



Globally Optimal Symbolic Regression

Vernon Austel, Sanjeeb Dash, Oktay Gunluk, Lior Horesh, Leo Liberti,
Giacomo Nannicini, Baruch Schieber

► To cite this version:

Vernon Austel, Sanjeeb Dash, Oktay Gunluk, Lior Horesh, Leo Liberti, et al.. Globally Optimal Symbolic Regression. Interpretable ML, satellite workshop of NIPS 2017, 2017, Long Beach, United States. hal-02105286

HAL Id: hal-02105286

<https://hal.science/hal-02105286v1>

Submitted on 20 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Globally Optimal Symbolic Regression

Vernon Austel* Sanjeeb Dash* Oktay Gunluk* Lior Horesh*

Leo Liberti† Giacomo Nannicini* Baruch Schieber*

Abstract

In this study we introduce a new technique for symbolic regression that guarantees global optimality. This is achieved by formulating a mixed integer non-linear program (MINLP) whose solution is a symbolic mathematical expression of minimum complexity that explains the observations. We demonstrate our approach by rediscovering Kepler’s law on planetary motion using exoplanet data and Galileo’s pendulum periodicity equation using experimental data.

1 Introduction

Discovering mathematical models that explain the behavior of a system has a broad range of applications. Symbolic regression is an important field in machine learning whose goal is to find a symbolic mathematical expression that explains a dependent variable in terms of a number of independent variables for a given data set, most commonly as an explicit function of the independent variables. Unlike traditional (numerical) regression schemes, the functional form of the expression is not assumed to be known *a priori* [KSH17]. The utility of the approach has been established for a broad range of applications [CT77, WHM⁺97, DSW99, DSW03], including the discovery of equations [Tod, Lan81, SL09, SL10].

Starting with [Koz92], symbolic regression problems have been typically solved with genetic programming [BNKF98, AB00], an evolutionary *metaheuristic* related to genetic algorithms; another such technique is grammatical evolution [?]. Even before [Koz92], heuristics to find explicit functional relationships were developed as part of the BACON system, see [PL87]. There has been much research into improving symbolic regression techniques. In [CL17], the separability of the desired functional form is exploited to speed up symbolic regression, whereas a set of candidate basis functions is used in [?] for this purpose. Other recent work has focused on finding accurate constants in the derived symbolic mathematical expression [?]. To discover meaningful functions in the context of physical systems [SL09] search for functions that not only match the data, but also have the property that partial derivatives also match the empirically computed partial derivatives. The same authors search for implicit functional relationships in [SL10]. Another approach populates a large hypothesis space of functional forms, on which sparse selection is applied [BPK16].

In this study we present a Mixed-Integer Non-Linear Programming (MINLP) formulation that produces the simplest symbolic mathematical expression that satisfies a prescribed upper bound on the prediction error. Our MINLP formulation can be solved to optimality using existing state-of-the-art global optimization solvers. The key advantage of our approach is that it produces a *globally optimal* mathematical expression while avoiding exhaustive search of the solution space. Another advantage is that it produces correct real-valued constants (within a tolerance) directly; most other methods use specialized algorithms to refine constants [?, ?] and cannot guarantee global optimality. In addition, our formulation can, in principle, seamlessly incorporate additional constraints and objectives to capture domain knowledge and user preferences.


*IBM T.J. Watson Research Center, Yorktown Heights, NY 10598.

. {austel, sanjeebd, gunluk, lhoresh, nannicini, sbar}@us.ibm.com

†Ecole Polytechnique, 91128 Palaiseau, Paris, France. liberti@polytechnique.fr

2 A MINLP Formulation for Free-Form Model Discovery

A symbolic regression scheme consists of a space of valid mathematical expressions together with a mechanism for its exploration. A mathematical expression can be represented by a (rooted) expression tree where each node is labeled by one of the following entities: operators (such as $+$, $-$, \times , \log), variables (i.e., the independent variables), and constants. The edges of the tree link these entities in a way that is consistent with a prescribed grammar. For example, an expression tree for $2(w_1 + w_2)^3 + 1$ is:



Here $+$, \times and the power function are the operators, w_1, w_2 are the variables, and numbers 1, 2 and 3 are the constants.

In our MINLP based approach we model the grammar of valid expression trees by a set of constraints. The discrete variables of the formulation are used to define the structure of the expression tree, and the continuous variables to evaluate the resulting symbolic expression for specific numerical values associated with the data [HLAN15, ?]. To choose among the many possible expressions that fit the data (within an error bound), we set the objective to minimize the description complexity of the expression. Consequently, the MINLP formulation has the form

$$\begin{aligned}
 \min \quad & \mathcal{C}(f_{cqwz}) && \text{(Complexity)} \\
 \text{s.t.} \quad & (c, q, w, z) \in \mathcal{T} && \text{(Grammar)} \\
 & v_i = f_{cqwz}(x^i), \quad \forall i \in I && \text{(Prediction)} \\
 & \mathcal{D}(v, y) \leq \epsilon && \text{(Error)}
 \end{aligned}$$

where c, q, w and z are decision variables, f_{cqwz} is the expression tree defined by these variables, \mathcal{T} is the set of values of the decision variables that define valid expression trees of bounded size, \mathcal{C} measures the description complexity of the expression tree, and, \mathcal{D} measures error of the predicted values v and (x, y) are the dataset observations. In practice, MINLP solvers (e.g., BARON [TS05], COUENNE[BLL⁺09], SCIP[MFG⁺17]) employ various convex relaxation schemes to obtain lower bounds on the objective function value and use these bounds in divide-and-conquer strategies to obtain *globally optimal solutions* [TS05].

We next give the details of our formulation. The input to the formulation consists of a rooted binary tree, a set of candidate operators and an observed dataset. For each observation $i \in I$, we denote the value of the independent variables by $x^i \in \mathbb{R}^m$ and the dependent variable values by $y^i \in \mathbb{R}$. Let $D = \{1, \dots, m\}$ denote the indices of the independent variables. Let $T = (N, E)$ be the input binary tree, let $s \in N$ denote root node and $L \subset N$ denote the leaf nodes. Each node $n \in N \setminus \{s\}$ in the tree has exactly one predecessor and each node $n \in N \setminus L$ has exactly two successors. Let $O = U \cup B$ denote the set of candidate operators where U contains the unary operators (such as squareroot) and B contains the binary operators (such as addition). Our formulation has two parts: The first part consists of constraints and decision variables that construct the expression tree and the second part is used to evaluate the difference between the estimated and actual label of each observation using the constructed expression tree.

Grammar. The first part of the formulation chooses a subtree of T and assigns either an operator, a constant, or an input variable to each node of the chosen subtree in a consistent way. For each $n \in N$, we have a decision variable $u_n \in \{0, 1\}$ to indicate if the node is used (active) in the expression tree. We also have a decision variable $z_{n,o} \in \{0, 1\}$ to denote if the operator $o \in O$ is assigned to the node, and variables $w_{n,d} \in \{0, 1\}$ to denote if the variable $d \in D$ is assigned to the node. Finally we have a decision variable $q_n \in \{0, 1\}$ if a constant value is assigned the node. The constraint

$$q_n + \sum_{d \in D} w_{n,d} + \sum_{o \in O} z_{n,o} = u_n, \quad \forall n \in N, \quad (1)$$

enforces that each active node is assigned an operator, a constant, or one of the variables. Let $n \in N \setminus L$ be a non-leaf node with successors $l, r \in N$ where l has a smaller index than r . Note that

if node n is assigned a binary operator then both nodes l, r have to be active in the tree; and if n is assigned a unary operator then exactly one of the nodes (say l) has to be active. Consequently, we have constraints $u_r = \sum_{o \in B} z_{n,o}$ and $u_l = \sum_{o \in B \cup U} z_{n,o}$. Notice that these constraints also enforce that if a node is inactive, or is assigned a constant or is a variable, then its successor nodes cannot be active. As the leaf nodes cannot be assigned operators, we also set $z_{n,o} = 0$ for all $o \in O, l \in L$. Finally, we define a continuous variable c_n for $n \in N$ to denote the value of the constant when $q_n = 1$. It is possible to show that any $c \in \mathbb{R}^{|N|}$ together with binary vectors q, w, z , and u (of appropriate dimensions) that satisfy the constraints above define an expression tree and vice versa.

In addition to the basic model described above, we impose some additional constraints to improve computational performance (without compromising optimality). For example we make sure that if a node is assigned a binary operator, then both its successor nodes cannot be assigned constants at the same time. Similarly the left successor of a unary operator cannot be a constant. We also have several constraints to deal with symmetry as expression trees are inherently symmetric objects in the sense that the same function can be represented with different trees, for example by flipping two branches succeeding a commutative binary operator. To avoid numerical problems, we also set bounds on the absolute value of the continuous variables.

Prediction. Using the first set of variables, the second part of the formulation computes the value of the function (defined by the expression tree) for each observation. More precisely, we define a variable $v_{n,i}$ to denote the value of node $n \in N$ for observation $i \in I$. Value of a node clearly depends on the operator assigned to it and the values of its successor nodes. For example, if the addition operator $a \in B$ assigned to node $n \in N$ that has successor nodes $r, l \in N$, then $v_{n,i}$ would be the sum of $v_{r,i}$ and $v_{l,i}$. Therefore we have a constraint:

$$v_{n,i} = q_n c_n + \sum_{d \in D} w_{n,d} x_d^i + \sum_{o \in B} z_{n,o} f_o(v_{r,i}, v_{l,i}) + \sum_{o \in U} z_{n,o} f_o(v_{l,i}) \quad \forall i \in I, n \in N, \quad (2)$$

where $f_o(\cdot)$ denotes the function on the argument(s) of the operator $o \in O$. Also note that $v_{n,i} \neq 0$ only if $u_n = 1$. Leaf nodes cannot be assigned operators, therefore the last two summations in (2) are zero for these nodes.

Error. Recall that $s \in N$ denotes the root node of the expression tree and therefore $v_{s,i}$ corresponds to the estimated output variable for observation $i \in I$. We define the error for observation $i \in I$ as $(v_{s,i} - y^i)$ and relative error as $(v_{s,i}/y^i - 1)$. Using these expressions, we can define a discrepancy function such as $\mathcal{D}(v_s, y) = \sum_{i \in I} (v_{s,i} - y^i)^2$. In our computational experiments, we control the relative error by a constraint added to the model $\sum_{i \in I} (v_{s,i}/y^i - 1)^2 \leq \epsilon$ for a given $\epsilon > 0$.

Complexity. We can define a function to measure description complexity of the model in various ways, for example $\mathcal{C}(f_{cqwz}) = \sum_{n \in N} \sum_{o \in O} c_o z_{n,o}$, where $c_o \in \mathbb{R}$ is a complexity weight assigned to operator $o \in O$ by the user. In our computational experiments, we use the expression $\sum_{n \in N} u_n$ to minimize the total number of nodes in the expression tree.

3 Computational Experiments

We report our computational experience on the discovery of physical laws using real-world datasets. We solve the optimization problem described in Section 2 with the MINLP solver BARON [TS05] v17.8.9, using IBM ILOG CPLEX and IPOPT [WB06] as subsolvers. We can only use operators that are handled by BARON. Furthermore, when multiple globally optimal solution exists, we do not have control over which one is returned. Experiments are run on the cloud, so CPU speed cannot be precisely quantified.

Exoplanet Data. The first set of experiments concerns the discovery of Kepler's law on planetary motion on a dataset taken from NASA [NAS17] reporting information on planetary systems. The dataset consists of tuples of the form $(\text{planet}, \text{star}, \tau, M, m, d)$ where τ is the orbital period of the planet, M the mass of the star, m the mass of the planet, d the major semi-axis of the orbit. Planet features are normalized to Earth's, star features to Sun's. τ is further normalized since its range is very large. We build three test problems from this dataset, including the systems listed in brackets: EP1 (Sol and Trappist), EP2 (HD 219134 and Kepler), EP3 (GJ 667C, HD 147018, HD 154857, HD 159243 and HD 159868).

Table 1: Results for the exoplanet dataset. Time limit is 6 hours. Grey cells exceed the time limit. c represents constants (we do not report its value for brevity).

Dataset	Maximum relative error					
	2%	5%	10%	20%	30%	50%
EP1	$\sqrt[3]{c\tau^2 M}$	$\sqrt[3]{c\tau^2 M}$	$\sqrt[3]{\tau^2(M+c)}$	$\sqrt[3]{\tau(\sqrt{\tau}+M)}$	$\sqrt{c\tau^2}$	$\sqrt{\tau M}$
EP2	$\sqrt[3]{c\tau^2 M}$	$c\sqrt[3]{\tau^2}$	$c\sqrt[3]{\tau^2}$	$\sqrt[3]{c\tau^2}$	$\sqrt{\tau}$	$\sqrt{\tau}$
EP3	$\sqrt[3]{c\tau^2 M}$	$\sqrt[3]{c\tau^2 M}$	$\sqrt{\tau M} + \tau$	$\sqrt[3]{c\tau^2 M}$	$\sqrt{c\tau} + c$	$\sqrt{\tau}$

Table 2: Results for the pendulum dataset. Time limit is 6 hours. All formulas are certified optimal.

Dataset	Maximum relative error			
	0.5%	1%	2%	5%
DS1	$j(\ell - c) - \ell$	$j\sqrt{\ell} + c$	$j\sqrt{\ell}$	cj
DS2	$j\sqrt{\ell}$	$j\sqrt{\ell}$	cj	cj
DS3	$j\sqrt{\ell}$	$j\sqrt{\ell}$	$\ell - cj$	cj
DS4	$j(\ell - c) - c$	$j\sqrt{\ell}$	cj	cj

We consider the set of operators $+, *, \sqrt{ }, \sqrt[3]{ }$, and full binary expression trees with depth at most 3 (the root has depth 0) in which each operator can appear at most 3 times. We aim to predict d in terms of the other input variables. Results are reported in Table 1, for different upper bounds applied to the relative model error. For small model errors, we always find a refined expression of Kepler’s third law: $d = \sqrt[3]{c\tau^2 M}$. Given the data, the more comprehensive formula would be $d = \sqrt[3]{c\tau^2(M+m)}$, but the dependency on m is not picked up because m is negligible compared to M . As model error upper bound increases, we find simpler formulas that do not accurately reflect Kepler’s third law.

Most of the formulas returned are certified globally optimal in a relatively short period of time: on average, 50 minutes for the instances solved to global optimality (two instances hit the time limit and thus are not certified globally optimal within the time limit), standard deviation 75. The number of nodes explored by the Branch-and-Bound algorithm varies greatly: from 147 for the simplest formula $\sqrt{\tau}$, to over 600k for one instance that is not solved to optimality within the allotted runtime. The geometric mean is 19917. By contrast, the number of binary expression trees of depth 3 with 4 candidate operations, 3 input variables, and numerical constants is $\approx 2 \cdot 10^7$ without accounting for symmetry. We remark that the number of nodes explored by Branch-and-Bound does not correspond to the number of candidate expression trees that have been examined: it is merely a measure of difficulty of the search.

Pendulum Data. The second set of experiments concerns ten pendulums with different link lengths, and timestamps for the time at which each pendulum crosses the midpoint of its arc. The data is obtained from a high-resolution video recording of the pendulums, yielding hundreds of timestamps. From the raw data, we randomly extract ten tuples (one for each link length) of the form (ℓ, i, t_i, j, t_j) , where ℓ is the link length, i and j are integers, t_i and t_j are the timestamps for the i -th and j -th crossing of the midpoint, and $j > i$. We aim to predict t_j .

The period of the pendulum is given by $\tau = 2\pi\sqrt{\ell/g}$. Hence, we aim to obtain the equation $t_j = \pi j \sqrt{\ell/g}$. The experimental setup is similar to the exoplanet dataset, but we add “ $-$ ” to the list of operators. Many experiments return the correct formula except the multiplicative constant, since $\pi/\sqrt{g} \approx 1.002$ and its discrepancy from 1 is too small to be significant, up to experimental error. With error bound 0.5%, we observe overfitting on DS1 and DS4: the simpler formula $j\sqrt{\ell}$ does not satisfy the error bound due to experimental error. Because the link lengths of the ten pendulum are close to each other (ranging from $0.28m$ to $0.307m$), for larger error bounds a simpler expression of the form cj suffices. The average runtime is 2 minutes, standard deviation 2.3. The number of nodes ranges from 57 to 8134, geometric mean 299. Despite the presence of redundant variables, the algorithm quickly identifies the simplest formula explaining the data.

In this extended abstract we focused on small datasets and did not address scalability, which is a known limitation of Branch-and-Bound-based methods for MINLP. This is left for future research.

References

- [AB00] Douglas Adriano Augusto and Helio JC Barbosa. Symbolic regression via genetic programming. In *Neural Networks, 2000. Proceedings. Sixth Brazilian Symposium on*, pages 173–178. IEEE, 2000.
- [BLL⁺09] P. Belotti, J. Lee, L. Liberti, F. Margot, and A. Wächter. Branching and bounds tightening techniques for non-convex MINLP. *Optimization Methods and Software*, 24(4-5):597–634, 2009.
- [BNKF98] Wolfgang Banzhaf, Peter Nordin, Robert E Keller, and Frank D Francone. *Genetic programming: an introduction*, volume 1. Morgan Kaufmann San Francisco, 1998.
- [BPK16] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.
- [CL17] Zonglin Jiang Changtong Luo, Chen Chen. A divide and conquer method for symbolic regression. *arXiv:1705.08061v2*, 2017.
- [CT77] JW Connor and JB Taylor. Scaling laws for plasma confinement. *Nuclear Fusion*, 17(5):1047, 1977.
- [CT06] Thomas M Cover and Joy A Thomas. Introduction and preview. *Elements of information theory*, pages 1–11, 2006.
- [DSW99] JW Davidson, D Savic, and GA Walters. Method for the identification of explicit polynomial formulae for the friction in turbulent pipe flow. *Journal of Hydroinformatics*, 1(2):115–126, 1999.
- [DSW03] JW Davidson, Dragan A Savic, and Godfrey A Walters. Symbolic and numerical regression: experiments and applications. *Information Sciences*, 150(1):95–117, 2003.
- [GML96] Murray Gell-Mann and Seth Lloyd. Information measures, effective complexity, and total information. *Complexity*, 2(1):44–52, 1996.
- [HLAN15] Lior Horesh, Leo S Liberti, Haim Avron, and David Nahamoo. Globally convergent system and method for automated model discovery, June 30 2015. US Patent App. 14/755,942.
- [KHS17] Paul Kroll, Alexandra Hofer, Ines V Stelzer, and Christoph Herwig. Workflow to set up substantial target-oriented mechanistic process models in bioprocess engineering. *Process Biochemistry*, 2017.
- [Koz92] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, 1992.
- [Lan81] Pat Langley. Data-driven discovery of physical laws. *Cognitive Science*, 5(1):31–54, 1981.
- [MFG⁺17] Stephen J. Maher, Tobias Fischer, Tristan Gally, Gerald Gamrath, Ambros Gleixner, Robert Lion Gottwald, Gregor Hendel, Thorsten Koch, Marco E. Lübbecke, Matthias Miltenberger, Benjamin Müller, Marc E. Pfetsch, Christian Puchert, Daniel Rehfeldt, Sebastian Schenker, Robert Schwarz, Felipe Serrano, Yuji Shinano, Dieter Weninger, Jonas T. Witt, and Jakob Witzig. The scip optimization suite 4.0. Technical Report 17-12, ZIB, Takustr.7, 14195 Berlin, 2017.
- [NAS17] NASA. NASA exoplanet archive, 2017. <https://exoplanetarchive.ipac.caltech.edu/docs/intro.html>.
- [Pie12] John R Pierce. *An introduction to information theory: symbols, signals and noise*. Courier Corporation, 2012.
- [PL87] H. A. Simon P Langley, G. L. Bradshaw. Heuristics for empirical discovery. In L. Bolc, editor, *Computational models of learning*. Springer-Verlag, Berlin, 1987.
- [R⁺61] Alfréd Rényi et al. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California, 1961.
- [Sah17] N. V. Sahinidis. *BARON 17.8.9: Global Optimization of Mixed-Integer Nonlinear Programs*, User’s Manual, 2017.
- [SL09] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(3):81–85, 2009.
- [SL10] Michael Schmidt and Hod Lipson. Symbolic regression of implicit equations. In *Genetic Programming Theory and Practice VII*, pages 73–85. Springer, 2010.
- [taw] A polyhedral branch-and-cut approach to global optimization.
- [Tod] Ljupco Todorovski. Declarative bias in equation discovery.
- [TS05] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103:225–249, 2005.

- [WB06] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
- [WHM⁺97] M-J Willis, Hugo G Hiden, Peter Marenbach, Ben McKay, and Gary A Montague. Genetic programming: An introduction and survey of applications. In *Genetic Algorithms in Engineering Systems: Innovations and Applications, 1997. GALESIA 97. Second International Conference On (Conf. Publ. No. 446)*, pages 314–319. IET, 1997.