket \lceil \log_2(x) \rceil \rrbracket_5, \dots \rangle \\ &= \langle \mathbf{5}, 5, 25, 125, 3, \dots \rangle \end{aligned}$$

2. Perform the sparse regression in two steps using the training set S created above. In the first step, we use linear regression with Lasso (ℓ_1) regularization (Hastie et al. 2015) on the coefficients. This is a penalty term that encourages coefficients whose associated base functions have a small correlation with the dependent value to be exactly zero. This way, typically most of the candidate terms in T will be discarded, and only those that are really needed to approximate our target function will be kept. The level of penalization is controlled by a hyperparameter $\lambda \geq 0$. As commonly done in machine learning (Hastie et al. 2015), the value of λ that generalizes optimally on unseen (test) inputs is found via cross-validation on a separate validation set (generated randomly in the same way as the training set). The result of this step is a (column) vector $\vec{\beta}$ of coefficients, and an independent coefficient β_0 . Finally, we generate a test set X_{test} (again, randomly in the same way as the training set) of input values to the recurrence function to obtain a measure R^2 of the accuracy of the estimation. Additionally, we discard those terms whose corresponding coefficient is less than a given threshold ϵ . The resulting closed-form expression that estimates the target function is

$$\hat{f}(\vec{x}) = rm_{\epsilon}(\vec{\beta}^T) \cdot E(T, \vec{x}) + \beta_0$$

where $E(T, \vec{x})$ is a vector of the terms in T with the arguments bound to \vec{x} , and rm_{ϵ} takes a vector of coefficients and returns another vector where the coefficients less than ϵ are rounded to zero. Both the Lasso regularization and the pruning function discard many terms from T in the final cost function.

3. Finally, our method performs again a standard linear regression (without Lasso regularization) on the training set S , but without using those base functions corresponding to the terms discarded previously by Lasso and the ϵ -pruning. In our example, with $\epsilon = 0.001$, we obtain:

$$\hat{f}(x) = 1.0 \ x$$

with a value $R^2 = 1$, which means that the estimation obtained predicts exactly the values for the test set, and thus, it is a candidate solution for the recurrence in Eq. 1. If R^2 were less than 1, it would mean that the function obtained is not a candidate (exact) solution, but a (possibly unsafe) approximation, as there are values in the test set that cannot be exactly predicted.

2.2 The “check” stage

Once a function that is a candidate solution for the recurrence has been obtained, the second step of our method tries to verify whether such a candidate is actually a solution. To do so, the recurrence is encoded as a first order logic formula where the references to the target function are replaced by the candidate solution whenever possible. Afterwards, we use an SMT-solver to check whether the negation of such formula is satisfiable, in which case we can conclude that the candidate is not a solution for the recurrence. Otherwise, if such formula is unsatisfiable, then the candidate function is an exact solution. Sometimes,

it is necessary to consider a precondition for the domain of the recurrence, which is also included in the encoding.

To illustrate this process, Expression (2) below shows the recurrence relation we target to solve, followed by the candidate solution obtained previously using linear regression:

$$\begin{aligned} f(x) &= 0 && \text{if } x = 0 \\ f(x) &= f(f(x-1)) + 1 && \text{if } x > 0 \\ \hat{f}(x) &= x && \text{if } x \geq 0 \end{aligned} \quad (2)$$

Now, Expression (3) below shows the encoding of the recurrence as a first order logic formula.

$$\forall x \cdot \left((x = 0 \implies \underline{f(x)} = 0) \wedge (x > 0 \implies \underline{f(x)} = \underline{f(\underline{f(x-1)})} + 1) \right) \quad (3)$$

Finally, Expression (4) below shows the negation of such formula, as well as the references to the function name substituted by the definition of the candidate solution. We underline both the subexpressions to be replaced, and the subexpressions resulting from the substitutions.

$$\exists x \cdot \neg((x = 0 \implies \underline{x} = 0) \wedge (x > 0 \implies \underline{x} = \underline{x-1} + 1)) \quad (4)$$

It is easy to see that Formula (4) is unsatisfiable. Therefore, $\hat{f}(x) = x$ is an exact solution for $f(x)$ in Eq. 1.

For some cases where the candidate solution contains transcendental functions, our implementation of the method uses a Computer Algebra System to perform simplifications and transformations, in order to obtain a formula supported by the SMT-solver. We find this combination of CAS and SMT-solver particularly useful, since it allows solving more problems than only using one of these systems in isolation.

3 Preliminaries

Recurrence relations. A recurrence relation of order k , $k > 0$, for a function f , is a set of equations that give k initial values for f , and an equation that recursively defines any other value of f as a function g that takes k previous values of f as parameters. For example, the following recurrence relation of second order ($k = 2$), with g being the arithmetic addition $+$, defines the Fibonacci function:

$$f(n) = \begin{cases} 1 & \text{if } n = 0 \text{ or } n = 1 \\ f(n-1) + f(n-2) & \text{if } n \geq 2 \end{cases} \quad (5)$$

A challenging class of recurrences that we can solve with our approach are “nested” recurrences, e.g., recurrences of the form $f(n) = g(f(f(n-1)))$.

We use the last letters from the alphabet to denote variables, and the first letters from the alphabet to denote constants and coefficients. We use f, g to represent functions, and e, t to represent arbitrary expressions. We use φ to represent arbitrary boolean constraints over a set of variables. Sometimes, we also use β to represent coefficients obtained with linear regression. In all cases, the symbols can be subscripted. We use \vec{x} to denote a finite sequence $\langle x_1, x_2, \dots, x_n \rangle$, for some $n > 0$. Given a sequence S and an element x , $\langle x|S \rangle$ is a new sequence with first element x and tail S .

Given a piecewise function:

$$f(\vec{x}) = \begin{cases} e_1(\vec{x}) & \text{if } \varphi_1(\vec{x}) \\ e_2(\vec{x}) & \text{if } \varphi_2(\vec{x}) \\ \vdots & \vdots \\ e_k(\vec{x}) & \text{if } \varphi_k(\vec{x}) \end{cases} \quad (6)$$

where $f \in \mathcal{D} \rightarrow \mathcal{R}^+$, with $\mathcal{D} = \{\vec{x} \mid \vec{x} \in \mathcal{Z}^m \wedge \varphi_{pre}(\vec{x})\}$ for some boolean constraint φ_{pre} , and $e_i(\vec{x}), \varphi_i(\vec{x})$ are arbitrary expressions and constraints over \vec{x} respectively. We say that φ_{pre} is the *precondition* of f , and that f is a *constrained recurrence relation* if and only if:

- $\exists i \in [1, k]$ such that e_i contains a call to f .
- $\exists i \in [1, k]$ such that e_i does not contain any call to f (i.e., it is in *closed form*).
- $\varphi_{pre} \models \bigvee_{1 \leq i \leq k} \varphi_i$.

Given a concrete input $\vec{d} \in \mathcal{D}$, we evaluate $f(\vec{d})$ deterministically, assuming the evaluation of f as a nested *if-then-else* control structure as follows:

```

if  $\varphi_1(\vec{d})$  then
  return  $e_1(\vec{d})$ 
else
  if  $\varphi_2(\vec{d})$  then
    return  $e_2(\vec{d})$ 
  else
    ...
  end if
end if

```

More formally, let $def(f)$ denote the definition of a (piecewise) constrained recurrence relation f represented as the sequence $\langle (e_1(\vec{x}), \varphi_1(\vec{x})), \dots, (e_k(\vec{x}), \varphi_k(\vec{x})) \rangle$, where each element of the sequence is a pair representing a case. The order of such sequence determines the evaluation strategy. Then, the evaluation of f for a concrete value \vec{d} , denoted $EvalFun(f(\vec{d}))$, is defined as follows:

$$EvalFun(f(\vec{d})) = EvalBody(def(f), \vec{d})$$

$$EvalBody(\langle (e, \varphi) \mid Ps \rangle, \vec{d}) = \begin{cases} \llbracket e \rrbracket_{\vec{d}} & \text{if } \varphi(\vec{d}) \\ EvalBody(Ps, \vec{d}) & \text{if } \neg \varphi(\vec{d}) \end{cases}$$

Our goal is to find a function $\hat{f} \in \mathcal{D} \rightarrow \mathcal{R}^+$ such that for all $\vec{d} \in \mathcal{D}$:

- If $EvalFun(f(\vec{d}))$ terminates, then $EvalFun(f(\vec{d})) = \llbracket \hat{f} \rrbracket_{\vec{d}}$, and
- \hat{f} does not contain any recursive call in its definition.

In particular, we look for a definition of the form:

$$\hat{f}(\vec{x}) = \beta_0 + \beta_1 t_1(\vec{x}) + \beta_2 t_2(\vec{x}) + \dots + \beta_n t_n(\vec{x}) \quad (7)$$

where $\beta_i \in \mathcal{R}$, and t_i are expressions over \vec{x} , not including recursive references to \hat{f} . If the above conditions are met, we say that \hat{f} is a *closed form* for f .

To illustrate the need of introducing an evaluation strategy for the recurrence that is consistent with the termination of the program, consider the following Prolog program which does not terminate for a call $p(X)$ where X is bound to an integer:

```

1 p(X) :- X > 0, X1 is X + 1, p(X1).
2 p(X) :- X = 0.

```

The following recurrence relation for its cost (in resolution steps) can be set up:

$$\begin{aligned} C_p(x) &= 1 & \text{if } x = 0 \\ C_p(x) &= 1 + C_p(x+1) & \text{if } x > 0 \end{aligned} \quad (8)$$

A CAS will give the closed form $C_p(x) = 1 - x$ for such recurrence, however, the cost analysis should give $C_p(x) = \infty$.

Linear Regression. Linear regression (Hastie et al. 2009) is a statistical technique used to approximate the linear relationship between a number of independent variables and a dependent (output) variable. Given a vector of independent (input) variables $X = (X_1, \dots, X_p)^T \in \mathcal{R}^p$, we predict the output variable Y using the formula

$$Y = \beta_0 + \sum_{i=1}^p \beta_i X_i \quad (9)$$

which is defined through the vector of coefficients $\beta = (\beta_0, \dots, \beta_p)^T \in \mathcal{R}^p$. Such coefficients are estimated from a set of observations $\{y_i, x_{i1}, \dots, x_{ip}\}_{i=1}^n$ so as to minimize a loss function, most commonly the sum of squares

$$\beta = \arg \min_{\beta \in \mathcal{R}^p} \sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j \right)^2 \quad (10)$$

Sometimes (as is our case) some of the input variables are not relevant to explain the output, but the above least-squares estimate will almost always assign nonzero values to all the coefficients. In order to force the estimate to make exactly zero the coefficients of irrelevant variables (hence removing them and doing *feature selection*), various techniques have been proposed. The most widely used one is the Lasso (Hastie et al. 2015), which adds an ℓ_1 penalty on β (i.e., the sum of absolute values of each coefficient) to Eq. 10:

$$\beta = \arg \min_{\beta \in \mathcal{R}^p} \sum_{i=1}^n \left(y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j \right)^2 + \lambda \sum_{j=1}^p |\beta_j| \quad (11)$$

where $\lambda \geq 0$ is a hyperparameter that determines the level of penalization: the greater λ , the greater the number of coefficients that are exactly equal to 0. The Lasso has two advantages over other feature selection techniques for linear regression. First, it defines a convex problem whose unique solution can be efficiently computed even for datasets where either of n or p are large (almost as efficiently as a standard linear regression). Second, it has been shown in practice to be very good at estimating the relevant variables.

4 Algorithmic Description of the Approach

In this section we describe our approach for generating and checking candidate solutions for recurrences that arise in resource analysis. Algorithms 1 and 2 correspond to the *guesser* and *checker* components, respectively, which are shown in Figure 1.

Algorithm 1: Candidate Solution Generation (*Guesser*).

Input : $F \in \mathcal{D} \rightarrow \mathcal{R}^+$: target recurrence relation.
 φ_{pre} : precondition defining \mathcal{D} .
 $T \subseteq \mathcal{D} \rightarrow \mathcal{R}^+$: candidate terms.
 Λ : range of values to automatically choose a Lasso hyperparameter $\lambda \in \mathcal{R}^+$.
 k : indicates performing k -fold cross-validation, $k > 2$.
 $\epsilon \in \mathcal{R}^+$: threshold for term ($t_i \in T$) selection.

Output: $\hat{F} \in Exp$: a candidate solution (or an approximation) for F .
 $S \in [0, 1]$: score, indicates the accuracy of the estimation (R^2).

```

1  $\mathcal{I} \leftarrow \{\vec{x}_i | \vec{x}_i \in \mathcal{Z}^m \wedge \varphi_{pre}(\vec{x}_i)\}_{i=1}^N$  ; // N Random inputs for F
2  $\mathcal{X} \leftarrow \{(F(\vec{x}) | E(T, \vec{x})) | \vec{x} \in \mathcal{I}\}$  ; // Training set
3  $(\vec{\beta}', \beta'_0) \leftarrow CVLassoRegression(\mathcal{X}, \Lambda, k)$ ;
4  $(T', \mathcal{X}') \leftarrow RemoveTerms(T, \mathcal{X}, \vec{\beta}', \beta'_0, \epsilon)$ ;
5  $(\vec{\beta}, \beta_0, S) \leftarrow LinearRegression(\mathcal{X}')$ ;
6  $\hat{F} \leftarrow \lambda \vec{x} \cdot \vec{\beta}^T \times E(T', \vec{x}) + \beta_0$ ;
7 return  $(\hat{F}, S)$ ;
```

Algorithm 1 receives a recurrence relation for a function F to solve, a set of candidate terms, and a threshold to decide when to discard irrelevant terms. The output is a closed-form expression \hat{F} for F , and a *score* S that reflects the accuracy of the approximation, in the range $[0, 1]$. If $S \sim 1$, the approximation can be considered a candidate solution. Otherwise, \hat{F} is a (possibly unsafe) approximation. In line 1 we start by generating a set \mathcal{I} of random inputs for F . Each input \vec{x}_i is a m -tuple verifying precondition φ_{pre} , where m is the number of arguments of F . In line 2 we produce the training set \mathcal{X} . The independent inputs are generated by evaluating the candidate terms in $T = \langle t_1, t_2, \dots, t_p \rangle$ with each tuple $\vec{x} \in \mathcal{I}$. This is done by using function E , defined as follows:

$$E(\langle t_1, t_2, \dots, t_p \rangle, \vec{x}) = \langle t_1(\vec{x}), t_2(\vec{x}), \dots, t_p(\vec{x}) \rangle$$

We also evaluate the recurrence equation for input \vec{x} , and add the observed output $F(\vec{x})$ as the first element in the vectors of the training set. In line 3 we generate a first linear model by applying function *CVLassoRegression* to the generated training set. *CVLassoRegression* performs a linear regression with Lasso regularization. As already mentioned, Lasso regularization requires a hyperparameter λ that determines the level of penalization for the coefficients. Instead of using a single value for λ , *CVLassoRegression* uses a range of possible values, applying cross-validation on top of the linear regression to automatically select the best value for that parameter, from the given range. The result of this function is the vector of coefficients $\vec{\beta}'$, together with the intercept β'_0 . These coefficients are used in line 4 to decide which candidate terms are discarded before the last regression step. Note that *RemoveTerms* removes the candidate terms from T together with their corresponding input values from the training set \mathcal{X} , returning the new set of candidate terms T' and its corresponding training set \mathcal{X}' . In line 5, standard linear regression (without regularization nor cross-validation) is applied, obtaining the final coefficients $\vec{\beta}$ and β_0 . Additionally, from this step we also obtain the score S of the resulting model. In line 6 we set up the resulting closed-form expression, given as a function on the variables in \vec{x} . Note that we use the function E to bind the variables in the candidate terms to the arguments of the closed-form expression. Finally, the closed-form expression and its corresponding score are returned as the result of the algorithm.

Algorithm 2 mainly relies on an SMT solver and a Computer Algebra System. Concretely, given the constrained recurrence relation $F \in \mathcal{D} \rightarrow \mathcal{R}^+$ defined as

$$F(\vec{x}) = \begin{cases} e_1(\vec{x}) & \text{if } \varphi_1(\vec{x}) \\ e_2(\vec{x}) & \text{if } \varphi_2(\vec{x}) \\ \vdots & \vdots \\ e_k(\vec{x}) & \text{if } \varphi_k(\vec{x}) \end{cases}$$

our algorithm constructs the logic formula:

$$\left\| \bigwedge_{i=1}^k \left(\left(\bigwedge_{j=1}^{i-1} \neg \varphi_j(\vec{x}) \right) \wedge \varphi_i(\vec{x}) \wedge \varphi_{pre}(\vec{x}) \implies Eq_i \right) \right\|_{smt} \quad (12)$$

where Eq_i is the result of replacing in $F(\vec{x}) = e_i(\vec{x})$ each occurrence of F , if possible, by the definition of the candidate solution \hat{F} (by using *replaceCalls* in line 4), and performing a simplification by the CAS (by using *simplifyCAS* in line 6). A goal of such simplification is to obtain (sub)expressions supported by the SMT-solver. The function *replaceCalls*(*expr*, $F(\vec{x}')$, \hat{F} , φ_{pre} , φ) replaces every subexpression in *expr* of the form $F(\vec{x}')$ by $\hat{F}(\vec{x}')$, if $\varphi_{pre}(\vec{x}') \wedge \varphi \implies \varphi_{pre}(\vec{x}')$. The operation $\llbracket e \rrbracket_{smt}$ is the translation of any expression e to a SMT-LIB expression. Although all variables appearing in 12 are declared as integers, we omit these details in Algorithm 2 and in Formula 12 for the sake of brevity. Note that this encoding is consistent with the evaluation (*EvalFun*) described in Section 3. Finally, the algorithm asks the SMT solver for models of the negated formula (line 17). If no model exists, then it returns *true*, concluding that \hat{F} is an exact solution to the recurrence, i.e., $\hat{F}(\vec{x}) = F(\vec{x})$ for any input $\vec{x} \in \mathcal{D}$ such that *EvalFun*($F(\vec{x})$) terminates. Otherwise, it returns *false*. Note that, if it is not possible to replace all occurrences of F by \hat{F} , or if after performing the simplification by *simplifyCAS* there are subexpressions not supported by the SMT solver, then the algorithm finishes returning *false*.

5 Our Approach in the Context of Static Cost Analysis

In this section, we describe how our approach could be used in the context of the motivating application, Static Cost Analysis. Although it is general, and could be integrated into any cost analysis system based on recurrence solving, we illustrate its use in the context of the CiaoPP system. Using a logic program, we first illustrate how the CiaoPP system sets up recurrence relations representing the sizes of output arguments of predicates and the cost of such predicates. Then, we show how our novel approach is used to solve a recurrence relation that cannot be solved by CiaoPP.

Example 1

Consider predicate `p/2` in Figure 2, and calls to it where the first argument is bound to a non-negative integer and the second one is a free variable. Upon success of these calls, the second argument is bound to a non-negative integer too. Such calling mode, where the first argument is input and the second one is output, is automatically inferred by CiaoPP (see (Hermenegildo et al. 2005) and its references).

The CiaoPP system first infers size relations for the different arguments of predicates, using a rich set of size metrics (see (Navas et al. 2007; Serrano et al. 2014) for details). Assume that the size metric used in this example, for the numeric argument `X` is the *actual*

Algorithm 2: Solution Checking (*Checker*).

Input : $F \in \mathcal{D} \rightarrow \mathcal{R}^+$: target recurrence relation.
 φ_{pre} : precondition defining \mathcal{D} .
 $\hat{F} \in Exp$: a candidate solution for F .
Output: *true* if \hat{F} is a solution for F , *false* otherwise.

```

1  $\varphi_{previous} \leftarrow true$  ;
2  $Formula \leftarrow true$  ;
3 foreach  $(e, \varphi) \in def(F)$  do
4    $Eq \leftarrow replaceCalls("F(\vec{x}) - e = 0", F(\vec{x}), \hat{F}, \varphi_{pre}, \varphi)$ ;
5   if  $\neg containsCalls(Eq, F)$  then
6      $Eq \leftarrow simplifyCAS(inlineCalls(Eq, \hat{F}, def(\hat{F})))$ ;
7     if  $supportedSMT(Eq)$  then
8        $Formula \leftarrow "Formula \wedge (\varphi_{pre} \wedge \varphi_{previous} \wedge \varphi \implies Eq)"$ ;
9        $\varphi_{previous} \leftarrow "\varphi_{previous} \wedge \neg \varphi"$  ;
10    else
11      return false;
12    end
13  else
14    return false;
15  end
16 end
17 return  $(\not\models_{SMT} [\neg Formula]_{SMT})$ ;

```

```

1 :- entry p/2: nnegint*var.
2 p(X,0):-
3   X=0.
4 p(X,Y):-
5   X>0,
6   X1 is X - 1,
7   p(X1,Y1),
8   p(Y1,Y2),
9   Y is Y2 + 1.

```

Fig. 2: A program with a nested recursion.

value of it (denoted $\text{int}(X)$). The system will try to infer a function $S_p(x)$ that gives the size of the output argument of $p/2$ (the second one), as a function of the size (x) of the input argument (the first one). For this purpose, the following size relations for $S_p(x)$ are automatically set up (the same as Eq. 1 used in Sec. 2 as example):

$$\begin{aligned} S_p(x) &= 0 & \text{if } x = 0 \\ S_p(x) &= S_p(S_p(x-1)) + 1 & \text{if } x > 0 \end{aligned} \quad (13)$$

The first and second recurrence correspond to the first and second clauses respectively (i.e., base and recursive cases). Once recurrence relations (either representing the size of terms, as the ones above, or the computational cost of predicates, as the ones that we will see later) have been set up, a solving process is started.

Nested recurrences, as the one that arise in this example, cannot be handled by most state-of-the-art recurrence solvers. In particular, the modular solver used by CiaoPP fails to find a closed-form function for the recurrence relation above. In contrast, the novel

approach that we propose, squetched in next section, obtains the closed form $\hat{S}_p(x) = x$, which is an exact solution of such recurrence (as shown in Section 2).

Once the size relations have been inferred, CiaoPP uses them to infer the computational cost of a call to $p/2$. For simplicity, assume that in this example, such cost is given in terms of the number of *resolution steps*, as a function of the size of the input argument, but note that CiaoPP’s cost analysis is parametric with respect to resources, which can be defined by the user by means of a rich assertion language, so that it can infer a wide range of resources, besides resolution steps. Also for simplicity, we assume that all builtin predicates, such as arithmetic/comparison operators have zero cost (in practice there is a “trust” assertion for each builtin that specifies its cost as if it had been inferred by the analysis).

In order to infer the cost of a call to $p/2$, represented as $C_p(x)$, CiaoPP sets up the following cost relations, by using the size relations inferred previously:

$$\begin{aligned} C_p(x) &= 1 & \text{if } x = 0 \\ C_p(x) &= C_p(x-1) + C_p(S_p(x-1)) + 1 & \text{if } x > 0 \end{aligned} \quad (14)$$

We can see that the cost of the second recursive call to predicate $p/2$ depends on the size of the output argument of the first recursive call to such predicate, which is given by function $S_p(x)$, whose closed form $S_p(x) = x$ is computed by our approach, as already explained. Pluging such closed form into the recurrence relation above, it can be solved now by CiaoPP, obtaining $C_p(x) = 2^{x+1} - 1$.

6 Implementation and Experimental Evaluation

We have implemented a prototype of our novel approach and performed an experimental evaluation in the context of the CiaoPP system, by solving recurrences generated during static cost analysis. Our prototype takes a recurrence and returns a closed form obtained together with two measures: 1) the accuracy of the estimation (*score*) of the candidate closed-form solution generated by the machine learning phase, and 2) an indication of whether such closed form is an exact solution of the recurrence (i.e., if it has been formally verified). It is implemented in Python 3, using Sympy (Meurer et al. 2017) as Computer Algebra System, and Scikit-Learn (Pedregosa et al. 2011) for the regression with Lasso regularization. We use Z3 (de Moura and Bjørner 2008) as SMT-Solver, and Z3Py (Z3Py 2010) as interface. The solver is pre-configured with a set of global parameters:

- An integer $k > 2$, to perform k -fold cross-validation. This means that the training set is split into k parts or *folds*. Then, each fold is taken as the validation set, training the model with the remaining $k - 1$ folds. Finally, the performance measure reported is the average of the values computed in the k iterations.
- A range of real values Λ , to automatically choose a λ for Lasso regularization that maximizes the performance of the model via cross-validation.
- A set of basic symbolic functions T , to form the candidate terms t_i to be used in the expression obtained by the algorithm.
- Optionally, a precondition φ_{pre} on the arguments of the recurrence to solve.

Our experimental results are shown in Table 1. Column **Bench** shows the name that we have assigned to each recurrence that we have chosen (which is inspired by the Horn-clause program such recurrence originated from during cost/size analysis), and Column **Recurrence** shows their definitions, where we use the same function symbol, f , for all of them. Such recurrences are challenging for CiaoPP, either because they cannot be solved by any of the back-end solvers, or because they are necessarily over-estimated in the solving

Table 1: Closed forms obtained with the previous (**CF**) and new solver (**CFNew**).

Bench	Recurrence	CF	CFNew	T (s)
merge-sz	$f(x, y) = \begin{cases} \max(f(x-1, y), & \text{if } x > 0 \wedge y > 0 \\ f(x, y-1)) + 1 & \text{if } x > 0 \wedge y \leq 0 \\ x & \text{if } x \leq 0 \wedge y > 0 \\ y & \end{cases}$	—	$x + y$	0.92
merge	$f(x, y) = \begin{cases} \max(f(x-1, y), & \text{if } x > 0 \wedge y > 0 \\ f(x, y-1)) + 1 & \text{otherwise} \\ 0 & \end{cases}$	—	$\max(0, x + y - 1)$	0.71
nested	$f(x) = \begin{cases} f(f(x-1)) + 1 & \text{if } x > 0 \\ 0 & \text{otherwise} \end{cases}$	—	x	0.13
open-zip	$f(x, y) = \begin{cases} f(x-1, y-1) + 1 & \text{if } x > 0 \wedge y > 0 \\ f(x, y-1) + 1 & \text{if } x \leq 0 \wedge y > 0 \\ f(x-1, y) + 1 & \text{if } y \leq 0 \wedge x > 0 \\ 0 & \text{otherwise} \end{cases}$	—	$\max(x, y)$	0.12
div	$f(x, y) = \begin{cases} f(x-y, y) + 1 & \text{if } x \geq y \\ 0 & \text{otherwise} \end{cases}$	—	$\left\lfloor \frac{x}{y} \right\rfloor$	0.13
div-ceil	$f(x, y) = \begin{cases} f(x-y, y) + 1 & \text{if } x \geq y \\ 1 & \text{if } x < y \wedge x > 0 \\ 0 & \text{otherwise} \end{cases}$	—	$\left\lceil \frac{x}{y} \right\rceil$	0.12
s-max	$f(x, y) = \begin{cases} \max(y, f(x-1, y)) + 1 & \text{if } x > 0 \\ y & \text{otherwise} \end{cases}$	$x + y$	$x + y$	0.12
s-max-1	$f(x, y) = \begin{cases} \max(y, f(x-1, y+1)) + 1 & \text{if } x > 0 \\ y & \text{otherwise} \end{cases}$	—	$2x + y$	0.14
sum-osc	$f(x, y) = \begin{cases} f(x-1, y) + 1 & \text{if } x > 0 \wedge y > 0 \\ f(x+1, y-1) + y & \text{if } x \leq 0 \wedge y > 0 \\ 1 & \text{otherwise} \end{cases}$	—	$x + \frac{y^2}{2} + \frac{3y}{2}$	0.13

process. Some recurrences, like **nested**, are problematic even for most of the current state-of-the-art solvers. Column **CF** shows the closed forms obtained by our previous recurrence solver, and Column **CFNew** shows the closed forms obtained by our approach, applying Algorithms 1 and 2. All of them have been verified as exact solutions to the recurrences by Algorithm 2.

Finally, Column **T(s)** shows the total time, in seconds (executing on a MacBook Pro machine, 2.4GHz Intel Core i7 CPU, 8 GB 1333 MHz DDR3 memory), needed to obtain the closed forms and verify them. For all the experiments, we have set $k = 2$, in order to perform 2-fold cross-validation. We have also set the range for λ to 100 values taken from the interval $[0.001 \ 1]$. Regarding the set T of candidate terms, for recurrences with one or two arguments, we provide a predefined set of representative functions of the most common complexity orders, as well as some compositions of them. For recurrences with three or more arguments, we provide an initial set of simple functions, that are combined automatically to generate the basic functions t_i for the set T . Finally, as a default precondition, we assume that the initial values for the variables that are the arguments of the recurrences are all

greater than or equal to zero, i.e., $\varphi_{pre} = \bigwedge_{x \in \text{Args}} (x \geq 0)$, where *Args* is the set of arguments of the recurrence.

As we can see, none of the recurrences are solvable by the current CiaoPP solver, except **s-max**. The specialized solver for such recurrence has been developed relatively recently. In contrast, our new solver is able to infer exact closed-forms functions for all the recurrences in a reasonable time.

7 Conclusions and Future Work

We have developed a novel approach for solving or approximating arbitrary, constrained recurrence relations. It consists of a *guess* stage that uses a sparse linear regression via Lasso regularization and cross-validation to infer a candidate closed-form solution, and a *check* stage that combines a SMT-solver and a Computer Algebra System to verify that such candidate is actually a solution. We have implemented a prototype and evaluated it with recurrences that are generated by the cost analysis module of the CiaoPP system, and are not solvable by it. The experimental results are quite promising, showing that our approach can find exact, verified, closed-form solutions, in a reasonable time, for such recurrences.

Since our technique uses linear regression with a randomly generated training set (by evaluating the recurrence to obtain the dependent value), it is not guaranteed that a solution can be found. Even if an exact solution is found in the first stage, it is not always possible to prove its correctness in the second stage. Therefore, in this sense, this approach is *not complete*. However, it is able to find some solutions that current state-of-the-art solvers are unable to find. As a proof of concept, we have considered a particular deterministic evaluation for constrained recurrence relations, and the verification of the candidate solution is consistent with this evaluation. However, it is possible to implement different evaluation semantics for the recurrences, adapting the verification stage accordingly. Note that we need to require the termination of the recurrence evaluation as a precondition for the conclusions obtained. This is also due to the particular evaluation strategy of recurrences that we are considering. In practice, non-terminating recurrences can be discarded in the first stage, by setting a timeout. Our approach can also be combined with a termination prover in order to guarantee such precondition. Finally, note that an alternative use of our tool is to omit the verification stage, using only the closed-form function inferred by the first stage, together with an error measure. This can be useful in some applications (e.g., granularity control in parallel/distributed computing) where it is enough to have good although unsafe approximations.

As a future work, we plan to fully integrate our novel solver into the CiaoPP system, combining it with its current set of back-end solvers in order to improve the static cost analysis. We also plan to further refine and improve our algorithms in several directions. As already explained, currently the set *T* of base functions is fixed, user-provided. We plan to automatically infer it by using different heuristics. We can perform an automatic analysis of the recurrence we are solving, to extract some features that allow selection of the terms that most likely are part of the solution. For example, if the recurrence has a nested, double recursion, then we can select a quadratic term, etc. Also, machine learning techniques may be applied to learn a good set of base functions from some features of the programs.

References

- ALBERT, E., ARENAS, P., GENAIM, S., AND PUEBLA, G. 2011. Closed-Form Upper Bounds in Static Cost Analysis. *Journal of Automated Reasoning*, 46, 2, 161–203.
- DE MOURA, L. M. AND BJØRNER, N. Z3: An Efficient SMT Solver. In RAMAKRISHNAN, C. R. AND REHOF, J., editors, *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008* 2008, volume 4963 of *Lecture Notes in Computer Science*, pp. 337–340. Springer.
- DEBRAY, S. K. AND LIN, N. W. 1993. Cost analysis of logic programs. *ACM TOPLAS*, 15, 5, 826–875.
- DEBRAY, S. K., LIN, N.-W., AND HERMENEGILDO, M. V. Task Granularity Analysis in Logic Programs. In *Proc. PLDI'90* 1990, pp. 174–188. ACM.
- DEBRAY, S. K., LOPEZ-GARCIA, P., HERMENEGILDO, M. V., AND LIN, N.-W. Lower Bound Cost Estimation for Logic Programs. In *ILPS'97* 1997, pp. 291–305. MIT Press.
- HASTIE, T., TIBSHIRANI, R., AND WAINWRIGHT, M. 2015. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman & Hall/CRC.
- HASTIE, T. J., TIBSHIRANI, R. J., AND FRIEDMAN, J. H. 2009. *The Elements of Statistical Learning—Data Mining, Inference and Prediction*. Springer-Verlag, second edition.
- HERMENEGILDO, M. V., PUEBLA, G., BUENO, F., AND LOPEZ-GARCIA, P. 2005. Integrated Program Debugging, Verification, and Optimization Using Abstract Interpretation (and The Ciao System Preprocessor). *Science of Computer Programming*, 58, 1–2, 115–140.
- LOPEZ-GARCIA, P., DARMAWAN, L., KLEMEN, M., LIQAT, U., BUENO, F., AND HERMENEGILDO, M. V. 2018. Interval-based Resource Usage Verification by Translation into Horn Clauses and an Application to Energy Consumption. *Theory and Practice of Logic Programming, Special Issue on Computational Logic for Verification*, 18, 2, 167–223.
- LOPEZ-GARCIA, P., KLEMEN, M., LIQAT, U., AND HERMENEGILDO, M. V. 2016. A General Framework for Static Profiling of Parametric Resource Usage. *TPLP (ICLP'16 Special Issue)*, 16, 5-6, 849–865.
- MEURER, A., SMITH, C. P., PAPROCKI, M., ČERTÍK, O., KIRPICHEV, S. B., ROCKLIN, M., KUMAR, A., IVANOV, S., MOORE, J. K., SINGH, S., RATHNAYAKE, T., VIG, S., GRANGER, B. E., MULLER, R. P., BONAZZI, F., GUPTA, H., VATS, S., JOHANSSON, F., PEDREGOSA, F., CURRY, M. J., TERREL, A. R., ROUČKA, V., SABOO, A., FERNANDO, I., KULAL, S., CIMRMAN, R., AND SCOPATZ, A. 2017. Sympy: symbolic computing in python. *PeerJ Computer Science*, 3, e103.
- NAVAS, J., MERA, E., LOPEZ-GARCIA, P., AND HERMENEGILDO, M. User-Definable Resource Bounds Analysis for Logic Programs. In *Proc. of ICLP'07* 2007, volume 4670 of *LNCS*, pp. 348–363. Springer.
- PEDREGOSA, F., VAROQUAUX, G., GRAMFORT, A., MICHEL, V., THIRION, B., GRISEL, O., BLONDEL, M., PRETTENHOFER, P., WEISS, R., DUBOURG, V., VANDERPLAS, J., PASSOS, A., COURNAPEAU, D., BRUCHER, M., PERROT, M., AND DUCHESNAY, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- PODELSKI, A. AND RYBALCHENKO, A. A Complete Method for the Synthesis of Linear Ranking Functions. In *5th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI'04)* 2004, Lecture Notes in Computer Science, pp. 239–251. Springer.
- ROSENDAHL, M. Automatic Complexity Analysis. In *Proc. of FPCA'89* 1989, pp. 144–156. ACM Press.
- SERRANO, A., LOPEZ-GARCIA, P., AND HERMENEGILDO, M. V. 2014. Resource Usage Analysis of Logic Programs via Abstract Interpretation Using Sized Types. *TPLP, ICLP'14 Special Issue*, 14, 4-5, 739–754.
- WEGBREIT, B. 1975. Mechanical Program Analysis. *Comm. of the ACM*, 18, 9, 528–539.
- Z3Py 2010. Z3 api in python. <https://ericpony.github.io/z3py-tutorial/guide-examples.htm>. Accessed: 2010-09-30.