Journal Pre-proofs

Analytic Continued Fractions for Regression: A Memetic Algorithm Approach

Pablo Moscato, Haoyuan Sun, Mohammad Nazmul Haque

PII: S0957-4174(21)00459-0
DOI: <https://doi.org/10.1016/j.eswa.2021.115018>
Reference: ESWA 115018

To appear in: *Expert Systems with Applications*

Received Date: 11 February 2020
Accepted Date: 6 April 2021

Please cite this article as: Moscato, P., Sun, H., Haque, M.N., Analytic Continued Fractions for Regression: A Memetic Algorithm Approach, *Expert Systems with Applications* (2021), doi: <https://doi.org/10.1016/j.eswa.2021.115018>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2021 Published by Elsevier Ltd.

Analytic Continued Fractions for Regression: A Memetic Algorithm Approach*

Pablo Moscato*

*College of Engineering, Science and Environment, The University of Newcastle, University
Drive, Callaghan, NSW 2308, Australia*

Haoyuan Sun

*California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, United
States*

Mohammad Nazmul Haque

*College of Engineering, Science and Environment, The University of Newcastle, University
Drive, Callaghan, NSW 2308, Australia*

*This work was supported by The University of Newcastle, Caltech SURF, the Maitland Cancer Appeal and the Australian Government through the Australian Research Council's Discovery Projects funding scheme (project DP200102364).

*Corresponding author. Postal Address: Room: ICT 316, ICT Building, The University of Newcastle, University Drive, Callaghan, NSW 2308, Australia. Phone: +61 (02) 492 16056

Email addresses: Pablo.Moscato@newcastle.edu.au (Pablo Moscato), hsun2@caltech.edu (Haoyuan Sun), Mohammad.Haque@newcastle.edu.au (Mohammad Nazmul Haque)

URL: <https://www.newcastle.edu.au/profile/pablo-moscato> (Pablo Moscato), <https://www.newcastle.edu.au/profile/mohammad-haque> (Mohammad Nazmul Haque)

Abstract

We present an approach for regression problems that employs analytic continued fractions as a novel representation. Comparative computational results using a memetic algorithm are reported in this work. Our experiments included fifteen other different machine learning approaches including five genetic programming methods for symbolic regression and ten machine learning methods. The comparison on training and test generalization was performed using 94 datasets of the Penn State Machine Learning Benchmark. The statistical tests showed that the generalization results using analytic continued fractions provide a powerful and interesting new alternative in the quest for compact and interpretable mathematical models for artificial intelligence.

Keywords: Symbolic Regression, Memetic Algorithm, Analytic Continued Fractions.

1. Introduction

Symbolic regression is a unique type of multivariate regression analysis in which the goal is to find the mathematical expression of an unknown target function that would fit a dataset $S = \{(\mathbf{x}^{(i)}, y^{(i)})\}$, i.e. a set of pairs of an unknown multivariate target function $f: \mathbb{R}^n \rightarrow \mathbb{R}$. It has been argued that when analysing experimental data for decision making symbolic regression methods should at least be used to complement standard multivariate analysis (Duffy & Engle-Warnick, 2002). Compared with the output of artificial neural network approaches, when the models generated by symbolic regressions are relatively smaller, they are perhaps more amenable to downstream studies via uncertainty propagation and sensitivity analysis and thus more “explainable” (Otte, 2013; Sun et al., 2019). The other benefit of symbolic regression is the lack of assumptions on prior knowledge on the underlying process or mechanism which produced the observed data (Moscato & de Vries, 2019b). This allows researchers to explore problem domains for which they have incomplete knowl-

edge and identifying underlying trends and patterns without subjecting human bias.

Over the past three decades, symbolic regression had produced an impressive number of results in many applications¹. For instance, symbolic regression has helped to extract physical laws using experimental data of chaotic dynamical systems without any knowledge of Newtonian mechanics (Schmidt & Lipson, 2009), which then motivated the data-driven discovery of hidden relationships in astronomy (Graham et al., 2013). More recent applications include prediction of friction systems performance (Kronberger et al., 2018), identification of nonlinear relationships in fMRI data (Märtens et al., 2017), radiotherapy dose reconstruction of childhood cancer survivors (Virgolin et al., 2018), also in the oncology field our own work on uncovering mechanisms of drug response in cancer cell lines using genomic and experimental data (Fitzsimmons & Moscato, 2018), predicting wind farm output from weather data (Vladislavleva et al., 2013), energy consumption forecasting (Delgado et al., 2018), computer game scene generation (Frade et al., 2009), Boolean classification (Muruzábal et al., 2000). They have also played a role in the elicitation of functional constructs from surveys (de Vries et al., 2014) and in the analysis of consumer and business data (Moscato & de Vries, 2019a).

One common approach to implementing symbolic regression is Evolutionary Computation (EC). EC is a family of optimization algorithms inspired by biological evolution, in particular, building upon Darwin’s theory of natural selection. In EC, a population of candidate solutions (of a problem, generally posed as an optimization one) is subject to a set of heuristics and exact algorithms to produce new solutions, while less desirable solutions are being removed from the population currently under consideration.

EC approaches to symbolic regression are commonly based on Genetic Programming (GP) with a tree-based representation. Karaboga et al. (2012) pro-

¹A ongoing updated online source of information about the field of genetic programming at: www.genetic-programming.org

posed Artificial Bee Colony Programming which also used the tree-based representation method for symbolic regression and the method showed competitive performance against GP-based methods. Each solution (a.k.a a mathematical expression) is written as a syntax tree and new solutions are produced by exchanging subtrees of two solutions (crossover) or modifying a syntax element, such as a binary operator (mutation) (Koza, 1990). Although highly popular, several researchers noted that recombination methods based on sub-tree crossovers have shown not to be better than some simple mutation of the sub-branches (Clegg et al., 2007). Clegg et al. (2007) cite previous contributions to this issue by Angeline (1997) and Luke & Spector (1997), in which they stated that “due to findings like these, some people now implement their GP’s without using crossover at all, i.e. using mutation only.”

We believe that this difficulty in symbolic regression could be addressed with generic problem-domain information about function approximation to search for better models. Like GP, Memetic Algorithms (MAs) are generic denomination for a population-based approach to solve optimization problems. However, MAs take explicit advantage of heuristic and exact methods in which solutions are individually optimized and also recombined and changed to improve the diversity of the population (Neri et al., 2012). First started at the California Institute of Technology three decades ago (Moscato, 1989, 2012), research in MAs has demonstrated over the past three decades that problem-domain information can be used to produce local search (LS) methods that can significantly accelerate the evolutionary process. Trujillo et al. (2017) recognize this fact and they point that, in contrast, local search has been underused in Genetic Programming. In their view, some of the problems faced by Genetic Programming are linked to the use of a tree-based representation of solutions. They conclude “*that numerical LS and memetic search is seldom integrated in most GP systems*” and that “*The fact that memetic approaches have not been fully explored in GP literature, opens up several areas (for) future lines of inquiry*”. We agree with this statement, in fact, of the 3918 publications we found about Memetic Algorithms (MAs) on the bibliographic database *Web of Science* (on 20/11/2019),

75 we have identified very few regarding the use of local search for symbolic regression. However, it is also true that some researchers have been trying to address the need of including individual optimization to existing Genetic Programming approaches, e.g. Cagnoni et al. (2005); Azad & Ryan (2014); Ffrancon & Schoenauer (2015); Semenkina & Semenkin (2015); Kommenda et al. (2020). While
 80 this list is probably not comprehensive, it is recognized that introducing individual optimization steps into EA methods based on current representations for solutions has been a challenge for symbolic regression approaches.

In this paper, we introduce a new approach to regression with a memetic algorithm and we analyze its performance against other existing implementa-
 85 tions of symbolic regression and machine learning approaches. In particular, our contributions are as follows:

- We introduce a novel method to represent mathematical expressions with analytic continued fractions by drawing inspirations from Padé approximants. We discuss the advantages of this particular representation over
 90 the more traditional syntax-tree based representation.
- We implement a MA for symbolic regression with the continued fraction representation, a hierarchical population structure to manage the quality of the population of solutions, and an individual search method based on the Nelder-Mead algorithm.
- 95 • We compare our MA-based approach with 15 other state-of-the-art implementations of symbolic regression with 94 benchmark data sets. We demonstrate that our algorithm is able to extrapolate well-fitting relationships and its performance is comparable to other methods.

Following this introduction the remainder of this paper is organized as follows: the datasets and methods used in Section 2, in particular, the memetic al-
 100 gorithm is described in Section 2.2; we then present an illustrative one-dimensional example of using symbolic regression to approximate an important special function in mathematics, the Gamma Function in Section 2.3. The computational

results are presented in Section 3 followed by their discussion in Section 4. Finally, Section 5 contains concluding remarks and discusses the possibility of future work in the area.

2. Data and Methods

In this section, we describe the proposed methods and datasets used to estimate its performance. We describe in detail the proposed symbolic regression method’s representation, followed by the memetic algorithm for model identification, and an illustrative example of the proposed method in approximating the gamma function. Next, the section will describe the experimental procedures and datasets used to measure the performance of the proposed method.

2.1. On representations and guiding functions

The paradigm of genetic programming (GP) (Koza, 1992) has been highly influential for the development of symbolic regression methods (e.g. P. Searson et al. (2010)). In most cases, these GP approaches work by maintaining a population of models that fit the data on the samples given as a training set, with “mutations” (i.e. small structural changes in the structures that code for these models) randomly changing the population. In addition, some “recombination” methods are used; generally these are simple heuristics involving the structures exchanging significant parts of them to create new individuals. These processes generate new structures and, consequently, new models and a “guiding function” is used to somehow bias the search for better models, by encouraging the better models to be used more frequently in future iteration of the evolutionary process. In the area of symbolic regression, the “guiding function” is often some sort of merit function, e.g. the *Mean Squared Error* (MSE) or the *Mean Absolute Difference* (MAD). Using a plethora of different techniques, in general some part of the population is eliminated and some of the new models obtained via mutation and recombination may remain in the population (while less performing models according to the guiding function are deleted). Some

other approaches to select the new population of solutions involve other considerations including, but not limited to, the use of metrics that quantify the population diversity to avoid a premature convergence of the population to the same type of model or very similar variants on the same theme (Szubert et al., 2016).

2.1.1. Representation of models by parsing trees

The syntax tree representation is simple and intuitive, and frequently used in evolutionary computation. Due to the influence of tree-based techniques for other problem domains, symbolic regression inherited this representation in GP. Mathematical expressions were created by parsing tree structures and the evolutionary operators of mutation and recombination were defined on them. In a syntax tree representation, each node represents a mathematical operator, a variable, or a constant; and for each node representing an operator, its arguments are given as the children of this node. This approach has become very popular and the open-source GPTIPS (P. Searson et al., 2010) and the commercial software Eureqa (Schmidt & Lipson, 2009, 2018) are just two good examples of the many implementations that use this technique.

These implementations generally required user-defined “building blocks”, the operands or functions that normally include the four arithmetic functions (i.e. addition, subtraction, multiplication and division), but other mathematical functions like logarithmic, exponentiation, trigonometric functions, etc., are sometimes needed as per implicit requirement of the problem characteristics. This breaks the assumption that we are not using problem domain information (i.e. we may be biasing the GP to bring solutions that involve some sinusoidal component for model building, by including the function $\sin()$, while it may not be required at all). In addition, a representation based on parsing syntax trees, even with only the four basic arithmetic functions, could be handicapped to uncover some particular types of models that frequently are the best explanation in many real world problems. We refer to those models that can be defined as the ratio of two polynomials on the variables. In our experience, we noticed that

some GP approaches do not produce those models. There may be several reasons to this, not generally related to the representation they use. For instance, in systems which also search for models of low complexity, like the commercial
 165 package Eureqa, the selection of a higher complexity for the division operator may bias the search to avoid models that employ it in equations. This may indicate that, after the population has evolved for some generations, most of the existing polynomial models are fitting the data relatively well, however, in general, it is unlikely that from those, via the use of division, we can produce a
 170 ratio of two polynomials that also fit the data well unless a powerful local search mechanism is added.

The subtle way by which a powerful representation can create an implicit bias is generally not discussed in the literature. Correctness and completeness, instead of evolvability, seems to be the main concern, and functions defined
 175 as the ratio of higher order polynomials in subsets of the set of domain variables are less likely to be generated during the evolutionary computation run. Consequently, since both recombination and mutation fail to create new models that are competitive to remain in the population, we reach a stagnation of the search process. As models become increasingly similar, we reach premature
 180 convergence, which limits the performance of GP-based approaches in many symbolic regression applications.

2.1.2. Representation via multivariate Padé approximants

One possible way to address this issue would be to find a representation that restricts all models to be rational functions (i.e. ratios of two polynomials). We
 185 first recall the following important definitions. An *analytic* function is defined as any function that can be written as a convergent power series in a neighborhood of each point in its domain D . A *holomorphic* function is a complex-valued function of one or more complex variables that is, at every point of its domain, complex differentiable in a neighborhood of the point. A *meromorphic* function
 190 on an open subset D is a function that is holomorphic on all of D except for a discrete set of isolated points, called the *poles* of the function. Every mero-

Figure 1: Approximation of $\sin(x)$ using the Padé approximant of (1) and Maclaurin series expansion of $\sin(x)$ for orders $=\{3, 5, 7, 9\}$. The values in x -axis are shown in $[-7, 7]$ and y -axis are shown within the range of $[-2.0, 2.0]$.

meromorphic function on D can be expressed as the ratio between two holomorphic functions defined on D (with the denominator not being the constant 0), and any pole must coincide with a zero of the denominator.

195 We also know that every meromorphic function has best approximating rational functions known as the *Padé approximants* (Baker Jr, 2012). This is very interesting since it is known that, numerically, Padé approximants are generally a better approximation to a function than truncating its Taylor series at the same order, and it may still be of use where the Taylor series does not converge.
200 For example, the Padé approximant of $\sin(x)$ at order $[5/6]$ is given by:

$$\sin(x) \approx \frac{\frac{12671}{4363920}x^5 - \frac{2363}{18183}x^3 + x}{\frac{121}{16662240}x^6 + \frac{601}{872784}x^4 + \frac{445}{12122}x^2 + 1}, \quad (1)$$

which provides a close fit in the interval $(-5, 5)$ of the real numbers (see Fig. 1). A Maclaurin series expansion (i.e. a Taylor series centered at zero) would need to be a polynomial of at least order 13 to approximate $\sin(x)$ to be at least competitive with expression given in (1), a polynomial of order 5 in the numerator
205 and one of order 6 in the denominator.

For a problem involving symbolic regression in a multi-dimensional domain D , a possible representation could be one based on *multivariate Padé approx-*

imants (Chaffy, 1986; Zhou et al., 2009; Akal & Lukashov, 2018). These representations may be useful since they may allow us to just use a basic mathematical form (ratio of multivariate polynomials). However, searching for a best fitting Padé approximant would require us to guess a priori the degree of two polynomials (in the numerator and denominator); we offer a novel approach to circumvent this problem. Most remarkably, this new method for machine learning and artificial intelligence has roots in eighteen-century mathematics as we will show in the following section.

2.1.3. Representation via analytic continued fractions

We shall employ a representation based on regular C-fractions (Backeljauw & Cuyt, 2009), which are the corresponding continued fractions of a Padé approximant (Lorentzen, 2010). To motivate this representation, we first note that the Padé approximant for $\sin(x)$ given by (1) can be written as follows:

$$\sin(x) \approx x - \frac{\frac{1}{6}x^3}{1 + \frac{\frac{1}{20}x^2}{1 - \frac{\frac{11}{420}x^2}{1 + \frac{\frac{25}{2772}x^2}{1 - \frac{\frac{11}{900}x^2}{1 + \frac{\frac{1331}{82650}x^2}{1}}}}}}. \quad (2)$$

Some readers may remember a result by Euler (1748) who first derived an equation that connects a finite sum of products with a finite continued fraction, as in (3) (Euler, 1748). Let $SP = a_0 + a_0a_1 + a_0a_1a_2 + \cdots + a_0a_1a_2 \cdots a_n$, then:

$$SP = \frac{a_0}{1 - \frac{a_1}{1 + a_1 - \frac{a_2}{\ddots \frac{a_{n-1}}{1 + a_{n-1} - \frac{a_n}{1 + a_n}}}}}. \quad (3)$$

225 Since the result can be proved by induction on n , this means that we can apply this result in the limit.

This fact takes directly to our proposed representation (e.g. the power of this representation for special functions can be found in Backeljauw & Cuyt (2009)). If, for instance, the left-hand side of the equation is a convergent
 230 and infinite series, then the right-hand side represents a *convergent and infinite continued fraction*. As an example, we consider the function $\tanh(x_1 + 5x_2) = \sinh(x_1 + 5x_2)/\cosh(x_1 + 5x_2)$. Unless $\sinh(x_1 + 5x_2)$ and $\cosh(x_1 + 5x_2)$ are given as “building blocks”, it would be very difficult for a system to evolve the ratio. However, Gauss proved in 1812 (see Olds (1963, p. 138), Wall (1948,
 235 p. 349)) that $\tanh(x_1 + 5x_2)$ can be well written as:

$$\tanh(x_1 + 5x_2) = \alpha_0 + \frac{f_0(\mathbf{x})}{\alpha_1 + \frac{f_1(\mathbf{x})}{\alpha_2 + \frac{f_2(\mathbf{x})}{\alpha_3 + \ddots}}} \quad (4)$$

with $\alpha_0 = 1$, $f_0(\mathbf{x}) = x_1 + 5x_2$, $f_i(\mathbf{x}) = (x_1 + 5x_2)^2$, for any natural number $i > 0$, and with the set $\{\alpha_i\}$ with $i > 0$ being the odd natural numbers in increasing values (i.e. $\alpha_1 = 1, \alpha_2 = 3, \alpha_3 = 5, \alpha_4 = 7, \alpha_5 = 9, \alpha_6 = 11, \dots$). Numerically, just a few levels would give a good approximation to the original
 240 function.

This shows that the representation naturally contains the Padé approximant as a finite sum in terms of the variables x and x^2 . Also, our representation allows
 255 a good truncated approximation to $\tanh()$ (e.g. by using Lambert’s continued fractions) even if the functions $\sinh()$ and $\cosh()$ have not been provided as “building blocks”. Since only the four fundamental arithmetic operations are used, the continued fraction representation is able to approximate a function to increasing precision with more “depth” and in the limit it will converge to
 260 the target function of interest. We will formally define the notion of depth in Section 2.1.4.

Analytic continued fractions are thus giving us a powerful representation; not only we can represent elementary functions (like $\sin()$, $\log()$, $\cos()$, $\tan^{-1}()$, etc.), but other special functions (like the Error function $\operatorname{erfc}()$, and the Incomplete Gamma Function, etc.) can also be represented in this way (Crandall,
 265 1994). We also note that there exists other possible representations (e.g. Stieltjes fractions, Thron fractions or Generalized T-Fractions, Thiele interpolating continued fractions, and Jacobi or J-Fractions) (see Backeljauw & Cuyt (2009) for their definitions). While exploring these alternative representations is an
 270 interesting research avenue with merits of its own, we concentrate our study in the one proposed here as a first benchmark on its performance for symbolic regression on a large dataset.

The use of analytic continued fractions would certainly power many other aspects of evolutionary computation due to its well-established mathematical
 275 properties and theoretical foundations. It has also the added advantage that it liberates from the need of sometimes “guessing building blocks” as continued fractions can represent an infinite number of functions that can be represented by an infinite but convergent series. Indeed, in addition to symbolic regression, we feel that analytic continued fractions will be a welcomed representation for
 280 the evolutionary computation community for many modelling approaches.

2.1.4. *Convergents and the depth of the representation*

While our representation is defined as an infinite sum, in practice, we expect that the sum should be truncated in a computational setting. Some solutions include: either we leave the truncation as another variable (since it directly
 285 relates to the “complexity” of the model that fits the data in question), or we leave it variable and somehow penalize those models that fit the data well but the truncation has a high value (e.g. for (9) that value is 6, and we will say that the “depth” is equal to 6). Our formal definition of “depth” is associated with the convergent of a continued fraction. This links with known results
 290 from analytic continue fractions since the convergents of regular C-fractions are exactly the Padé approximants (Lorentzen, 2010), so the *depth of n* for a model is defined to be the n th convergent of our representation (i.e. a model of depth of n also contains a model of depth n' for all $0 \leq n' < n$). This said, if we decide to have a representation with a fixed “depth” of six (or the third convergent),
 295 this means that in (5), we have $h_6(\mathbf{x}) = 0$ but $g_5(\mathbf{x})$ is not constrained to be equal to zero. Clearly, the selection of a higher order (and fixed) depth would guarantee a more precise fit, but may lead to overfitting on the training set and poor generalization in the test sets.

Our representation is then quite general and, in principle, may be adopted
 300 for any arbitrarily defined maximum depth. For some problem domains, a multi-objective approach could be employed (e.g. following similar experiences in GP-based symbolic regression (Vladislavleva et al., 2009)), in which minimization of depth is one of the objectives since it directly contributes to model complexity. We note that the use of model complexity as a second objective is common
 305 in GP and in other domains (e.g. to control so-called the undesired effect of “bloating” in the solutions obtained). Since solving the high-dimensional non-linear optimization problem we face, we need to employ a powerful and robust optimization solving technique, one that enforces both individual and global non-linear optimization. Towards this end, we present our memetic approach
 310 in the following section.

2.2. A memetic algorithm for model identification

In this section, we present a memetic algorithm (Neri et al., 2012) to search for the best coefficients for (8), and deliver a best model for our training data. We call this regression method using this memetic algorithm and continued
 315 fraction representation as ‘Continued Fraction Regression’ (CFR). Our selection and design of this optimization technique in part stems from our familiarity with the technique (Moscato, 1989) and the robustness shown in thousands of different applications, including many in the area of non-linear optimization. We also aim here to show how it can be easily implemented from the combined
 320 use of non-linear optimization solvers and a few primitives that organize the search via interaction of a set of agents.

The “guiding function” of a solution is its mean squared error (MSE) and the goal of the population is to find the model that minimizes the MSE. Clearly, others can be used as well. However, the comparison of different guiding func-
 325 tions is outside of the scope of this paper.

The advantage of the memetic algorithm is that it combines problem domain information (i.e. the use of CFR, based on established mathematical theory for function approximation), together with local optimizers which are robust and global search mechanisms to search for good subsets of variables to create models
 330 with.

Unlike the GP’s tree-based representation, we have a fixed representation. We do not search in a space of trees but in the high-dimensional space defined by all the coefficients in the formula of a CFR. In some sense, we are solving two optimization problems, one nested in the other one: (1) identify which are
 335 the variables that are needed in a model, and (2), given those variables what are the values of the coefficients in (8).

The decision of including or not a variable in turn decides if its associated coefficients should be optimized or not. This is handled by the recombination and mutation parts of the algorithm. A population of models is then maintained, and we refer as a “generation” the period of the MA evolutionary process in
 340 which we perform the operations of mutation, recombination and individual

search optimization to find the best coefficients. We employ a very simple recombination approach needed for variable selection, but our recombination could also be considered “memetic” (Eskridge & Hougen, 2004). We use a
 345 direct search method for individual model optimization and it is described in Sec. 2.2.3. Coefficients are optimized by the use of a variant of a Nelder-Mead algorithm (which provides a kind of individual search mechanism).

One key aspect we maintain from previous successful implementations of memetic algorithms (Buriol et al., 2004) is the use a population structure, which
 350 is explained in Sec. 2.2.1.

We then start with the discussion of the individual model optimization and the population structure which may be the most uncommon feature of an MA for some of our readers. To improve the readability of the paper, we have shown the high-level view of the memetic framework for the symbolic regression in
 355 Algorithm 1 and it summarizes the description of the algorithm presented in later sections.

2.2.1. A memetic algorithm with a hierarchical population structure

Tree-structured MAs. In fact, a tree-based population structure of agents was initially proposed in the early 90s (Moscato & Tinetti, 1994) and it has been
 360 subsequently used for the solution of many combinatorial optimization problems such as the asymmetric traveling salesman problem (Moscato & Tinetti, 1994; Buriol et al., 2004), hierarchical clustering and phylogenetics (Cotta & Moscato, 2002, 2003), multistage capacitated lot-sizing problem (Berretta & Rodrigues, 2004), gene expression linear ordering microarray data (Cotta et al., 2003b),
 365 Lot Sizing and Scheduling (Toledo et al., 2012), total tardiness single machine scheduling (Mendes et al., 2002a), Quadratic Assignment Problem (Harris et al., 2015), Gate Matrix Layout problem (Mendes et al., 2002b), 3D Protein Structure Prediction (de Lima Correa et al., 2018) to mention just a few very challenging applications. Other problem domains include the number partitioning
 370 problem (Berretta et al., 2004), and many problems in bioinformatics and visualization (Cotta et al., 2003a; Inostroza-Ponta et al., 2011) dealt with by our

Algorithm 1: CFR Algorithm

Input : Num. of Vars $nVars$, Mutation Rate μ_r , training data D_{trn}
Output: Best solution to fit the problem, $best$

```

/* Generate Initial Population by a randomized algorithm */
1  $pop \leftarrow \text{InitialPopulation}(nVars, D_{trn})$ 
2  $best \leftarrow \text{pocket}(pop.root)$ 
3 for  $gen \in 1 \rightarrow numGen$  do
    /* Mutate each current solution in the population */
4      $pop_\mu \leftarrow \text{Mutate}(pop, \mu_r)$ 
    /* Generate new Population by recombination mechanism */
5      $pop_r \leftarrow \text{RecombinePopulation}(pop_\mu)$ 
    /* Local Search Optimization of current solutions */
6     foreach  $agent \in pop_r$  do
7         |  $\text{LocalSearch}(agent)$ 
8     end
9     if guiding function stagnates for consecutive 5 gens then
10        |  $\text{reset}(pop_r.root)$ 
11    end
    /* Replace old population with evolved population */
12     $pop \leftarrow pop_r$ 
    /* Keep track of the best solution */
13    if  $\text{guiding\_function}(best) < \text{guiding\_function}(\text{pocket}(pop.root))$ 
        then
14        |  $best \leftarrow \text{pocket}(pop.root)$ 
15    end
16 end
17 return  $best$ 

```

group over many years (in which tree-based memetic algorithms were used to solve particular sub-problems of interest). The number of publications that have used a ternary-tree topology exceeds 40 (according to Google Scholar) and it has been employed during the research work of several PhD Theses in different countries being an established alternative to other topologies in use. In Buriol et al. (2004), a comprehensive study of 40 different tree topologies was conducted with population sizes varying from 85 to 7 agents (each having two solutions being considered). Many of them included binary trees of different lengths. One measure of merit for these experiments was the “gap” to the known optimal

Figure 2: The structure of a population consisting of 13 agents. Each of the agent contains two solutions: named ‘pocket’ and ‘current’.

solutions on the set of instances, and the CPU time employed. Interestingly, from this analysis, the ternary tree topology was shown to have the best quality trade-off between solution quality and time, and we have subsequently adopted this topology in several other applications.

385 *Agents.* The basic unit of population in our system is called an *agent*. Each one has two possible different solutions (models), which are referred as the *pocket* and the *current* (Buriol et al., 2004; Mendes et al., 2002b; Harris et al., 2015). Our population contains 13 agents and they are organized as a depth-3 ternary tree (Fig. 2). This approach has enabled in many cases good performances with
390 a small population size of 13 agents (thus having a total of 26 solutions).

Initial Population. The initial population is decided by a randomized algorithm that works as follows. For each agent in the population we need to initialize two solutions. Each one is done the same and it is as follows. We select with equal probability the inclusion of a variable in a model, and when this is done we select
395 at random with uniform probability a coefficient for each of the occurrences of a variable in a model (i.e. if we are working with **depth** equal to 8 it will be 17 coefficients to set to a value selected uniformly at random between -3 and 3 for each variable occurrence in the model). We also set in the same way any α_i and β_i values in from (6) and (7).

400 Premature convergence in evolutionary algorithms needs to be addressed by introducing some techniques. Some of them are numerical, adaptive algorithms and heuristics that aim to control for loss of diversity. One that is non-numerical

in nature is “isolation by distance”, the introduction which are rules that prevent certain solutions to be recombined. One way of doing this is via the use of a
 405 “population structure”, thus by restricting the recombination of solutions.

Algorithm 2 shows the process of generating an initial population. Here, we generate 13 agents. We generate two solutions per agent (one ‘pocket’ and one ‘current’ solution), where each solution consists of a set of variables taken uniformly at random. Associated coefficients of the variables are also generated
 410 uniformly at random in the range of -3 to 3 . Then we evaluate the guiding function score of both solutions. We place the ‘pocket’ and ‘current’ in such a manner so that it maintains the invariants of pocket solution having better guiding function value than the one of the current solution.

Invariants. Our algorithm maintains the following invariants in the population
 415 structure. First, within each agent, the *pocket* solution always has a better value of the guiding function value than its *current* solution. In addition, the pocket solution of each non-leaf node is always better or equal in quality to the three in the level immediately below to it. Maintaining these two invariants guarantees that there is a flow of good models upwards in the tree structure. A
 420 tree-based population structure avoids in part premature convergence (Moscato et al., 2007).

2.2.2. Recombination and mutation to identify sets of variables for the models

The evolutionary operators of mutation and recombination are responsible for the selection of the variables that will have nonzero coefficients and be in-
 425 cluded in a model. For example, the following function

$$f(w, x, y, z) = 2.1 w + \frac{4.7 x + w + 1.01}{x + \frac{1.3 + 5.7 y}{3.9 x}} \quad (10)$$

only has the variable z missing, so it uses three.

Algorithm 2: InitialPopulation**Input** : Number of Variables $nVars$, training data D_{trn} **Output:** Randomly initialized Population pop

```

/* Instantiate 13 Agents in the population */
1  $pop \leftarrow Agent(13)$ 
2 foreach  $agent \in pop$  do
    /* Randomly generate two solutions per agent */
3      $current(agent) \leftarrow ContinuedFraction(nVars)$ 
4      $pocket(agent) \leftarrow ContinuedFraction(nVars)$ 
    /* evaluate and save guiding function of both solutions */
5      $guiding\_function(current(agent)) \leftarrow evalGF(current(agent), D_{trn})$ 
6      $guiding\_function(pocket(agent)) \leftarrow evalGF(pocket(agent), D_{trn})$ 
    /* within each agent, the pocket solution always holds
       better value of guiding function than its current
       solution */
7      $MaintainInvariant(agent)$ 
8 end
9 return  $pop$ 

```

The recombination operator takes two solutions and generates an offspring by combining the set of variables present in two parent individuals. At each step, one of the three possible recombination operators is chosen uniformly at random. We choose one recombination by considering the *union* ($\cup$), *intersection* ($\cap$), and *symmetric difference* (Δ) between the set of variables present in the models to be recombined.

The recombination operator then acts on each subpopulation, from top to bottom. Within each subpopulation, the operation is performed on each agent as follows (leader labelled as ℓ , and supporters are labelled s_1, s_2 and s_3 from left to right):

1. $current(\ell) \leftarrow Recombine(pocket(\ell), current(s_1))$
2. $current(s_3) \leftarrow Recombine(pocket(s_3), current(\ell))$
3. $current(s_1) \leftarrow Recombine(pocket(s_1), current(s_2))$
4. $current(s_2) \leftarrow Recombine(pocket(s_2), current(s_3))$

The `Recombine()` process of two agents to create an offspring is shown in

Algorithm 3. Here, we select the variables in the offspring by applying the operator chosen uniformly at random on the variables of both parents. For each depth of the continued fraction, we compute the associated coefficient value
 445 for each variable. **In the case of the new variable being selected** only from a single parent, we copy the parent’s associated coefficient into the offspring solution. However, if this variable is already **a part of both parents** (a and b), we compute the new **value of the coefficient** for corresponding variables as: $cof_a + rand(-1, 4) * (cof_b - cof_a)/3$ (historically, this randomized approach has been
 450 previously used in early Scatter Search methods for non-linear optimization). We use the same formula to compute the value of the constant (β) portion of the continued fraction. Then the guiding function is reevaluated. The new model then will contain a variable with nonzero coefficient only if that variable belongs to the set being preserved and at least one of the parents contains the
 455 same variable in the model. The individual search step is then the one in charge of optimizing the non-zero coefficients of the newly created model.

Like in other evolutionary computation algorithms, mutation is a random mechanism used to increase the diversity of the population. Here we have decided to incorporate two forms of mutation operators (‘major mutation’ and
 460 ‘soft mutation’) in our implementation to toggle random variables. We choose a mutation operation depending on the guiding function value of current and pocket solutions. We do a ‘major mutation’ on the current *if the guiding function value is either within 120% of the guiding function value of pocket* ($current.guiding_function < 1.2 * pocket.guiding_function$) *or greater than*
 465 *twice of the pocket’s guiding function value* ($current.guiding_function > 2 * pocket.guiding_function$). A ‘soft mutation’ is taken place in other cases.

For the ‘major mutation’ (detail of this mutation operation, `toggleVariables()`, is shown in Algorithm 4), we select a variable uniformly at random for toggle and modify associated coefficients in all depths of the continued fraction. If the
 470 randomly selected variable was already incorporated (“switched on”) for the depth, we remove the variable. However, either we “remember” the existing co-

Algorithm 3: Recombine

Input : Two agents take part in recombination a, b , choice of an operator uniformly at random from $\{\cup, \cap, \Delta\}$ as OP

Output: Offspring as the Current solution of agent a

```

/* Initialize parents ( $p_1, p_2$ ) and offspring  $ch$  */
1  $p_1 \leftarrow pocket(a), p_2 \leftarrow current(b)$ 
2  $ch \leftarrow \phi$ 

/* Apply a recombination operator chosen uniformly at random
   on variables of two parents into offspring */
3  $ch.vars() \leftarrow p_1.vars() OP p_2.vars()$ 

/* Recombine the coefficients for each term ( $h$ ) of the
   continued fraction */
4 for  $i = 0$  to  $(2 * depth)$  do
5    $coef_a \leftarrow p_1.h_i.coef(), vars_a \leftarrow p_1.h_i.vars()$ 
6    $coef_b \leftarrow p_2.h_i.coef(), vars_b \leftarrow p_2.h_i.vars()$ 

   /* recombine coefficient values for variables */
7    $coef_\chi \leftarrow \phi$ 
8   for  $vi = 1$  to  $nVars$  do
9     if  $ch.feasAt(vi) = true$  then
10      if  $vars_a[vi] = true$  and  $vars_b[vi] = true$  then
11         $coef_\chi[vi] \leftarrow coef_a[vi] + rand(-1, 4) * (coef_b[vi] - coef_a[vi]) / 3$ 
12      else if  $vars_a[vi] = true$  then
13         $coef_\chi[vi] \leftarrow coef_a[vi]$ 
14      else if  $vars_b[vi] = true$  then
15         $coef_\chi[vi] \leftarrow coef_b[vi]$ 
16      end
17    end
18  end

  /* Update new coefficients of the term in offspring */
19   $ch.h_i.coef() \leftarrow coef_\chi$ 
  /* compute new value of constant ( $\beta$ ) for term  $h_i$  in the
     offspring solution  $ch$  using  $\beta$  of  $p_1.h_i$  and  $p_2.h_i$  */
20   $ch.h_i.\beta \leftarrow p_1.h_i.\beta + rand(-1, 4) * (p_2.h_i.\beta - p_1.h_i.\beta) / 3$ 
21 end

/* Update current solution and apply Local Search */
22  $current(a) \leftarrow ch$ 
23  $a \leftarrow localSearch(a)$ 
24 return  $a$ 

```

efficient value or assign ‘zero’ to “remove” with probability 50%. Alternatively, if the variable was not incorporated, we switch it on and replace the coefficient either by ‘zero’ or by a uniformly random number in the range of -3 to 3 with the probability of 50%. In the case of ‘soft mutation’ (shown in Algorithm 5 as `modifyVariable()`), we select a depth and a variable, both uniformly at random. If the variable was already incorporated, toggle the variable selection and ‘remove’ the coefficient by assigning zero as value. Otherwise, we ‘modify’ the coefficient of the random variable by a uniformly generated random number in the range of -3 to 3 and toggle the variable selection. Finally, the local search operation is executed on the mutated solution in order to optimize non-zero coefficients. We do not apply mutation on pocket solutions because we consider them as a “collective memory” of good models visited in the past. They influence the search process via recombination only.

2.2.3. Individual model optimization via a direct search method

A period of individual search operation is performed every generation on all current solutions. We remind again that each solution corresponds to a single model, this means that if a current model becomes better than its corresponding pocket model (in terms of the guiding function of the solution), then an individual search optimization step is also performed on the pocket solution (or model) before we swap it with the current solution (or model). Individual search can then make a current model better than the pocket model (again, according to the guiding function), and in that case they switch positions within the agent that contains both of them.

To do these optimizations, we used a modified version of Nelder-Mead algorithm recently proposed by Fajfar et al. (2016) to optimize the coefficients of a model. Nelder-Mead methods, also known as the *downhill simplex algorithm*, is a derivative-free nonlinear optimization algorithm known for its simplicity and relatively good empirical performance (Singer & Nelder, 2009).

To run the Nelder-Mead algorithm, the list of constants and coefficients in a model is mapped to a vector by the order in which they appear. The initial

Algorithm 4: toggleVariables

Input : A Continued Fraction Solution $cfrac$
Output: The modified solution $cfrac$

```

/* select a variable index uniformly at random */
1  $N \leftarrow cfrac.nVars$ 
2  $vIdx \leftarrow \text{randChoice}(N)$ 

/* for each depth of continued fraction, toggle the
   selection of variables of the term ( $h$ ) */
3 foreach  $h \in cfrac$  do
    /* Case 1: cfrac variable is turned ON: Turn OFF the
       variable, and either 'Remove' or 'Remember' the
       coefficient value at random */
    4 if  $cfrac.varAt(vIdx) = \text{true}$  then
    5      $h.varAt(vIdx) \leftarrow \text{false}$ 
    6      $h.coefAt(vIdx) \leftarrow \text{coinToss}(0, h.coefAt(vIdx))$ 
    7 else
        /* Case 2: term variable is turned OFF: Turn ON the
           variable, and either 'Remove' or 'Replace' the
           coefficient with a random value between -3 to 3 at
           random */
    8     if  $h.varAt(vIdx) = \text{false}$  then
    9          $h.varAt(vIdx) \leftarrow \text{true}$ 
    10         $h.coefAt(vIdx) \leftarrow \text{coinToss}(0, \text{rand}(-3\ 3))$ 
    11    end
    12 end
13 end

/* Toggle the randomly selected variable */
14  $cfrac.varAt(vIdx) \leftarrow \neg cfrac.varAt(vIdx)$ 

```

Algorithm 5: modifyVariable**Input** : A Continued Fraction Solution $cfrac$ **Output:** The modified solution $cfrac$

```

/* Randomly select a variable which is turned ON */
1  $candVars \leftarrow \{\forall i : cfrac.varAt(i) = true\}$ 
2  $vIdx \leftarrow randChoice(candVars)$ 

/* Randomly select a term ( $h$ ) of continued fraction */
3  $h \leftarrow randChoice(\{\forall term \in cfrac\})$ 

/* Modify the coefficient value */
4 if  $h.varAt(vIdx) = true$  then
5 |  $h.coefAt(vIdx) \leftarrow 0$ 
6 else
7 |  $h.coefAt(vIdx) \leftarrow rand(-3\ 3)$ 
8 end

/* Toggle the randomly selected variable */
9  $h.varAt(vIdx) \leftarrow \neg h.varAt(vIdx)$ 

```

simplex is generated by adding a unit step to each dimension of said vector. This heuristic of initializing the simplex is described in Singer & Nelder (2009). The individual search stops when the best and worst vertex of the simplex are within numerical tolerance, or when a maximum number of iterations is reached. In our implementation, we set the numerical tolerance as 10^{-3} and the maximum iterations of 250.

While our memetic algorithm does individual search optimization with a modified Nelder-Mead algorithm (Fajfar et al., 2016), the nature of the non-linear optimization problem indicates that more powerful solvers could later be used. We have chosen to start with a Nelder-Mead solver as a way of ensuring and promoting reproducibility, keeping the core memetic algorithm as simple as possible (Sun & Moscato, 2019) for this first in-depth test of performance of the new representation in real-world problems.

2.2.4. Small batch learning

To deal with datasets having more than 200 samples, during an execution of individual search, a small subset of the set of training samples (20%) is selected uniformly at random (from the whole training set) and during the

individual search optimization, the guiding function value is computed only
 520 using these samples. This selection is conducted every time an individual search
 optimization is required. For each model, 4 independent local searches are
 performed (using the variant of the Nelder-Mead algorithm discussed before);
 the result with the best guiding function value (on the entire dataset) is chosen.
 This allows to have some sort of sampling of the quality of the variables in
 525 independent trials of the individual search process.

2.2.5. Diversity Management

We have set up a time-limit on the best model currently in the population
 (i.e. the one represented by the pocket solution of the root agent of our tree
 hierarchy) to influence the search. Our criterion for relevance has been fairly
 530 strict: if no better model has been produced for five (5) straight generations,
 then the pocket of the root agent is removed and a new solution is created at
 random. This is a fairly strong requirement, but for a depth-3 ternary tree, a
 solution that has climbed to the top and has been “the best seen” for five gen-
 erations had already the opportunity to influence the search procedure and we
 535 avoid being trapped into a local minimum of the search space and not exploring
 other combinations of variables.

2.2.6. Model complexity management

In the area of symbolic regression, and in some GP implementations, the
 tendency of fitting better the data at the cost of producing more complex models
 540 is called “bloating”. This is an undesirable characteristic as one of the objectives
 of symbolic regression is to have easy to interpret, small and useful models and
 with better generalization (Dick, 2014).

We decided to set up an adaptive control for the complexity of our feasible
 solutions by penalizing the number of variables being used by a model. At ini-
 545 tialization of the set of models, a whitelist is generated uniformly at random and
 independently done for each individual. Each input variable then has a proba-
 bility $p = 1/3$ to be present in the whitelist. Only variables on the whitelist may

Table 1: Parameter value of the CFR-based memetic algorithm for regression.

Parameter	Value
Delta	0.10
Fraction's Depth	4
Reset root of population after stuck for generations	5
Number of Generations	200
Mutation Rate	0.10
Number of Nelder-Mead Instances	4
Number of Iterations in Nelder-Mead	250
Nelder-Mead terminates if stagnates for consecutive iterations of	10
Percentage of Samples used to evaluate a model in local search	20%

be assigned with a value so that no bloated solution appear in the initial population. To bias the search towards solutions of lesser complexity, we penalize the complexity.

As mentioned above, the *Mean Squared Error* (MSE) metric is used to quantify solution quality. The solution quality is the average of the squared errors of prediction (y') with the measured (or observed) value (y) of the target for all n samples $\{(\mathbf{x}^{(i)}, y^{(i)})\}$ in the dataset (S) according to:

$$MSE = \frac{1}{n} \sum_{i=1}^n (y^{(i)} - y'^{(i)})^2. \quad (11)$$

Instead of comparing the quality of the solutions by the goodness of fit alone (measured by MSE), we aim at minimizing an “adjusted MSE” given by:

$$\text{adjusted MSE} = MSE \times (1 + \Delta \times \# \text{ of variables used}) \quad (12)$$

where $\Delta > 0$ is the scale of the penalty. Different values of Δ help to achieve different balances between goodness of fit and complexity depending on the workload.

Table 1 shows the value of the parameters used for our experiments.

2.3. Learning the Gamma Function

The Gamma function $\Gamma(x)$ is one of a number of analytic extensions of the factorial function to real and complex numbers. Many common integral

(a) $Depth = 2$ (b) $Depth = 4$ (c) $Depth = 6$

Figure 3: Results of our memetic algorithm (red +) learning the $\Gamma(x)$ function with $x \in [-2.683, 4.5]$ (blue o) producing models restricted to a maximum value of Depth: (a) $Depth = 2$, (b) $Depth = 4$ and (c) $Depth = 6$ in CFR. The values in Y-axis are shown in the range of $\{-30, \dots, 30\}$. However, we have two points at $x = -1.978$ and $x = 0.072$ out of this range ($y = 62.00324026$ and $y = 139.476602$, respectively) for $Depth = 6$ (in Sub-Fig. (c)) that have been omitted for clarity. Understandably, with a deeper model we can approximate the function better, although we have observed some runs for which $Depth = 4$ obtaining results remarkably similar to those of $Depth = 6$ in this figure. This indicates that global optimization methods will have a remarkable role in using small depth analytic continued fraction representations for symbolic regression.

calculations in applied mathematics can be expressed in terms of the Gamma
 565 function. For complex numbers with a positive real part, the Gamma function
 is defined via the convergent improper integral:

$$\Gamma(z) = \int_0^{\infty} x^{z-1} e^{-x} dx. \quad (13)$$

The only complex numbers for which the $\Gamma(z)$ is not defined are the non-positive
 integers for which it has simple poles.

The Gamma function occurs in many areas of statistics and data analysis
 570 and it frequently appears in the study of natural phenomena where there exists
 a process in which an decay in time is present and this decay follows a law of the
 form $f(t)e^{-g(t)}$ where t represents time (Miller, 2015). For instance, when $f(t)$
 is a power function and $g(t)$ is a linear function of time, respectively, thanks to
 a change of variables we can write:

$$\int_0^{\infty} t^b e^{-at} dt = \frac{\Gamma(b+1)}{a^{b+1}}. \quad (14)$$

575 We have chosen to show the performance of the technique on a testbed
 problem before presenting the computational results in a large benchmarking
 dataset. This illustrates the role of the depth of the continued fraction represen-
 tation in increasing the approximation to the true values. This said, we propose
 that learning the values of $\Gamma(x)$ on an interval on the reals around the zero is an
 580 interesting challenge for symbolic regression methods and could be now added
 to the list of benchmarking functions of one variable. It is also an interesting
 test function since for non-positive integers the target function values alternate
 signs between its poles, while for positive integers there are no poles but its val-
 ues increase faster than an exponential function (since for any positive integer
 585 n , $\Gamma(n) = (n-1)!$).

As an example of how we can approximate the Gamma Function in the
 interval in which the function alternates signs between some poles, we have
 chosen the interval $[-2.683, 4.5]$ (see Fig 3). We have generated a set $S =$

$\{(x, y)\}$ of 873 samples using a uniform separation step. To simulate a situation
 590 similar to the one we will observe in the experimental setting that was used
 later in the paper (in which we compare our results with other methods having
 a high-dimensional input space), we have created an input space consists of the
 variable x and all of its powers up to degree 6. This then creates an instance
 of symbolic regression of 7 dimensions since now we have $(x) = [1, x, x^2, \dots, x^6]$.
 595 We have run our memetic algorithm for several depths of the representation,
 i.e. 2, 4 and 6. Fig. 3 shows the learning outcome of $\Gamma(x)$ for different depths.

We note that the GP based software Eureqa version 1.24.0 (Schmidt & Lipson, 2009, 2018) was able to very quickly find a solution of the form:

$$\Gamma(x) \approx ax^3 + \frac{b}{x + x^2}, \quad (15)$$

with $a \approx 0.096350$ and $b \approx 1.00808$. However, further computation using nearly
 600 one million generations (approximately 40 minutes) was required to finally deliver a better solution:

$$\Gamma(x) \approx c_1 + c_2x + \frac{c_3}{x} + c_4x^6 + \frac{c_5}{c_6x + x^3 + c_7x^2}, \quad (16)$$

with $c_1 \approx 0.201148$, $c_2 \approx 0.216747$, $c_3 \approx 0.492681$, $c_4 \approx 0.001204$, $c_5 \approx 1.004562$,
 $c_6 \approx 2.000514$, $c_7 \approx 3.000165$.

In terms of time, our method required only 87.5 seconds, a fraction of the
 605 time employed by Eureqa on the same laptop, to obtain the solution shown in
 Fig. 3(c). Since the academic version of Eureqa we have, can not be used for
 large-scale batch experimentation in our grid, we resort to a battery of tests
 recently proposed to study the performance of genetic programming and other
 machine learning methods in a systematic manner. The next section presents
 610 the benchmarking tests and results.

2.4. Experimental settings and Datasets

To compare the performance of our proposed algorithms with other state-of-the-art algorithms we have used a standard *benchmark* dataset (Olson et al.,

2017). In addition to MSE score, we will also compute the normalized MSE
 615 score (NMSE) that provides information about the deviation (loss) of the perfor-
 mances and which allows comparisons between different databases. The NMSE
 score is calculated by normalizing the MSE score with respect to the variance
 (Var) of the measured/observed value (y) (Quiñonero-Candela et al., 2006) ac-
 cording to: $\text{Var}(y) = \frac{1}{n-1} \sum_{i=1}^n |y_i - \bar{y}|^2$. The NMSE score is given by:

$$NMSE = \frac{1}{n} \sum_{i=1}^n \frac{(y_i - y'_i)^2}{\text{Var}(y)}. \quad (17)$$

620 2.4.1. The 94 Datasets from the Penn Machine Learning Benchmarks

Olson et al. (2017) compiled a collection of datasets to evaluate machine
 learning problems, known as *Penn Machine Learning Benchmarks* (PMLB)².
 Later on, Orzechowski et al. (2018) selected a subset of 94 datasets from PMLB
 which are suitable for regression problems. We have employed these 94 datasets
 625 for the evaluation of our memetic algorithm.

Since each of those 94 datasets only has a single file of samples associated to
 the regression problem, we have selected uniformly at random 75% of the entries
 as a training set and the remaining 25% as testing for the algorithm to measure
 the performance (i.e. an independent test set). To make the results more
 630 significant and free from bias, that split of train-test has been taken uniformly
 at random for each of the 100 independent executions of the memetic algorithm.

2.4.2. Computational Environment

We have compiled the algorithm with gcc compiler version 4.8.5 in 64 bit
 Red Hat Enterprise Linux Server 7.5 (Maipo). We have executed the algorithm
 635 on The University of Newcastle's High Performance Computing (HPC) grid³;
 we split the execution of the CFR on 94 datasets into 10 batches. For each
 batch, we assigned 2 CPU cores, 8GB RAM, and a total of 100 hours of CPU

²<https://github.com/EpistasisLab/penn-ml-benchmarks>

³<https://www.newcastle.edu.au/research-and-innovation/resources/research-computing-services/advanced-computing>

wall time to execute 100 independent runs of the algorithm.

3. Results

3.1. Results on Penn Machine Learning Benchmarks Datasets

We measured the MSE and NMSE scores obtained on the training set for each of the datasets. The model found during the training phase was also evaluated on the testing data. We tabulated the median scores of MSE and NMSE for both in training ($\mathfrak{t}$) and testing ($\mathfrak{T}$) scores of the algorithm for the 100 independent runs. Table 2 summarizes the scores for all benchmark datasets. We have sorted the results from the smallest to the largest of the median NMSE scores achieved in testing.

Table 2: Median value of the proposed approach (Computed Over 100 runs) for the MSE and NMSE scores on Training ($\mathfrak{t}$.) and Testing ($\mathfrak{T}$.) splits on the 94 Benchmark Datasets.

Dataset	t_{mse}	T_{mse}	t_{nmse}	T_{nmse}	Dataset	t_{mse}	T_{mse}	t_{nmse}	T_{nmse}
rabe_266	8.381	9.641	0.003	0.004	sl_ex1714	1.11e6	1.86e6	0.145	0.299
a_elec2000	4.09e6	1.42e7	0.001	0.005	fri_c3_500_10	0.238	0.278	0.244	0.307
a_neavote	0.725	0.903	0.046	0.059	fri_c2_500_50	0.270	0.310	0.284	0.307
cpu	6.37e2	1.33e3	0.028	0.077	fri_c0_250_5	0.249	0.308	0.246	0.311
fri_c2_1000_5	0.069	0.078	0.069	0.080	fri_c0_1000_50	0.275	0.313	0.279	0.314
fri_c1_1000_5	0.079	0.088	0.079	0.088	fri_c2_250_25	0.249	0.326	0.252	0.317
fri_c3_1000_5	0.081	0.094	0.080	0.096	fri_c0_500_10	0.285	0.324	0.281	0.324
v_galaxy	8.00e2	8.86e2	0.090	0.099	fri_c0_500_25	0.262	0.307	0.262	0.325
fri_c2_500_5	0.083	0.105	0.083	0.105	fri_c0_500_50	0.285	0.349	0.285	0.346
fri_c4_1000_10	0.103	0.115	0.106	0.112	fri_c1_500_50	0.333	0.355	0.342	0.347
fri_c1_500_5	0.095	0.116	0.096	0.117	fri_c1_1000_50	0.330	0.350	0.333	0.349
a_apnea2	1.17e6	1.17e6	0.114	0.129	sl_case1202	2.17e3	2.80e3	0.262	0.355
fri_c1_1000_10	0.129	0.135	0.130	0.134	fri_c0_250_10	0.298	0.376	0.301	0.373
a_apnea1	1.16e6	1.21e6	0.109	0.138	fri_c0_100_5	0.240	0.377	0.243	0.388
chatfield_4	2.47e2	2.96e2	0.123	0.141	fri_c3_250_10	0.343	0.413	0.339	0.420
fri_c3_1000_25	0.128	0.141	0.129	0.150	fri_c4_250_25	0.333	0.394	0.321	0.421
ESL	0.265	0.309	0.133	0.150	fri_c0_250_25	0.293	0.411	0.295	0.422

Continue on the next page

Table 2 – continued from previous page

Dataset	t_{mse}	T_{mse}	t_{nmse}	T_{nmse}	Dataset	t_{mse}	T_{mse}	t_{nmse}	T_{nmse}
fri.c2_1000_10	0.136	0.157	0.138	0.151	sl_ex1605	95.070	103.057	0.411	0.423
fri.c1_500_10	0.138	0.153	0.140	0.153	fri.c0_250_50	0.345	0.431	0.340	0.423
fri.c2_1000_25	0.154	0.156	0.154	0.160	fri.c3_250_25	0.323	0.407	0.322	0.426
fri.c3_500_5	0.128	0.159	0.131	0.160	rmftsa_ladata	3.135	3.208	0.384	0.433
fri.c1_1000_25	0.155	0.175	0.156	0.167	a.vehicle	2.17e4	3.73e4	0.279	0.440
fri.c4_1000_25	0.145	0.176	0.145	0.179	fri.c1_250_10	0.379	0.458	0.384	0.449
fri.c3_1000_10	0.156	0.168	0.154	0.180	fri.c3_1000_50	0.383	0.449	0.387	0.453
FacultySalary	1.751	3.127	0.083	0.180	fri.c2_100_5	0.328	0.500	0.326	0.459
machine_cpu	2.24e3	3.74e3	0.089	0.190	fri.c1_100_10	0.365	0.472	0.369	0.462
fri.c1_250_5	0.128	0.187	0.126	0.190	LEV	0.409	0.418	0.444	0.465
fri.c2_250_5	0.147	0.187	0.145	0.190	vineyard	4.734	7.993	0.242	0.475
cloud	0.095	0.165	0.083	0.195	fri.c4_1000_50	0.423	0.445	0.417	0.476
fri.c0_1000_5	0.182	0.197	0.182	0.200	fri.c2_100_10	0.382	0.537	0.399	0.482
fri.c2_500_10	0.178	0.217	0.178	0.211	fri.c0_100_10	0.277	0.504	0.279	0.487
fri.c4_500_25	0.164	0.212	0.163	0.218	fri.c4_500_10	0.436	0.484	0.407	0.488
auto_price	6.54e6	7.62e6	0.187	0.219	fri.c4_1000_1000	0.526	0.541	0.520	0.546
elusage	86.700	1.19e2	0.147	0.219	fri.c1_250_50	0.457	0.532	0.452	0.551
autoPrice	6.35e6	7.56e6	0.184	0.220	no2	0.303	0.320	0.531	0.563
fri.c3_500_25	0.177	0.242	0.177	0.225	fri.c3_500_50	0.582	0.605	0.582	0.595
USCrime	2.02e2	3.72e2	0.127	0.227	fri.c1_100_5	0.397	0.579	0.395	0.598
fri.c2_1000_50	0.214	0.230	0.214	0.230	pollution	1.29e3	2.40e3	0.319	0.623
chs_geyser1	36.539	38.619	0.223	0.237	ERA	2.553	2.665	0.653	0.671
fri.c2_500_25	0.219	0.228	0.217	0.242	sl_case2002	55.722	59.521	0.591	0.674
vinnie	2.375	2.371	0.253	0.242	fri.c4_500_50	0.625	0.664	0.623	0.680
fri.c2_250_10	0.204	0.239	0.210	0.248	v.env.	7.380	8.412	0.574	0.687
fri.c3_250_5	0.202	0.274	0.196	0.252	fri.c4_250_10	0.594	0.702	0.587	0.697
fri.c0_500_5	0.210	0.258	0.206	0.258	fri.c3_100_5	0.469	0.679	0.463	0.698
fri.c0_1000_10	0.254	0.260	0.253	0.262	fri.c0_100_25	0.450	0.778	0.461	0.745
fri.c0_1000_25	0.222	0.258	0.223	0.266	SWD	0.484	0.488	0.736	0.749
fri.c1_500_25	0.204	0.278	0.203	0.275	pm10	0.681	0.693	0.870	0.885

3.2. Performance Comparison with State-of-the-art Algorithms

We compared our method with a selection of Genetic Programming (GP) and Machine Learning (ML) based regression approaches presented in Orzechowski et al. (2018) for the ranking computed over MSE scores. In their work

each of the GP-based algorithms was ran for 100,000 evaluations (i.e. population size $\times$ number of generations, see details in Orzechowski et al. (2018)), except the **eplex-1m** that was allowed to run for 1 million generations. We
 655 labelled our proposal as ‘CFR’ (standing for the fact that it is based on ‘Continued Fraction Regression’) and it was ran 10 times to compare with the outcomes presented in (Orzechowski et al., 2018) (that also run the methods 10 times). A detailed description of the algorithms that are compared against CFR can be found in (Orzechowski et al., 2018), however, we are listing them
 660 to enhance the readability of the paper. The methods based on Genetic Programming (GP) presented in the comparison are **afp**: Age-fitness Pareto Optimization (Schmidt & Lipson, 2011), **eplex**: ϵ -Lexicase selection, **eplex-1m**: a variation of **eplex** with stopping criteria of one million evaluations (population size $\times$ generations) (La Cava et al., 2016), **gsgp**: Geometric Semantic Genetic
 665 Programming (Moraglio et al., 2012) and **mrpgp**: Multiple Regression Genetic Programming (Arnaldo et al., 2014).

The next category of regressors is based on Machine Learning Algorithms. Those are **adaboost**: Adaptive Boosting (AdaBoost) Regression (**ada-b**) (Drucker, 1997), **gradboost**: Gradient Boosting Regression (**grad-b**) (Friedman, 2000),
 670 **kernel-ridge**: Kernel Ridge (**krnl-r**) (Murphy, 2012), **lasso-lars**: Least-Angle Regression with Lasso (**lasso-l**) (Tibshirani, 1994), **linear-regression**: Linear Regression (**l-regr**) (Efron et al., 2004), **linear-svr**: Linear Support Vector Regression (**l-svr**) (Smola & Schölkopf, 2004), **mlp**: Multilayer Perceptrons (MLPs) Regressor (Kingma & Ba, 2014), **rf**: Random Forests Regression
 675 (**sgd-r**) (Pedregosa et al., 2011) and **xgboost**: Extreme Gradient Boosting (**xg-b**) (Chen & Guestrin, 2016).

To illustrate the differences in the performances of the algorithms, we will use *violin plots*⁴. These plots combine the *box-and-whisker plot* with the quantitative distribution of the results. The box-and-whisker plot represents the five
 680

⁴<https://seaborn.pydata.org/generated/seaborn.violinplot.html>

(a) Training Performance of Algorithms

(b) Testing Performance of Algorithms

Figure 4: Median Ranking of the performance of CFR compared with GP-based algorithms (as reported in Orzechowski et al. (2018)) for the MSE score on a) Training and b) Testing datasets.

number descriptive statistics. The ends of the box represent the *upper* and *lower quartiles*, a vertical line inside the box represent the *median* and two whiskers outside the box extend to the *highest* and *lowest* value of the observations. In addition to these statistics, the colored violin shows the quantitative distribution of the results.

3.2.1. Performance Comparison with GP-based algorithms

Fig. 4 shows the median ranking of the GP-based approaches (including the proposed CFR) for 10 repeated runs on the 94 benchmark datasets. From the performance in Fig. 4(a) on the training set, we can observe that **mrgp** achieved the best median ranking among six algorithms. The proposed CFR appeared the 3rd in median ranking for the training set. However, observing the generalization

performance of those algorithms on the testing set (Fig. 4(b)), our method has improved to 2nd rank. We note now that **mrGP**, the best-performing approach in the training set, is now in the 3rd place for testing (and very close to **eplex**). In addition, **mrGP**'s 75th percentile ranking is the 2nd worst performance exhibited by all GP-based methods in the testing set. If we consider the 75th percentile of the median ranking on testing, CFR appears as the best regression method among the GP-based approaches. In terms of generalization performances, we can claim that the proposed CFR exhibited better performance than many of the state-of-the-art GP-based regression approaches.

3.2.2. Performance Comparison with ML-based algorithms

Fig. 5 shows the median ranking of the 10 ML-based approaches and the proposed CFR for 10 repeated runs on the 94 benchmark datasets. On the training set (Fig. 5(a)), we can observe that the **gradboost** achieved the best middle quartile performance. The proposed CFR is, at least, better than four rightmost machine learning algorithms (**linear-svr**, **linear-regression**, **lasso-lars** and **sgd-regression**).

When we look at the generalization performance of those Machine Learning (ML)-based algorithms on the testing set (Fig. 5(b)), **CFR** achieved the joint top position with the **xgboost** and **gradboost** while considering the middle quartile. We note that **gradboost** was unable to achieve 1st ranking position for any of the testing sets, despite it being ranked 1st for some experiments in testing sets.

3.3. Statistical Significance Testing

To check whether there are any significant differences in the median rankings of the algorithms based on their MSE performances, we have used a modification of the Friedman test (Friedman, 1937) by Iman & Davenport (1980). The obtained p -value (p -value $< 2.2e-16$, Corrected Friedman's chi-squared = 68.089 with $df1 = 15$, $df2=1395$) indicates that the null hypothesis "*all the algorithms perform the same*" can safely be rejected. Therefore, we proceeded with the post-hoc test.

(a) Training Performance of Algorithms

(b) Testing Performance of Algorithms

Figure 5: Median Ranking of the performance of CFR compared with ML-based algorithms (as reported in Orzechowski et al. (2018)) for the MSE score on a) Training and b) Testing datasets. The leftmost three are the boosting-based methods.

Figure 6: Critical Difference (CD) diagram proposed in Demšar (2006), which is based on the Nyemeni post-hoc test. We have plotted the CD graph using the implementation of Calvo & Santafé Rodrigo (2016) for our experiments with a significance threshold of $p = 0.05$. From the plot it is clear that there is *no significant differences* found in the rankings of CFR, eplex-1m, xgboost, gradboost algorithms on 94 benchmark datasets. Moreover, the median ranking of CFR is within 3 to 4, which is the best result among all of the algorithms.

To better understand the post-hoc results, we looked at the Critical Difference diagram using the implementation of Calvo & Santafé Rodrigo (2016) (in Fig. 6) of rankings. The algorithms are placed on the x-axis at the place of their median ranking, then statistically “non significant” (at $p \leq 0.05$) algorithms are connected with horizontal lines.

3.4. Runtime Comparison of the Algorithms

To compare the running time requirements of CFR with GP-based methods, we used an Ubuntu 18.04 Computer with 4 CPU cores and 20 GB RAM to execute all algorithms. We also selected 10 datasets (USCrime, Fac.Salaries, autoPrice, elusage, anlc_vehicle, anlc_neavote, anlc_elect.2000, sl_ex1714, rabe_266, chatfield_4) for which the GP methods required the least amount of reported runtime in Orzechowski et al. (2018) to facilitate this benchmarking. We have limited the running time for each of the algorithms to 72 hours in our experimental setup.

It is notable that the eplex-1m was able to finish the execution in only one of the datasets (USCrime) using 3115.36 times more CPU than CFR. Now we compare the ratio of the required running time by the reported value. We found that the running time of afp in the experimental computer is higher by

Figure 7: Total running time of the algorithms in seconds (in logarithmic scale) for 10 datasets as in a) Total running time from Orzechowski et al. (2018) b) Total running time required for the algorithms in our experimental computer for the 10 datasets. In the experimental computer, **eplex-1m** only completed one run in a dataset so the total time is expected to be at least ten times more than the one in the graphic (we indicated it with an arrow that shows this fact).

a factor of 4.93 than the reported values using the hardware of Orzechowski et al. (2018). Similarly, the runtime of **eplex** has increased by a factor of 5.17 in the experimental computer. Besides, **gsgp** required an average CPU times by a factor of 0.93 than the values reported in Orzechowski et al. (2018). Finally, **mrqp** has exhibited a runtime increment by a factor of 4.18 in the experimental environment. In the experimental setup, we have found that the CFR outperformed the GP-based algorithms in terms of running time.

3.5. The performance profile of all algorithms

Dolan & Moré (2002) proposed an approach to visualize the information that results from running multiple optimization algorithms over a set of datasets. This approach leads to the creation of a ‘performance profile’. We will use it here to show the results of the 16 learning algorithms over the 94 datasets studied.

The *performance profile* of an optimization is an (x, y) -plot and there are

as many discontinuous functions in the graph as algorithms are compared. To create it, we first need to identify, for each of the 94 datasets on which general-
 755 ization has been tested, the minimum average MSE value observed (i.e. the best average MSE result obtained by one of the 16 algorithms, where the average is obtained using the 10 independent runs). This value is then taken as a useful reference; it then used to calculate the percentual relative error differences observed. Accordingly, assume that we have the curves of two algorithms A_1
 760 and A_2 that respectively cross the points (500, 30) and (500, 70). The x value of 500 indicates that for a relative error difference corresponding to 500 percent to the minimum average MSE value observed, algorithm A_1 has obtained such relative error difference (on average) in only 30 percent of the 94 datasets. In comparison, the algorithm A_2 has achieved such a performance in 70 percent of
 765 the datasets. Varying the value of the percentual relative error difference, we can generate a curve that will go from 0 to 100 percent (for all algorithms). On a given interval on the x -axis, an algorithm A would dominate an algorithm B if the y value of any point of its performance profile curve is larger than the of A for the same value of x , regardless of which value of x we are choosing on
 770 that interval.

Fig 8 shows that few of the 16 learning algorithms achieved the top 100 percent mark. Sub-fig. 8(b) shows the profiling on the x -axis interval of [0,500]. CFR is the first in achieving it and does it at a percentual relative error of 451. `xgboost` achieves the same feat at 547 later followed by `eplex-1m` at 615.

775 3.6. Median MSE scores of Top 4 algorithms for 94 Datasets

We tabulated the performance (measured in MSE score) of the top four methods in Table 3 that showed no significant differences of performances in Fig.6.

(a) All Algorithms on the x -axis interval from 0 to 2000 (b) All Algorithms on the x -axis interval from 0 to 500

Figure 8: Performance Profiles for all algorithms on the x -axis on two different intervals: a) $[0, 2000]$ and b) $[0, 500]$. The performance profiles of **CFR**, **eplex-1m**, **xgboost**, and **grad-boost**, relative to the other curves, indicate that for some datasets the other algorithms present errors in generalization that are 5 times larger than the best one observed. This speaks of the relative robustness of **CFR**, **eplex-1m**, **xgboost** and **grad-boost** in generalization performances.

Table 3: The median value of MSE score for 10 runs obtained in Test sets of the best four algorithms. A bold-face in ‘CFR’ column indicates it has achieved the best score among four algorithms for corresponding dataset.

Dataset	eplex- 1m	grad- b	xg-b	CFR	Dataset	eplex- 1m	grad- b	xg-b	CFR
ESL	0.274	0.319	0.272	0.268	fri_c3_1000_10	0.077	0.066	0.062	0.068
SWD	0.390	0.405	0.408	0.432	fri_c0_1000_5	0.042	0.074	0.071	0.106
LEV	0.425	0.424	0.422	0.353	fri_c3_100_5	0.250	0.182	0.093	0.125
ERA	2.514	2.585	2.567	2.450	fri_c1_1000_5	0.049	0.047	0.047	0.050
USCrime	3.93e2	2.57e2	3.78e2	2.20e2	fri_c3_250_5	0.126	0.110	0.108	0.086
FacultySalry	4.035	8.071	4.111	1.277	fri_c4_250_10	0.148	0.173	0.172	0.133
vineyard	6.010	8.223	7.825	4.222	fri_c4_500_50	0.079	0.141	0.114	0.169
auto_price	5.89e6	3.89e6	4.03e6	6.01e6	fri_c3_500_5	0.103	0.085	0.071	0.078
autoPrice	4.17e6	5.29e6	2.87e6	4.83e6	fri_c3_1000_50	0.068	0.082	0.070	0.120
cloud	0.110	0.208	0.144	0.095	fri_c1_1000_25	0.057	0.070	0.067	0.076
elusage	1.35e2	1.99e2	1.37e2	65.797	fri_c0_100_10	0.149	0.319	0.307	0.226
machine_cpu	3.80e3	2.23e3	2.69e3	2.10e3	fri_c2_1000_50	0.063	0.069	0.076	0.118
a.vehicle	4.14e4	2.41e4	4.20e4	1.55e4	fri_c4_1000_10	0.050	0.074	0.060	0.063
vinnie	2.287	2.860	2.659	1.934	fri_c0_100_5	0.152	0.193	0.164	0.202
pm10	0.640	0.431	0.399	0.621	fri_c2_500_50	0.049	0.101	0.109	0.130
a.neavote	1.180	0.818	0.917	0.401	fri_c2_500_10	0.064	0.088	0.097	0.081
a.elect2000	4.33e7	3.40e8	7.72e8	5.09e5	fri_c3_1000_5	0.059	0.049	0.048	0.046
pollution	1.87e3	2.19e3	1.67e3	1.42e3	fri_c1_500_5	0.068	0.077	0.074	0.075
no2	0.272	0.227	0.210	0.295	fri_c0_500_25	0.047	0.132	0.127	0.144
a.apnea2	1.12e6	9.42e5	7.86e5	6.09e5	fri_c2_100_10	0.750	0.321	0.233	0.320
a.apnea1	8.16e5	9.98e5	5.28e5	6.96e5	fri_c0_250_10	0.047	0.183	0.170	0.259
cpu	1.75e2	2.36e3	8.83e2	1.64e2	fri_c1_500_50	0.091	0.111	0.133	0.116
fri_c0_250_5	0.065	0.162	0.165	0.176	fri_c1_500_10	0.062	0.086	0.069	0.075
fri_c3_500_25	0.064	0.090	0.097	0.081	fri_c2_500_25	0.077	0.097	0.110	0.121
fri_c1_500_25	0.117	0.108	0.107	0.111	fri_c4_250_25	0.157	0.206	0.178	0.110
fri_c1_1000_50	0.066	0.073	0.077	0.104	fri_c3_500_50	0.101	0.132	0.131	0.156
fri_c4_500_25	0.127	0.111	0.095	0.111	fri_c3_500_10	0.062	0.079	0.066	0.075
fri_c3_1000_25	0.056	0.069	0.059	0.064	fri_c1_250_10	0.114	0.123	0.105	0.130
fri_c4_1000_100	0.047	0.087	0.072	0.227	fri_c1_250_50	0.098	0.171	0.154	0.140
fri_c2_1000_25	0.051	0.064	0.070	0.077	fri_c0_500_5	0.048	0.099	0.106	0.124

Continue on the next page

Table 3 – continued from previous page

Dataset	eplex- 1m	grad- b	xg-b	CFR	Dataset	eplex- 1m	grad- b	xg-b	CFR
fri_c0_1000_50	0.044	0.111	0.113	0.141	fri_c0_500_50	0.053	0.131	0.170	0.211
fri_c1_100_10	0.485	0.245	0.154	0.233	fri_c0_100_25	0.259	0.419	0.383	0.365
fri_c4_1000_25	0.058	0.076	0.066	0.081	fri_c0_250_25	0.064	0.209	0.209	0.224
fri_c1_1000_10	0.047	0.051	0.049	0.059	fri_c0_500_10	0.048	0.153	0.138	0.202
fri_c2_100_5	0.509	0.277	0.255	0.259	fri_c1_100_5	0.490	0.208	0.176	0.241
fri_c0_1000_10	0.047	0.096	0.093	0.155	fri_c2_250_10	0.128	0.123	0.116	0.140
fri_c2_250_5	0.104	0.113	0.094	0.086	fri_c3_250_25	0.116	0.194	0.184	0.111
fri_c2_500_5	0.079	0.068	0.073	0.054	sl_ex1714	1.42e6	1.57e6	2.31e6	6.83e5
fri_c0_1000_25	0.039	0.092	0.093	0.074	rabe_266	7.107	7.316	3.043	2.643
fri_c2_1000_5	0.046	0.048	0.045	0.044	sl_case2002	75.772	56.240	72.363	41.742
fri_c1_250_5	0.084	0.103	0.093	0.093	rmftsa_ladata	3.012	3.511	3.212	2.764
fri_c3_250_10	0.086	0.150	0.120	0.084	v.env.	9.624	9.798	9.540	4.700
fri_c0_250_50	0.071	0.238	0.252	0.254	sl_ex1605	101.824	98.440	91.986	83.998
fri_c4_500_10	0.065	0.084	0.072	0.072	v.galaxy	3.13e2	2.68e2	2.24e2	4.34e2
fri_c2_250_25	0.179	0.173	0.156	0.149	chatfield_4	2.82e2	3.84e2	2.88e2	1.89e2
fri_c2_1000_10	0.069	0.062	0.058	0.073	sl_case1202	3.29e3	3.42e3	3.20e3	1.39e3
fri_c4_1000_50	0.051	0.088	0.068	0.092	chs_geyser1	38.888	39.902	42.887	31.053

4. Discussion

After analyzing the median ranking of state-of-the-art regression algorithms (both of the GP and ML-based) we found that our proposed approach comparable performance while compared with all state-of-the-art approaches. The results show that the use of continued fractions is a promising new idea due to its representational power.

The statistical test on the results (in Fig. 6) showed that there are no statistically significant differences for the performance of CFR with the **eplex-1m**, **xgboost** and **grad-boost** algorithm. Nonetheless, running the **eplex** for a million evaluations (in **eplex-1m**) improves the performance as expected. In contrast, we note that we have executed the CFR for only 200 generations. Moreover, the average ranking of CFR (which was consistently between 3 and

4) for 94 datasets makes it ahead of all of the algorithms in the CD plot.

Recently (Kommenda et al., 2020) proposed the hybridization of a Genetic Programming (GP) approach by using the Levenberg–Marquardt algorithm for parameter optimization of the regression models and calculated gradients via automatic differentiation; this is well in resonance within the spirit of a memetic computing approach. They used the same benchmark datasets as ours and employed different GP variants. The methods were executed 50 times on a 75-25 train-test split and the results are available from the website of the Heuristic and Evolutionary Algorithms Laboratory (HEAL)⁵. For each of the 94 problems they have tuned the *maximum tree length nodes*, *population size* and *maximum generations* using grid search and five-fold cross-validation (in (Kommenda et al., 2020), Sec.4. p.484). This contrasts with our approach in which no problem-specific hyperparameters have been tuned. If we disregard this important difference, we can consider the results of the first 10 runs per dataset and we can compare these with performance on the 10 runs in this paper. We have found that their two best methods, GP NLS and OSGP NLS had no significant differences in performances with CFR, `eplex-1m`, `xgboost` and `grad-bost`. This speaks very well of this new method and it also points that the introduction of gradient information may also be beneficial for CFR in future memetic approaches for numerical regression when the choice of loss function allows differentiability.

Another important issue to consider is that the CFR results have been obtained with minimal level of parameter tuning. In contrast, the hyperparameters of the state-of-the-art algorithms in the experiment were extensively tuned using grid-search for the experiments as reported by Orzechowski et al. (2018). Similarly, the performance of the CFR could be improved if the parameters of the algorithms are tuned per datasets or by including techniques such as *boosting*. Hence, we believe there is a clear avenue for further research in this area with the joint optimization of parameters for the inclusion of boosting in

⁵Results.csv file is downloaded from HEAL website at: <https://dev.heuristiclab.com/trac.fcgi/wiki/AdditionalMaterial#GPEMSpecialIssue-IntegratingNumericalOptimizationMethodswithGeneticProgramming>

a new type of memetic algorithm.

820 5. Conclusion

In this paper, we present a comprehensive experimental comparison of an alternative method for multivariate regression problems that uses analytic continued fractions as a representation of the mathematical models. This has led to a challenging type of nonlinear optimization problem and we have presented a memetic algorithm for this problem using a hierarchical structured population. A variant of the classical Nelder-Mead-based search was proposed as the individual optimization. We compared the performances of our proposed method with 10 machine-learning and five GP-based regression methods on 94 benchmark datasets. The results indicated a better generalization performance than many state-of-the-art regression approaches.

Fig. 8 shows that in comparison with the top 4 algorithms, our proposed method is the first at reaching ceiling of 100 percent making it more robust at a fraction of the computational time required to obtain the models with the best GP-based approach (**eplex-1m**) (e.g. for one dataset (US Crime) **mr**gp and **eplex-1m** was 1040 times and 3120 times slower, respectively).

The proposed approach can be extended by including adaptations that could explore parameter optimization of the memetic algorithm to improve the performance in generalization. We foresee that more sophisticated gradient descent methods could be used in place of the Nelder-Mead solver as an optimizer in our memetic algorithm.

In terms of running time, our approach with GP-based methods, the memetic algorithm exhibited better performances compared against all of the GP-based methods. **CFR** required the least amount of CPU time to execute on 10 selected datasets. The **eplex-1m** was the most compute extensive GP-based method. We encourage the computing community to explore the possibilities of this new approach for regression involving analytic continued fractions and to extend to other domains of machine learning and artificial intelligence (e.g. classification).

Acknowledgements

We thank Dr Markus Wagner, School of Computer Science at The University
 850 of Adelaide, Australia for his thoughtful comments that helped us to improve an
 earlier version of the manuscript. We also thank the members of Prof. Jason H.
 Moore’s research lab at the University of Pennsylvania, USA, for making both
 the source code of their experiments and the Penn Machine Learning Bench-
 marks datasets available. M.N.H. and P.M. thank Renata Sarment from the
 855 Universidade Federal de São Carlos, Sao Paulo, Brazil for discussion about the
 performance profile plot and sharing her Python code to produce one of the
 figures.

References

- Akal, C., & Lukashov, A. (2018). Newton-padé approximations for multivariate
 860 functions. *Applied Mathematics and Computation*, 334, 367–374.
- Angeline, P. J. (1997). Subtree crossover: Building block engine or macromu-
 tation. *Genetic Programming*, 97, 9–17.
- Arnaldo, I., Krawiec, K., & O’Reilly, U.-M. (2014). Multiple regression genetic
 programming. In *Proceedings of the 2014 Annual Conference on Genetic and*
 865 *Evolutionary Computation GECCO ’14* (pp. 879–886). New York, NY, USA:
 ACM.
- Azad, R. M. A., & Ryan, C. (2014). A Simple Approach to Lifetime Learning
 in Genetic Programming-Based Symbolic Regression. *Evolutionary Compu-
 tation*, 22, 287–317. PMID: 24079729.
- 870 Backeljauw, F., & Cuyt, A. A. M. (2009). Algorithm 895: A continued fractions
 package for special functions. *ACM Trans. Math. Softw.*, 36, 15:1–15:20.
- Baker Jr, G. A. (2012). Padé approximant. URL: [http://scholarpedia.org](http://scholarpedia.org/article/Pad%C3%A9_approximant)
 /article/Pad%C3%A9_approximant Accessed April 15, 2019.

- 875 Berretta, R., Cotta, C., & Moscato, P. (2004). Enhancing the performance of
memetic algorithms by using a matching-based recombination algorithm. In
M. G. C. Resende, J. P. de Sousa, & A. Viana (Eds.), *Metaheuristics: Com-
puter Decision-Making* (pp. 65–90). Norwell, MA, USA: Kluwer Academic
Publishers.
- 880 Berretta, R., & Rodrigues, L. F. (2004). A memetic algorithm for a multi-
stage capacitated lot-sizing problem. *International Journal of Production
Economics*, 87, 67 – 81.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32.
- Buriol, L., França, P. M., & Moscato, P. (2004). A new memetic algorithm
for the asymmetric traveling salesman problem. *Journal of Heuristics*, 10,
885 483–506.
- Cagnoni, S., Rivero, D., & Vanneschi, L. (2005). A purely evolutionary memetic
algorithm as a first step towards symbiotic coevolution. In *2005 IEEE
Congress on Evolutionary Computation (CEC), Vols 1-3, Proceedings IEEE
Congress on Evolutionary Computation* (pp. 1156–1163). IEEE; IEEE Com-
890 putat Intelligence Soc; IEE; Evolut Programming Soc. IEEE Congress on
Evolutionary Computation, Edinburgh, SCOTLAND, SEP 02-05, 2005.
- Calvo, B., & Santafé Rodrigo, G. (2016). scmamp: Statistical comparison of
multiple algorithms in multiple problems. *The R Journal*, Vol. 8/1, Aug.
2016, .
- 895 Chaffy, C. (1986). How to compute multivariate pade approximants. In B. W.
Char (Ed.), *SYMSAC 1986, Proceedings of the Symposium on Symbolic and
Algebraic Manipulation, Waterloo, Ontario, Canada, July 21-23, 1986* (pp.
56–58). ACM.
- Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In
900 *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowl-*

edge Discovery and Data Mining KDD '16 (pp. 785–794). New York, NY, USA: ACM.

Clegg, J., Walker, J. A., & Miller, J. F. (2007). A new crossover technique for cartesian genetic programming. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation GECCO '07* (pp. 1580–1587). New York, NY, USA: ACM.

Cotta, C., Mendes, A., Garcia, V., França, P., & Moscato, P. (2003a). Applying memetic algorithms to the analysis of microarray data. In S. Cagnoni, C. G. Johnson, J. J. R. Cardalda, E. Marchiori, D. W. Corne, J.-A. Meyer, J. Gottlieb, M. Middendorf, A. Guillot, G. R. Raidl, & E. Hart (Eds.), *Applications of Evolutionary Computing* (pp. 22–32). Berlin, Heidelberg: Springer Berlin Heidelberg.

Cotta, C., Mendes, A., Garcia, V., França, P. M., & Moscato, P. (2003b). Applying memetic algorithms to the analysis of microarray data. In G. R. Raidl, J. Meyer, M. Middendorf, S. Cagnoni, J. J. R. Cardalda, D. Corne, J. Gottlieb, A. Guillot, E. Hart, C. G. Johnson, & E. Marchiori (Eds.), *Applications of Evolutionary Computing, EvoWorkshop 2003: EvoBIO, EvoCOP, EvoIASP, EvoMUSART, EvoROB, and EvoSTIM, Essex, UK, April 14-16, 2003, Proceedings* (pp. 22–32). Springer volume 2611 of *Lecture Notes in Computer Science*.

Cotta, C., & Moscato, P. (2002). Inferring phylogenetic trees using evolutionary algorithms. In J. J. Merelo Guerv'os, P. Adamidis, H. Beyer, J. L. F. Martín, & H. Schwefel (Eds.), *Parallel Problem Solving from Nature - PPSN VII, 7th International Conference, Granada, Spain, September 7-11, 2002, Proceedings* (pp. 720–729). Springer volume 2439 of *Lecture Notes in Computer Science*.

Cotta, C., & Moscato, P. (2003). A memetic-aided approach to hierarchical clustering from distance matrices: application to gene expression clustering and phylogeny. *Biosystems*, 72, 75 – 97. Computational Intelligence in Bioinformatics.

- 930 Crandall, R. E. (1994). *Projects in Scientific Computation*. Berlin, Heidelberg: Springer-Verlag.
- Delgado, R. R., Ruíz, L. G. B., Cuéllar, M. P., Calvo-Flores, M. D., & del Carmen Pegalajar Jiménez, M. (2018). A comparison between NARX neural networks and symbolic regression: An application for energy consumption forecasting. In J. Medina, M. Ojeda-Aciego, J. L. V. Galdeano, I. Perfilieva, B. Bouchon-Meunier, & R. R. Yager (Eds.), *Information Processing and Management of Uncertainty in Knowledge-Based Systems. Applications - 17th International Conference, IPMU 2018, Cádiz, Spain, June 11-15, 2018, Proceedings, Part III* (pp. 16–27). Springer volume 855 of *Communications in Computer and Information Science*.
940
- Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of Machine learning research*, 7, 1–30.
- Dick, G. (2014). Bloat and generalisation in symbolic regression. In G. Dick, W. N. Browne, P. A. Whigham, M. Zhang, L. T. Bui, H. Ishibuchi, Y. Jin, X. Li, Y. Shi, P. Singh, K. C. Tan, & K. Tang (Eds.), *Simulated Evolution and Learning - 10th International Conference, SEAL 2014, Dunedin, New Zealand, December 15-18, 2014. Proceedings* (pp. 491–502). Springer volume 8886 of *Lecture Notes in Computer Science*.
945
- Dolan, E. D., & Moré, J. J. (2002). Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91, 201–213.
950
- Drucker, H. (1997). Improving regressors using boosting techniques. In *Proceedings of the Fourteenth International Conference on Machine Learning ICML '97* (pp. 107–115). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- 955 Duffy, J., & Engle-Warnick, J. (2002). Using symbolic regression to infer strategies from experimental data. In S.-H. Chen (Ed.), *Evolutionary Computation in Economics and Finance* (pp. 61–82). Heidelberg: Physica-Verlag HD.

- Efron, B., Hastie, T., Johnstone, I., & Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, 32, 407–499.
- 960 Eskridge, B. E., & Hougen, D. F. (2004). Memetic crossover for genetic programming: Evolution through imitation. In K. Deb (Ed.), *Genetic and Evolutionary Computation – GECCO 2004* (pp. 459–470). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Euler, L. (1748). *Introductio in analysin infinitorum*. chapter 18. volume 1.
 965 Reprinted as *Opera* (1)8.
- Fajfar, I., Puhan, J., & Bürmen, A. (2016). Evolving a Nelder-Mead algorithm for optimization with genetic programming. *Evolutionary computation*, 25.
- Ffrancon, R., & Schoenauer, M. (2015). Memetic semantic genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2015, Madrid, Spain, July 11-15, 2015* (pp. 1023–1030).
 970
- Fitzsimmons, J., & Moscato, P. (2018). Symbolic regression modelling of drug responses. In *First IEEE Conference on Artificial Intelligence for Industries, Sep 26, 2018 - Sep 28, 2018, Laguna Hills, CA, USA*.
- Frade, M., de Vega, F. F., & Cotta, C. (2009). Breeding terrains with genetic
 975 terrain programming: The evolution of terrain generators. *Int. J. Computer Games Technology*, 2009, 125714:1–125714:13.
- Friedman, J. H. (2000). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29, 1189–1232.
- Friedman, M. (1937). The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the american statistical association*, 32, 675–701.
 980
- Graham, M., Djorgovski, S., Mahabal, A., Donalek, C., & Drake, A. (2013). Machine-assisted discovery of relationships in astronomy. *Monthly Notices of the Royal Astronomical Society*, 431, 2371–2384.

- 985 Harris, M., Berretta, R., Inostroza-Ponta, M., & Moscato, P. (2015). A memetic algorithm for the quadratic assignment problem with parallel local search. In *IEEE Congress on Evolutionary Computation, CEC 2015, Sendai, Japan, May 25-28, 2015* (pp. 838–845).
- Iman, R. L., & Davenport, J. M. (1980). Approximations of the critical region
990 of the fbietkan statistic. *Communications in Statistics-Theory and Methods*, 9, 571–595.
- Inostroza-Ponta, M., Berretta, R., & Moscato, P. (2011). Qapgrid: A two level qap-based approach for large-scale data analysis and visualization. *PLOS ONE*, 6, 1–18.
- 995 Karaboga, D., Ozturk, C., Karaboga, N., & Gorkemli, B. (2012). Artificial bee colony programming for symbolic regression. *Information Sciences*, 209, 1–15.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, .
- 1000 Kommenda, M., Burlacu, B., Kronberger, G., & Affenzeller, M. (2020). Parameter identification for symbolic regression using nonlinear least squares. *Genetic Programming and Evolvable Machines*, 21, 471–501.
- Koza, J. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: The MIT Press.
- 1005 Koza, J. R. (1990). Non-linear genetic algorithms for solving problems. United States Patent 4935877. Filed may 20, 1988, issued june 19, 1990, 4,935,877. Australian patent 611,350 issued september 21, 1991. Canadian patent 1,311,561 issued december 15, 1992.
- 1010 Kronberger, G., Kommenda, M., Promberger, A., & Nickel, F. (2018). Predicting friction system performance with symbolic regression and genetic programming with factor variables. In *Proceedings of the Genetic and Evolution-*

ary Computation Conference, *GECCO 2018, Kyoto, Japan, July 15-19, 2018* (pp. 1278–1285).

La Cava, W., Spector, L., & Danai, K. (2016). Epsilon-lexicase selection for
 1015 regression. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 GECCO '16* (pp. 741–748). New York, NY, USA: ACM.

de Lima Correa, L., Borguesan, B., Krause, M. J., & Dorn, M. (2018). Three-dimensional protein structure prediction based on memetic algorithms. *Computers & OR*, *91*, 160–177.

1020 Lorentzen, L. (2010). Padé approximation and continued fractions. *Applied Numerical Mathematics*, *60*, 1364 – 1370.

Luke, S., & Spector, L. (1997). A comparison of crossover and mutation in genetic programming. *Genetic Programming*, *97*, 240–248.

Märtens, M., Kuipers, F. A., & Mieghem, P. V. (2017). Symbolic regression
 1025 on network properties. In *Genetic Programming - 20th European Conference, EuroGP 2017, Amsterdam, The Netherlands, April 19-21, 2017, Proceedings* (pp. 131–146).

Mendes, A., Franca, P., & Moscato, P. (2002a). Fitness landscapes for the total tardiness single machine scheduling problem. *Neural Network World*, *12*,
 1030 165–180.

Mendes, A., França, P. M., Moscato, P., & Garcia, V. (2002b). Population studies for the gate matrix layout problem. In *Advances in Artificial Intelligence - IBERAMIA 2002, 8th Ibero-American Conference on AI, Seville, Spain, November 12-15, 2002, Proceedings* (pp. 319–339).

1035 Miller, S. J. (2015). The gamma function and related distributions. In *The Probability Lifesaver* book chapter 15. (pp. 413–438). Princeton University Press, Princeton, NJ.

- Moraglio, A., Krawiec, K., & Johnson, C. G. (2012). Geometric semantic genetic programming. In C. A. C. Coello, V. Cutello, K. Deb, S. Forrest, G. Nicosia, & M. Pavone (Eds.), *Parallel Problem Solving from Nature - PPSN XII* (pp. 21–31). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Moscato, P. (1989). *On Evolution, Search, Optimization, Genetic Algorithms and Martial Arts: Towards Memetic Algorithms*. C3P Report 826 Caltech Concurrent Computation Program.
- Moscato, P. (2012). Memetic algorithms: The untold story. In *Handbook of Memetic Algorithms* (pp. 275–309).
- Moscato, P., Mendes, A., & Berretta, R. (2007). Benchmarking a memetic algorithm for ordering microarray data. *Biosystems*, 88, 56–75.
- Moscato, P., & Tinetti, F. (1994). *Blending heuristics with a population-based approach: A memetic algorithm for the traveling salesman problem*. Technical Report CeTAD, Report 92-12. Universidad Nacional de La Plata, C.C. 75, 1900 La Plata, Argentina.
- Moscato, P., & de Vries, N. J. (2019a). *Business and Consumer Analytics: New Ideas*. Springer.
- Moscato, P., & de Vries, N. J. (2019b). Marketing meets data science: Bridging the gap. In *Business and Consumer Analytics: New Ideas* (pp. 3–117).
- Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. The MIT Press.
- Muruzábal, J., Cotta, C., & Fernández, A. (2000). Some probabilistic modelling ideas for boolean classification in genetic programming. In *Genetic Programming* (pp. 133–148).
- Neri, F., Cotta, C., & Moscato, P. (Eds.) (2012). *Handbook of Memetic Algorithms* volume 379 of *Studies in Computational Intelligence*. Springer.
- Olds, C. D. (1963). *Continued fractions*. Technical Report Random House.

- 1065 Olson, R. S., La Cava, W., Orzechowski, P., Urbanowicz, R. J., & Moore, J. H.
(2017). Pmlb: a large benchmark suite for machine learning evaluation and
comparison. *BioData Mining*, 10, 36.
- Orzechowski, P., La Cava, W., & Moore, J. H. (2018). Where are we now?: A
large benchmark study of recent symbolic regression methods. In *Proceedings*
1070 *of the Genetic and Evolutionary Computation Conference GECCO '18* (pp.
1183–1190). New York, NY, USA: ACM.
- Otte, C. (2013). Safe and interpretable machine learning: A methodological
review. In C. Moewes, & A. Nürnberger (Eds.), *Computational Intelligence in*
Intelligent Data Analysis (pp. 111–122). Berlin, Heidelberg: Springer Berlin
1075 Heidelberg.
- P. Searson, D., Leahy, D., & Willis, M. (2010). GPTIPS: An open source genetic
programming toolbox for multigene symbolic regression. In *Int. Multiconfer-*
ence of Engineers and Computer Scientists (pp. 77–80). volume 1.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O.,
1080 Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos,
A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-
learn: Machine learning in python. *J. Mach. Learn. Res.*, 12, 2825–2830.
- Quiñonero-Candela, J., Rasmussen, C. E., Sinz, F., Bousquet, O., & Schölkopf,
B. (2006). Evaluating predictive uncertainty challenge. In J. Quiñonero-
1085 Candela, I. Dagan, B. Magnini, & F. d’Alché Buc (Eds.), *Machine Learning*
Challenges. Evaluating Predictive Uncertainty, Visual Object Classification,
and Recognising Textual Entailment (pp. 1–27). Berlin, Heidelberg: Springer
Berlin Heidelberg.
- Schmidt, M., & Lipson, H. (2009). Distilling free-form natural laws from exper-
1090 imental data. *Science*, 324, 81–85.
- Schmidt, M., & Lipson, H. (2011). Age-fitness pareto optimization. In R. Riolo,

- T. McConaghy, & E. Vladislavleva (Eds.), *Genetic Programming Theory and Practice VIII* (pp. 129–146). New York, NY: Springer New York.
- Schmidt, M., & Lipson, H. (2018). Eureqa (version 1.24.0). URL: <http://www.nutonian.com> Software, Accessed April 15, 2019.
- Semenkina, M., & Semenko, E. (2015). Memetic self-configuring genetic programming for solving machine learning problems. In *IIAI 4th International Congress on Advanced Applied Informatics, IIAI-AAI 2015, Okayama, Japan, July 12-16, 2015* (pp. 599–604). IEEE Computer Society.
- Singer, S., & Nelder, J. (2009). Nelder-Mead algorithm. *Scholarpedia*, 4, 2928. Revision #91557.
- Smola, A. J., & Schölkopf, B. (2004). A tutorial on support vector regression. *Statistics and Computing*, 14, 199–222.
- Sun, H., & Moscato, P. (2019). A memetic algorithm for symbolic regression. In *2019 IEEE Congress on Evolutionary Computation (CEC)*. Wellington, New Zealand: IEEE.
- Sun, S., Ouyang, R., Zhang, B., & Zhang, T.-Y. (2019). Data-driven discovery of formulas by symbolic regression. *Materials Research Society Bulletin*, 44, 559–564.
- Szubert, M., Kodali, A., Ganguly, S., Das, K., & Bongard, J. C. (2016). Reducing antagonism between behavioral diversity and fitness in semantic genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 GECCO '16* (pp. 797–804). New York, NY, USA: ACM.
- Tibshirani, R. (1994). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, 58, 267–288.
- Toledo, C. F. M., da Silva Arantes, M., França, P. M., & Morabito, R. (2012). A memetic framework for solving the lot sizing and scheduling problem in soft

- drink plants. In R. Chiong, T. Weise, & Z. Michalewicz (Eds.), *Variants of Evolutionary Algorithms for Real-World Applications* (pp. 59–93). Springer.
- 1120 Trujillo, L., Z-Flores, E., Juarez Smith, P. S., Legrand, P., Silva, S., Castelli, M., Vanneschi, L., Schütze, O., & Munoz, L. (2017). Local Search is Underused in Genetic Programming. In A. Arbor (Ed.), *Genetic Programming Theory and Practice XIV*. Springer.
- 1125 Virgolin, M., Alderliesten, T., Bel, A., Witteveen, C., & Bosman, P. A. N. (2018). Symbolic regression and feature construction with GP-GOMEA applied to radiotherapy dose reconstruction of childhood cancer survivors. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2018, Kyoto, Japan, July 15-19, 2018* (pp. 1395–1402).
- 1130 Vladislavleva, E., Friedrich, T., Neumann, F., & Wagner, M. (2013). Predicting the energy output of wind farms based on weather data: Important variables and their correlation. *Renewable Energy*, 50, 236 – 243.
- Vladislavleva, E., Smits, G., & den Hertog, D. (2009). Order of nonlinearity as a complexity measure for models generated by symbolic regression via pareto genetic programming. *IEEE Trans. Evolutionary Computation*, 13, 333–349.
- 1135 de Vries, N. J., Carlson, J., & Moscato, P. (2014). A data-driven approach to reverse engineering customer engagement models: Towards functional constructs. *PLOS ONE*, 9, 1–19.
- Wall, H. S. (1948). *Analytic theory of continued fractions*. Courier Dover Publications.
- 1140 Zhou, P., Cuyt, A. A. M., & Tan, J. (2009). General order multivariate padé approximants for pseudo-multivariate functions. II. *Math. Comput.*, 78, 2137–2155.