Delft University of Technology

Symbolic regression driven by training data and prior knowledge

Kubalik, Jia; Derner, Erik; Babuska, Robert

DOI

[10.1145/3377930.3390152](https://doi.org/10.1145/3377930.3390152)

Publication date

2020

Document Version

Final published version

Published in

Proceedings of the 2020 Genetic and Evolutionary Computation Conference, GECCO 2020

Citation (APA)

Kubalik, J., Derner, E., & Babuska, R. (2020). Symbolic regression driven by training data and prior knowledge. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference, GECCO 2020* (pp. 958-966). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3377930.3390152>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Symbolic Regression Driven by Training Data and Prior Knowledge

Jiří Kubalík

Czech Institute of Informatics,
Robotics and Cybernetics
Czech Technical University in Prague
16000 Prague, Czech Republic
jiri.kubalik@cvut.cz

Erik Derner

Czech Institute of Informatics,
Robotics and Cybernetics
and
Department of Control Engineering
Faculty of Electrical Engineering
Czech Technical University in Prague
16000 Prague, Czech Republic
erik.derner@cvut.cz

Robert Babuška

Cognitive Robotics
Delft University of Technology
Delft, 2628 CD, The Netherlands
and
Czech Institute of Informatics,
Robotics and Cybernetics
Czech Technical University in Prague
16000 Prague, Czech Republic
r.babuska@tudelft.nl

ABSTRACT

In symbolic regression, the search for analytic models is typically driven purely by the prediction error observed on the training data samples. However, when the data samples do not sufficiently cover the input space, the prediction error does not provide sufficient guidance toward desired models. Standard symbolic regression techniques then yield models that are partially incorrect, for instance, in terms of their steady-state characteristics or local behavior. If these properties were considered already during the search process, more accurate and relevant models could be produced. We propose a multi-objective symbolic regression approach that is driven by both the training data and the prior knowledge of the properties the desired model should manifest. The properties given in the form of formal constraints are internally represented by a set of discrete data samples on which candidate models are exactly checked. The proposed approach was experimentally evaluated on three test problems with results clearly demonstrating its capability to evolve realistic models that fit the training data well while complying with the prior knowledge of the desired model characteristics at the same time. It outperforms standard symbolic regression by several orders of magnitude in terms of the mean squared deviation from a reference model.

CCS CONCEPTS

• **Computing methodologies** → Genetic algorithms; Model development and analysis; • **Theory of computation** → Design and analysis of algorithms; Genetic programming; • **Applied computing** → Multi-criterion optimization and decision-making;

KEYWORDS

symbolic regression, genetic programming, multi-objective optimization, model learning

ACM Reference Format:

Jiří Kubalík, Erik Derner, and Robert Babuška. 2020. Symbolic Regression Driven by Training Data and Prior Knowledge. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancún, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390152>

1 INTRODUCTION

Many model-learning approaches have been described in the literature: time-varying linear models [20, 22], Gaussian processes and other probabilistic models [7, 11], deep neural networks [9, 21] or local linear regression [14]. All these approaches suffer from drawbacks induced by the use of the specific approximation technique, such as a large number of parameters (deep neural networks), local nature of the approximator (local linear regression), computational complexity (Gaussian process), etc. Symbolic regression (SR) is an approach that generates models in the form of analytic equations that can be constructed by using even very small training data sets. SR has been used in nonlinear data-driven modeling with quite impressive results [1, 2, 12, 13, 23, 25, 26].

In standard SR, the search for analytic models is driven purely by the prediction error observed on the training data samples. However, the training data may not provide a sufficient guidance towards desired models, for instance, when the data set does not sufficiently cover the input space or even when some parts of the input space are completely omitted in the data set. SR techniques then yield models that are partially incorrect, for instance, in terms of their steady-state characteristics or local behavior. On the other hand, some information about the desired properties of the modelled system is often available. If these properties were considered already during the search process, more accurate and relevant models could be produced.

There are very few SR approaches in the literature that take into account information about the model sought other than just the minimum training error. Perhaps the most promising and the most relevant is the Counterexample-Driven Symbolic Regression [6], where Counterexample-Driven Genetic Programming [17] is used to synthesize regression models that not only comply with

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '20, July 8–12, 2020, Cancún, Mexico

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7128-5/20/07...\$15.00

<https://doi.org/10.1145/3377930.3390152>

the training data set, but also meet formal constraints imposed on the model. A Satisfiability Modulo Theories (SMT) solver is used to verify whether a given model meets the formal specification. However, this method suffers from several deficiencies. Queries to the SMT solver are computationally very costly. Only a limited function set $\{+, -, *, /\}$ can be used since transcendental functions are not fully supported by contemporary SMT solvers. There are also SMT solvers that can handle functions like *sine* and *log*, however, they are very computationally expensive for larger models. Moreover, general nonlinear inequalities over the real numbers represent a non-decidable problem, therefore, delta-decidability is used instead. This means, one has to supply the solvers with a proper value of the delta precision parameter. In the end, the solver can return both false positive and false negative answers.

In this paper, we propose a multi-objective SR approach that is driven by the training data as well as by the prior knowledge on the desired properties the model should exhibit. Various types of constraints can be used such as the monotonicity of the model's output on a given interval, odd symmetry of the model, symmetry w.r.t. the input variables, steady-state characteristics of the model, etc. The properties, given in the form of formal constraints, are internally represented by a set of discrete constraint samples on which the validity of candidate models is checked. Both aspects of model performance are treated with equal importance. Consequently, the method produces models that fit the training data as well as possible while complying with the prior knowledge of the desired model characteristics at the same time.

We use a variant of Single Node Genetic Programming (SNGP) [16, 19] that generates models in the form of a linear combination of possibly nonlinear features. It has been shown that SR methods producing such compound models outperform SR methods generating single tree models [3, 4, 18, 24]. In the standard version of the SNGP algorithm, the coefficients of the linear model are estimated using least squares. In the proposed multi-objective SNGP, this may not be the most efficient way since the estimation of the coefficients is biased in the direction of models fitting well the training data. So, we propose an alternative way to derive the coefficients using a multi-objective local search procedure.

The paper is organized as follows. Section 2 defines the problem. Section 3 describes the proposed method. In Section 4, the method is experimentally evaluated on three test problems. Section 5 concludes the paper and suggests topics for future research.

2 PROBLEM DEFINITION

We solve the problem of constructing an optimal analytic model $f: \mathcal{X} \rightarrow \mathbb{R}$ operating in the input space $\mathcal{X} \subset \mathbb{R}^n$ given two optimization goals:

- The model fits the training data as accurately as possible.
- The model is consistent with the constraints imposed on the model that capture the desired model's properties.

Two data sets are used to search for the model, *standard training data set* and *constraint data set*.

Standard training data set $D = \{d_1, \dots, d_m\}$. Each training sample d_i is a tuple

$$d_i = \langle \mathbf{x}_i, y_i \rangle, \quad (1)$$

where $\mathbf{x}_i \in \mathcal{X}$ is a particular input vector and $y_i \in \mathbb{R}$ is the corresponding desired target value.

Constraints and constraint data set. We assume that all constraints can be written as nonlinear inequality and equality constraints. Inequality constraints are:

$$g_i^f(\mathbf{x}) \leq 0, \quad i = 1, \dots, p \quad (2)$$

where function g_i has a specific form for each particular type of inequality constraint and in general may have more arguments: $g_i(x_1, x_2, \dots)$. For instance, to specify a monotonically increasing function, we can define $g_i^f(x_1, x_2) = f(x_1) - f(x_2)$ and then, when checking whether the constraint is satisfied, evaluate it for any pair of data points $x_1 \leq x_2, x_1, x_2 \in \mathcal{X}$. In the sequel, to avoid notational clutter we will write g_i with a single argument as in (2). Inequality constraint violation for model f is calculated as follows:

$$E_g^f = \sum_{i=1}^p \sum_{\forall \mathbf{x}_\ell \in X_i^g} (\max(g_i^f(\mathbf{x}_\ell), 0))^2 \quad (3)$$

where X_i^g is a set of data points on which the violation of constraint g_i is calculated. Analogously, equality constraints have the form:

$$h_j^f(\mathbf{x}) = 0, \quad j = 1, \dots, q \quad (4)$$

where function h_j is specific to the particular type of equality constraint. Also this function may have more arguments. Equality constraint violation for model f is calculated as follows:

$$E_h^f = \sum_{j=1}^q \sum_{\forall \mathbf{x}_\ell \in X_j^h} (h_j^f(\mathbf{x}_\ell))^2 \quad (5)$$

The constraint data set is given by $C = X^g \cup X^h$, with $X^g = \bigcup_{i=1}^p X_i^g$ and $X^h = \bigcup_{j=1}^q X_j^h$.

The two aforementioned optimization goals are formally defined as follows:

- C_t – minimize the mean-squared error calculated for model f on the training data set D

$$C_t = \frac{1}{m} \sum_{i=1}^m (f(\mathbf{d}_i) - y_i)^2. \quad (6)$$

- C_c – minimize the mean-squared error calculated for model f on the constraint data set C

$$C_c = \frac{E_g^f + E_h^f}{|C|}. \quad (7)$$

Importantly, both aspects of the model's performance – i.e., its accuracy as well as its formal validity – are treated as equally important through the optimization process.

3 METHOD

In this section, the proposed multi-model SR method based on the SNGP algorithm is described. Firstly, the base SNGP algorithm and its population structure for storing and operating with a single analytic model are briefly described. Then, we introduce an extended population structure that allows for operating with multiple independent models. Finally, the algorithm itself is described with the focus on the multi-objective aspect of the search process.

3.1 Base SNGP

The idea of the proposed multi-objective symbolic regression method is applicable to any population-based approach. Here, we adopt a variant of SNGP [16] and particularly the variant proposed in [19]. In the following text we use the term “*base SNGP*” to refer to this algorithm.

Standard SNGP is a tree-based genetic programming (GP) technique that evolves a population of individuals, i.e. program nodes, organized in an ordered linear array structure. The nodes are interconnected in the left-to-right manner, meaning that a node can act as an input operand only of those nodes which are positioned to its right in the population. Thus, the population of nodes represents a whole set of tree-based programs rooted in its individual nodes. In the context of SR, the population starts with constant nodes and variables followed by general function nodes chosen from a set $\mathcal{F}$ of elementary functions defined by the user for the problem at hand, see Figure 1a. The expression trees rooted in function nodes provide a capacity to represent complex and possibly non-linear analytic functions. The population is evolved through a first-improvement iterative local search procedure using a mutation operator that varies the input links of the function nodes.

An important property of the *base SNGP* is that it evolves linear-in-parameters nonlinear analytic models of the form

$$f(x) = \beta_0 + \sum_{i=1}^{n_f} \beta_i \varphi_i(x) \quad (8)$$

where the nonlinear functions $\varphi_i(x)$ are *features* constructed by means of GP operations using a predefined set of elementary functions $\mathcal{F}$. The coefficients β_i are not evolved using genetic operators. Instead, they are estimated using some multiple regression technique, e.g. the least squares one. Importantly, the whole population represents a single analytic model whose features, $\varphi_i(x)$, are rooted in so-called *identity nodes*, where each identity node just refers to some non-constant-output node in the population, see Figure 1a.

The complexity of evolved analytic models is constrained by two user-defined parameters: n_f is the maximum number of features the analytic model can be composed of, and δ is the maximal depth of the feature’s tree representation.

We chose for this variant of GP since it has recently been shown in [3, 4, 18, 24] that GP methods evolving this kind of compound regression models outperform conventional GP evolving a single-tree structure representing the whole model. In particular, the *base SNGP* has been successfully used for several SR tasks from the reinforcement learning and robotics domains [1, 2, 12, 13]. A detailed description of the *base SNGP* is beyond the scope of this paper. For more details please refer to [19].

3.2 Multi-objective SNGP

We propose a multi-objective variant of SNGP for the bi-objective SR that simultaneously optimizes both optimization criteria, C_t and C_c . First, we adapt the population architecture to allow for operating with a set of M independent models. For this purpose, we use a set of *base SNGP* populations, each representing a unique model, see Figure 1b. From now on, we will use the term *population* in the sense of the population of models.

Figure 1: (a) Structure of the *base SNGP* population with a set of identity nodes defining features of a single model. (b) Population of models, each represented by a unique *base SNGP* population.

The proposed algorithm is based on the NSGA-II algorithm [10] that uses the following *domination* principle: A solution $\mathbf{x}^{(1)}$ is said to dominate another solution $\mathbf{x}^{(2)}$, if $\mathbf{x}^{(1)}$ is not worse than $\mathbf{x}^{(2)}$ in any objective and $\mathbf{x}^{(1)}$ is strictly better than $\mathbf{x}^{(2)}$ in at least one objective.

The outline of the multi-objective SNGP algorithm is shown in Algorithm 1. It starts with a random initialization of the population of models, *population*. Each model is first assigned its coefficients β . In the original *base SNGP*, these are estimated using least squares. However, this might not be the best choice when solving multi-objective SR as will be discussed later in this section. Complete models are evaluated on both data sets D and C .

The algorithm then iterates through a specified number of generations, lines 7–26. In each generation, an intermediate population of models, *interPop*, is created from models of the current *population*, lines 10–23. First, a parent model is selected from the *population* and its copy is assigned as the initial value to the offspring model, *child*. A standard tournament selection uses the *crowded-comparison operator* [10] to choose parental models to be mutated. The crowded-comparison operator takes two models and returns the one that is from the better non-dominated front or if both are from the same non-dominated front the more unique one is returned. Thus, well-performing and unique models are preferred.

The *child* then undergoes a predefined number of optimization iterations, lines 14–22. In each iteration, the *child* is mutated and its coefficients β are recalculated. The mutated model *temp* is then evaluated and it becomes the *child* for the next iteration if it is not dominated by the current version of *child* nor by the *parent* model, lines 21–22. Final version of the *child* is added to *interPop*.

Once *interPop* has been completed, it is merged with the current *population* resulting in a new version of the *population*. This is done using the NSGA-II replacement strategy that again prefers non-dominated solutions to the dominated ones and among solutions of the same non-dominated front the more unique ones are preferred. For more details refer to [10]. In the end of the generation, the C data set can optionally be updated, see Section 3.4.

Algorithm 1: Multi-objective SNGP algorithm

Input: M ... size of the population of models
 C ... set of constraint samples
 D ... training data set
 $MAXGENS$... maximum number of generations
 $MAXITERS$... maximum number of iterations
carried out to produce an offspring model from the parent one
 $PERIOD$... number of generations between updates
of C
Output: S ... set of final models

```

1 init(population)
2 for  $\forall model \in population$  do
3    $model.calculateBetaCoeffs()$ 
4    $model.calculateC\_train(D)$ 
5    $model.calculateC\_constraint(C)$ 
6  $generation \leftarrow 0$ 
7 while  $generation < MAXGENS$  do
8    $generation \leftarrow generation + 1$ 
9    $interPop \leftarrow \{\}$ 
10  while  $interPop.size() < M$  do
11     $parent \leftarrow selectModel(population)$ 
12     $child \leftarrow parent.clone()$ 
13     $i \leftarrow 0$ 
14    while  $i < MAXITERS$  do
15       $i \leftarrow i + 1$ 
16       $temp \leftarrow child.clone()$ 
17       $temp.applyMutations()$ 
18       $temp.calculateBetaCoeffs()$ 
19       $temp.calculateC\_train(D)$ 
20       $temp.calculateC\_constraint(C)$ 
21      if  $\neg parent.dominates(temp) \wedge$ 
          $!child.dominates(temp)$  then
22         $child \leftarrow temp$ 
23     $interPop.add(child)$ 
24   $population \leftarrow NSGAII\_merge(population, interPop)$ 
25  if  $generation \% PERIOD == 0$  then
26     $C.update()$ 
27  $S \leftarrow population.getNondominatedModels()$ 
28 return  $S$ 

```

Finally, a set of final models is selected as the output of the run. Since this is a multi-objective optimization approach, the population contains a whole set of non-dominated solutions in the end. So, the question is how to choose the best solutions to be returned as the output of the run? Definitely, the *extreme* model with the best value of C_c should be in the final set of solutions. However, this may not necessarily be the most interesting one as it can do poorly in the other objective. We will demonstrate this in Section 4. On the other hand, it is very likely that the model with the best value of C_t does not belong to the most useful ones unless it coincides with the *extreme* model. The rationale for it is that such a model probably over-fits the training data while ignoring the constraints imposed on the model's properties, which is not what we want to get. So, we need to take into consideration also the high-quality trade-off solutions. Here, we take the whole set of non-dominated solutions of the final population.

3.3 Alternative way to estimate coefficients β

As mentioned above, the coefficients β weighting the model's features in (8) are fitted using the least squares method. In particular, the coefficients are found such that the final model minimizes the sum of squared residuals over the training data set D . Thus, just the C_t objective is considered at that moment, the C_c is ignored. This means, there is no pressure towards coefficients that would make the model better in terms of C_c , even if it was attainable with the given set of features. Clearly, this may adversely affect the performance of the whole method. In order to remedy this issue, we propose an alternative way to calculate the coefficients so that both objectives are optimized simultaneously.

We adopt a simple local search method to tune the coefficients β . It initializes the coefficients with values uniformly sampled from interval $(-1, 1)$ and the performance measures of the initial model, C_t and C_c , are calculated. Then, it iterates for the specified number of iterations. In each iteration, the vector β is perturbed by adding values sampled from the normal distribution according to

$$\beta' \leftarrow \beta + N(0, 0.1).$$

The new vector of coefficients β' is accepted if the model using it dominates the model with the current values β . Otherwise, the current vector β remains for the next iteration.

3.4 Constraint Data Set Management

The content of the constraint sample set C is crucial for the success of the method. Initial constraint samples, i.e., the points x of the input space where the constraint will be checked, are drawn randomly with a uniform distribution from the whole input space $\mathcal{X}$. Similarly, new constraint samples that are added to C during the optimization process are generated at random. Note that even such a simple method can be beneficial for the optimization process as the newly added constraint samples can alter the ranking of models in the population. Such an intervention can boost the exploration towards different regions of the search space and can help to prevent the population from stagnating.

Obviously, more sophisticated sampling strategies could improve the performance of this method. However, this is out of the scope of this paper. We leave this for the future research.

4 EXPERIMENTS

Three methods were compared:

- **baseSNGP** – *base SNGP* minimizing only the mean squared error on the training data set as described in Section 3.1,
- **mSNGP** – the proposed multi-objective SNGP using the least squares method to estimate coefficients β ,
- **mSNGP-ls** – the proposed multi-objective SNGP using the local search procedure to estimate coefficients β .

The methods were experimentally evaluated on three problems:

- **resistance2** – This is a test problem originally proposed in [6]. It uses a sparse set of noisy samples derived using the equivalent resistance of two resistors in parallel, $r = r_1 r_2 / (r_1 + r_2)$, denoted as a *reference model*. The goal is to find such a model $f(r_1, r_2)$ that fits the training data and has the same properties as the *reference model*. For the sake of a unified notation, we define $x = (r_1, r_2)$ and $y = r = f(x_1, x_2)$.

- **magman** – The magnetic manipulation system consists of an iron ball moving along a rail and an electromagnet at a static position under the rail. The goal is to find a model of the nonlinear magnetic force affecting the ball, $f(x)$, as a function of the horizontal distance, x , between the iron ball and the activated coil given a constant current through the coil, i . We use data measured on a real system and an empirical model $\hat{f}(x) = -ic_1x/(x^2 + c_2)^3$ proposed in the literature [15] as the *reference model*. Parameters c_1 and c_2 were found empirically for the given system and this model was used to design well-performing nonlinear controllers in [1, 8]. For this example, we define $y = f(x)$.
- **pressure** – In this problem, highly nonlinear pressure dynamics in a laboratory fermenter is modelled. The process under consideration is a 40 l laboratory fermenter which contains 25 l of water. At the bottom of the fermenter, air is fed into the water at a specified flow rate which is kept at a desired value by a local mass-flow controller. The air pressure p in the head space can be controlled by opening or closing an outlet valve u at the top of the fermenter. The goal is to find a dynamic model $p_{k+1} = f(p_k, u_k)$. The exact form of the nonlinear target function is unknown. For more details, please refer to [5]. For this example, we define $x = (p_k, u_k)$ and $y = f(x_1, x_2)$.

We chose these three problems since we possess detailed knowledge of the data and they have a potential to demonstrate advantages of the proposed approach. For all the problems, the training data set D is either very sparse or its samples are unevenly distributed in the input space. Standard symbolic regression has a very small chance of converging to an acceptable model as it is likely to over-fit the data. In addition, it is easy to visualize the models.

Fifty independent runs were carried out with each method on each problem. In order to assess statistical significance of the differences among the algorithms we used the Wilcoxon rank sum test, which rejects the null hypothesis that the two compared sets are sampled from continuous distributions with equal medians at the 1% significance level. In the tables, two cases are highlighted – whether baseSNGP is significantly better or worse than both new methods and whether mSNGP performance is significantly different from that of mSNGP-1s.

The algorithms were tested with the following parameter setting:

- Population size of each *base SNGP* population: 400
- Maximum number of features:
 - resistance2: $n_f = 3$
 - magman, pressure: $n_f = 5$
- Maximum feature's depth:
 - resistance2: $\delta = 5$
 - magman, pressure: $\delta = 7$
- Elementary functions:
 - resistance2: $\mathcal{F} = \{+, -, *, /\}$
 - magman, pressure: $\mathcal{F} = \{+, -, *, \text{square}, \text{cube}, \text{sine}, \text{tanh}\}$
- Population size: $M = 50$
- Tournament size: 3
- Maximum number of generations: $\text{MAXGENS} = 40$
- Maximum number of iterations: $\text{MAXITERS} = 50$
- Number of local search iterations: $\text{LS_ITERS} = 50$

- Number of generations between C updates: $\text{PERIOD} = 2$
- Total number of fitness evaluations in baseSNGP: $M \times \text{MAXGENS} \times \text{MAXITERS} = 10^5$.

For the resistance2 problem, just elementary arithmetic operators are used and rather low-complexity models are allowed. However, such a configuration is not sufficient for more difficult problems like magman and pressure. Thus, a richer set of elementary functions as well as parameters allowing for more complex models were used for these two problems.

4.1 Resistance2

4.1.1 Training data. We use the data set with 10 training samples that was used in [6]. The values of x_1 and x_2 are sampled uniformly from the interval $[0.0001, 20]$, see Figure 2. The variables as well as the target value of each training sample are disturbed with a noise randomly generated with a normal distribution $\mathcal{N}(0, 0.1\sigma_X)$, where X is a given variable. Such a noisy training data set is generated anew for each independent symbolic regression run.

4.1.2 Constraints. We used the following three constraints as defined in [6]:

- symmetry with respect to arguments: $f(x_1, x_2) = f(x_2, x_1)$,
- domain-specific constraint: $x_1 = x_2 \implies f(x_1, x_2) = \frac{x_1}{2}$,
- domain-specific constraint: $f(x_1, x_2) \leq x_1, f(x_1, x_2) \leq x_2$.

The initial set C contained 60 constraint samples.

4.1.3 Performance evaluation. We use a two-phase procedure to evaluate the obtained models. First, a model is checked whether it is “reasonably” close to the *reference model* $r(\cdot)$. For this purpose, a high-resolution grid of 200×200 validation points sampled in the input space $[0.0001, 20]^2$ was generated. The response values of the model on all grid points are calculated. If the maximum absolute deviation, MAD, of the model's response $f(\cdot)$ from the *reference model* $r(\cdot)$ is less than $\epsilon = 0.1\sigma_y$ over all validation points then the model is considered as *acceptable*. From each run of the multi-objective SNGP algorithm, only the acceptable model with the least MAD value is selected to the set of acceptable models.

Then, the following four performance measures are defined:

- succ – the number of runs, which yielded an acceptable model.
- $\text{MSE}_{\text{train}}$ – median C_t value over the set of acceptable models.
- MSE_{ref} – median of the mean squared error between the model's output and the *reference model* on the validation points calculated over the set of acceptable models.
- violation – median C_c value over the set of acceptable models.

4.1.4 Results. The results achieved on this problem are presented in Table 1. The *reference model* and examples of the models obtained are in Figure 2.

4.2 Magman

4.2.1 Training data. The portion of the input space of interest spans over the interval $-0.075 \text{ m} \leq x \leq 0.075 \text{ m}$. However, only its small part, $[-0.027 \text{ m}, 0.027 \text{ m}]$, is covered by the 858 data samples collected for this task, see Figure 3. The data were measured on

Table 1: Comparison on resistance2 problem. Since the baseSNGP algorithm did not find any acceptable model, its $\text{MSE}_{\text{train}}$, MSE_{ref} and violation values are calculated over all fifty resulting models. $\text{MSE}_{\text{train}}$ values shown in the brackets are the median C_t values achieved by the *reference model* on the training data sets used in the runs where the acceptable models were produced. Bold values indicate that baseSNGP is significantly better or worse than both mSNGP and mSNGP-1s.

	baseSNGP	mSNGP	mSNGP-1s
succ	0 / 50	8 / 50	16 / 50
$\text{MSE}_{\text{train}}$	3.6×10^{-3}	0.078 (0.11)	0.14 (0.15)
MSE_{ref}	1.5×10^3	1.7×10^{-3}	1.6×10^{-4}
violation	2.3×10^3	4.5×10^{-4}	1.1×10^{-6}

a real system [8]. The whole data set was split into the training and test data sets, D and D_{test} , in the ratio 7:3. Properties of the model sought outside the sampled interval are specified purely by the additional constraints.

4.2.2 Constraints. The following constraints were defined for the magman problem. The model sought is an odd function that is positive on the interval $[-0.075, 0]$ and negative on the interval $[0, 0.075]$. Furthermore, it is monotonically increasing on the intervals $[-0.075, -0.008]$ and $[0.008, 0.075]$ and monotonically decreasing on the interval $[-0.008, 0.008]$. Finally, we define the exact output value $f(0) = 0$ and two exact output values at the boundary points of the input space as $f(-0.075) = 1 \times 10^{-3}$ and $f(0.075) = -1 \times 10^{-3}$.

The initial set C contained 90 constraint samples. An example of a constraint sample of the decreasing monotonicity constraint is

$$x = \langle x_1, x_2 \rangle, \quad x_1, x_2 \in [-0.01, 0.01] \wedge 0 < x_2 - x_1 < \eta,$$

and the constraint violation is calculated according to (5) as

$$\max(f(x_2) - f(x_1), 0)^2.$$

This means, two distance values, y_1 and y_2 , are sampled from a narrow interval of size η . Here, $\eta = 0.0001$ was used.

4.2.3 Performance evaluation. Similarly to the resistance2 problem, the models are evaluated in two steps. First, we use 30000 validation points evenly sampled from the whole input domain $[-0.075 \text{ m}, 0.075 \text{ m}]$ to check whether the model output lies within a tolerance margin around the *reference model*, see Figure 3. If so, the model is considered acceptable and its mean squared deviation MSD from the *reference model* is calculated. From each run of the multi-objective SNGP algorithm, only the acceptable model with the least MSD value is selected to the set of acceptable models. Then, the succ, $\text{MSE}_{\text{train}}$, MSE_{ref} and violation performance values are calculated in the same way as for resistance2 problem. In addition, a MSE_{test} performance measure is calculated as a median of the mean squared error on D_{test} over the set of acceptable models. Note, the scope of the model's validation on the D_{test} is limited since it applies only to the portion of the input space that was covered by the data.

4.2.4 Results. The results achieved on this problem are presented in Table 2. Figure 3 shows an example of a trivial model that

perfectly satisfies the constraints, but poorly fits the training data and an example of an overall well-performing model.

Figure 2: Models for resistance2 problem. (a) The *reference model* $r(r_1, r_2) = r_1 r_2 / (r_1 + r_2)$, from which the training data were sampled. (b) An example of a model evolved with baseSNGP that perfectly fits D , but violates the physical law. (c) An example of a high-quality model evolved with mSNGP-1s.

Table 2: Comparison on magman problem. The values presented in column baseSNGP are calculated over all fifty resulting models. Bold values in the first column indicate that baseSNGP is significantly better or worse than both mSNGP and mSNGP-1s. Bold values in the second and third column indicate that the respective method is significantly better than the other proposed method.

	baseSNGP	mSNGP	mSNGP-1s
succ	0 / 50	15 / 50	27 / 50
MSE _{train}	2.78×10^{-3}	2.80×10^{-3}	2.84×10^{-3}
MSE _{test}	3.14×10^{-3}	3.14×10^{-3}	3.22×10^{-3}
MSE _{ref}	12.8	1.38×10^{-4}	1.03×10^{-4}
violation	5.6	7.3×10^{-6}	2.7×10^{-9}

Figure 3: Models for magman problem. (a) A trivial model perfectly satisfying constraints, but poorly fitting the training data. (b) An example of well-performing model.

4.3 Pressure

4.3.1 Training data. A set of 756 data samples unevenly distributed in the input space $[1, 2] \times [0, 100]$ were measured on the real system, see Figure 4. The data were split into the training and test data sets, D and D_{test} , in the ratio 7:3.

Table 3: Comparison on pressure problem. The bold value in the first column indicates that baseSNGP is significantly better than both mSNGP and mSNGP-1s. Bold values in the second column indicate that mSNGP is significantly better than mSNGP-1s.

	baseSNGP	mSNGP	mSNGP-1s
succ	13 / 50	42 / 50	49 / 50
MSE _{train}	4.43×10^{-6}	4.51×10^{-6}	7.28×10^{-6}
MSE _{test}	8.83×10^{-6}	7.83×10^{-6}	1.28×10^{-5}

4.3.2 Constraints. Two types of constraints were defined for this problem. The model sought is monotonically increasing w.r.t. both inputs on the whole input space domain. The model's output is bounded in the interval $[1, 2.2]$. The initial set C contained 80 constraint samples.

4.3.3 Performance evaluation. In this case, we use a grid of 200×200 validation points evenly sampled from the input space $[1, 2] \times [0, 100]$. Models that have zero violation on all of those validation points are considered acceptable. For each run, the model with the least C_t is chosen. Only succ, MSE_{train} and MSE_{test} values are calculated since violation is by definition zero for all acceptable models and MSE_{ref} can not be calculated since we do not have any reference model, neither theoretical nor empirical one.

4.3.4 Results. The results achieved on this problem are presented in Table 3. Figure 4 shows an example of a trivial model that perfectly satisfies the constraints, but poorly fits the training data and an example of an overall well-performing model.

4.4 Discussion

4.4.1 resistance2. baseSNGP was not able to find any acceptable model in the fifty independent runs. The models it produces fit the training data very well, but they largely violate the constraints, as illustrated in Figure 2b. The proposed multi-objective method already finds acceptable models, an example is in Figure 2c. However, none of the acceptable models found has zero violation, none of them is an identical version of the *reference model*. Our hypothesis is that SNGP searches the space of excessively complex models and misses the simpler ones. Among these models, it is hard to find those that perfectly satisfy the equality constraints. This issue can be resolved adjusting the algorithm to search for models of varying complexity. We leave this for the future work.

In addition, both mSNGP and mSNGP-1s outperform baseSNGP by six orders of magnitude in terms of the MSE_{ref}. Interestingly, both mSNGP and mSNGP-1s, achieve better C_t than the *reference model* as illustrated for one particular run in Figures 2a and 2c. In 14 out of the 16 cases mSNGP-1s found an acceptable model with the C_t value better than that of the *reference model*.

An important observation is that mSNGP-1s outperforms mSNGP in terms of all four performance indicators, though the p-values ranging from 0.022 to 0.071 returned by the rank sum test do not suggest that the observed differences are statistically significant. This can be attributed to very small size of tested samples, 8 and 16.

Figure 4: Models for pressure problem. (a) Trivial model perfectly satisfying constraints, but poorly fitting the training data. (b) An example of an overall well-performing model.

4.4.2 magman. The baseSNGP method finds the most precise models w.r.t. C_t , but they are effectively useless due to large constraint violation. It generated no acceptable model. The proposed methods can find acceptable models, mSNGP-ls is better than mSNGP. Both mSNGP and mSNGP-ls outperform baseSNGP by five orders of magnitude in terms of the MSE_{ref} . Similarly to *resistance2*, it is hard to find a model perfectly satisfying the constraints defined for this problem. Again, the small constraint violations of the acceptable models are due to the equality constraints, which is no problem, as the purpose of these constraints is to force the function asymptotically approach zero, rather than attain an exact value. Interestingly, the acceptable models are very close to the empirical one in the regions of the input space where only the constraints were specified, see Figure 3b.

As mentioned in Section 3.2, all non-dominated models are considered as the output of the mSNGP and mSNGP-ls run. The reason is that the models with the best C_t or C_c are often useless. A typical example of the model with zero constraint violation is in Figure 3a. Clearly, the model is bad as it does not fit the training data well.

4.4.3 pressure. This is the only problem where the baseSNGP method succeeded in finding acceptable models. Again, both variants of the proposed approach find an acceptable model significantly

more frequently than baseSNGP, with mSNGP-ls having higher success rate than mSNGP.

4.4.4 Constraint samples. One thing that still remains an open issue is that even when a model successfully passes all constraint samples, the model may not necessarily be valid on the whole constraint domain. To increase the efficacy of the validity checks, the right constraint samples, i.e., the most informative ones, should be used during the whole run. In general, the ability to continuously generate arbitrary constraint samples has a great potential to direct the search to better models on the fly and it is something that should effectively be utilized. Various strategies can be used to attain this goal. For example, new samples of a given constraint can be generated in the vicinity of the sample that has been found hard to be satisfied by models in the current population. Similarly, a set of candidate constraint samples can be randomly generated and the one with the highest fail ratio over the models in the current population is chosen. This is in accordance with the observation presented in [6] that counterexamples are more beneficial for evolving correct models than just random samples. New samples can also be generated so that the coverage of the constraint domain increases as much as possible.

5 CONCLUSIONS

We proposed a new multi-objective symbolic regression method where both aspects of the model performance, the model's accuracy as well as its formal validity, are treated equally. This is a general approach that is applicable whenever some information about the desired properties of the modelled system in the form of explicit samples is available. The results achieved through experiments on three test problems clearly demonstrate its capability to evolve realistic models that fit well the training data while complying with the prior knowledge of the desired model characteristics at the same time. We also proposed an alternative method for estimating coefficients of the linear model. This simple yet effective local search method proved to be better than the least squares method.

An advantage of the proposed method over the validation methods based on the use of SMT solvers is that it checks the model validity on discrete samples, which is fast (even for large models) and exact. Moreover, arbitrary functions can be used to build the models. However, the selection of constraint samples is an open issue. We will investigate various strategies to maintain the most relevant constraint samples during the whole run. Here, we intend to take an inspiration from the field of active learning.

When evaluating the validity of a candidate model, its cumulative constraint violation is calculated over all constraint samples. Note, the violations can be of a very different scale for the different constraints. Consequently, some constraints can dominate the others within the constraint violation objective. The normalization of the constraint violations is another future research line.

6 ACKNOWLEDGEMENTS

This work was supported by the European Regional Development Fund under the project Robotics for Industry 4.0 (reg. no. CZ.02.1.01/0.0/0.0/15_003/0000470).

REFERENCES

- [1] ALIBEKOV, E., KUBALÍK, J., AND BABUSKA, R. Policy derivation methods for critic-only reinforcement learning in continuous spaces. *Eng. Appl. of AI* 69 (2018), 178–187.
- [2] ALIBEKOV, E., KUBALÍK, J., AND BABUŠKA, R. Symbolic method for deriving policy in reinforcement learning. In *2016 IEEE 55th Conference on Decision and Control (CDC)* (Dec 2016), pp. 2789–2795.
- [3] ARNALDO, I., KRAWIEC, K., AND O'REILLY, U.-M. Multiple regression genetic programming. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation* (New York, NY, USA, 2014), GECCO '14, Association for Computing Machinery, p. 879–886.
- [4] ARNALDO, I., O'REILLY, U.-M., AND VEERAMACHANENI, K. Building predictive models via feature synthesis. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation* (New York, NY, USA, 2015), GECCO '15, Association for Computing Machinery, p. 983–990.
- [5] BABUŠKA, R. *Fuzzy Modeling for Control*. Kluwer Academic Publishers, Boston, USA, 1998.
- [6] BŁĄDEK, I., AND KRAWIEC, K. Solving symbolic regression problems with formal constraints. In *Proceedings of the Genetic and Evolutionary Computation Conference* (New York, NY, USA, 2019), GECCO '19, ACM, pp. 977–984.
- [7] BOEDECKER, J., SPRINGENBERG, J. T., WÜLFING, J., AND RIEDMILLER, M. Approximate real-time optimal control based on sparse gaussian process models. In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)* (Dec 2014), pp. 1–8.
- [8] DAMSTEEG, J., NAGESHRAO, S., AND BABUŠKA, R. Model-based real-time control of a magnetic manipulator system. In *Proceedings 56th IEEE Conference on Decision and Control (CDC)* (Melbourne, Australia, Dec. 2017), pp. 3277–3282.
- [9] DE BRUIN, T., KOBER, J., TUYLS, K., AND BABUŠKA, R. Integrating state representation learning into deep reinforcement learning. *IEEE Robotics and Automation Letters* 3, 3 (July 2018), 1394–1401.
- [10] DEB, K., PRATAP, A., AGARWAL, S., AND MEYARIVAN, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (April 2002), 182–197.
- [11] DEISENROTH, M. P., AND RASMUSSEN, C. E. PILCO: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011* (2011), pp. 465–472.
- [12] DERNER, E., KUBALÍK, J., AND BABUSKA, R. Data-driven construction of symbolic process models for reinforcement learning. In *Proceedings IEEE International Conference on Robotics and Automation (ICRA)* (Brisbane, Australia, May 2018), pp. 5105–5112.
- [13] DERNER, E., KUBALÍK, J., AND BABUSKA, R. Reinforcement learning with symbolic input-output models. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2018), pp. 3004–3009.
- [14] GRONDMAN, I., VAANDRAGER, M., BUSONTU, L., BABUSKA, R., AND SCHUITEMA, E. Efficient model learning methods for actor-critic control. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 42, 3 (June 2012), 591–602.
- [15] HURAK, Z., AND ZEMANEK, J. Feedback linearization approach to distributed feedback manipulation. In *American control conference* (Montreal, Canada, 2012), pp. 991–996.
- [16] JACKSON, D. A new, node-focused model for genetic programming. In *Proceedings of the 15th European Conference on Genetic Programming* (Berlin, Heidelberg, 2012), EuroGP'12, Springer-Verlag, p. 49–60.
- [17] KRAWIEC, K., BŁĄDEK, I., AND SWAN, J. Counterexample-driven genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference* (New York, NY, USA, 2017), GECCO '17, ACM, pp. 953–960.
- [18] KUBALÍK, J., ALIBEKOV, E., ŽEGKLITZ, J., AND BABUŠKA, R. Hybrid single node genetic programming for symbolic regression. In *Transactions on Computational Collective Intelligence XXIV - Volume 9770* (Berlin, Heidelberg, 2016), Springer-Verlag, p. 61–82.
- [19] KUBALÍK, J., DERNER, E., AND BABUŠKA, R. Enhanced symbolic regression through local variable transformations. In *Proceedings of the 9th International Joint Conference on Computational Intelligence - Volume 1: IJCCI* (2017), INSTICC, SciTePress, pp. 91–100.
- [20] LEVINE, S., AND ABBEEL, P. Learning neural network policies with guided policy search under unknown dynamics. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 1* (Cambridge, MA, USA, 2014), NIPS'14, MIT Press, p. 1071–1079.
- [21] LILLICRAP, T. P., HUNT, J. J., PRITZEL, A., HEES, N., EREZ, T., TASSA, Y., SILVER, D., AND WIERSTRA, D. Continuous control with deep reinforcement learning, 2015.
- [22] LIOUTIKOV, R., PARASCHOS, A., PETERS, J., AND NEUMANN, G. Sample-based information-theoretic stochastic optimal control. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (May 2014), pp. 3896–3902.
- [23] SCHMIDT, M., AND LIPSON, H. Distilling free-form natural laws from experimental data. *Science* 324, 5923 (2009), 81–85.
- [24] SEARSON, D. P. Gptips 2: An open-source software platform for symbolic data mining. In *Handbook of Genetic Programming Applications* (2014).
- [25] STAELENS, N., DESCHRIJVER, D., VLADISLAVLEVA, E., VERMEULEN, B., DHAENE, T., AND DEMEESTER, P. Constructing a no-reference h.264/avc bitstream-based video quality metric using genetic programming-based symbolic regression. *IEEE Trans. Cir. and Sys. for Video Technol.* 23, 8 (Aug. 2013), 1322–1333.
- [26] VLADISLAVLEVA, E., FRIEDRICH, T., NEUMANN, F., AND WAGNER, M. Predicting the energy output of wind farms based on weather data: Important variables and their correlation. *Renewable Energy* 50 (2013), 236 – 243.