

Dimension Reduction for Symbolic Regression

Paul Kahlmeyer, Markus Fischer, Joachim Giesen

Friedrich Schiller University Jena

paul.kahlmeyer@uni-jena.de, markus.fischer@uni-jena.de, joachim.giesen@uni-jena.de

Abstract

Solutions of symbolic regression problems are expressions that are composed of input variables and operators from a finite set of function symbols. One measure for evaluating symbolic regression algorithms is their ability to recover formulae, up to symbolic equivalence, from finite samples. Not unexpectedly, the recovery problem becomes harder when the formula gets more complex, that is, when the number of variables and operators gets larger. Variables in naturally occurring symbolic formulas often appear only in fixed combinations. This can be exploited in symbolic regression by substituting one new variable for the combination, effectively reducing the number of variables. However, finding valid substitutions is challenging. Here, we address this challenge by searching over the expression space of small substitutions and testing for validity. The validity test is reduced to a test of functional dependence. The resulting iterative dimension reduction procedure can be used with any symbolic regression approach. We show that it reliably identifies valid substitutions and significantly boosts the performance of different types of state-of-the-art symbolic regression algorithms.

1 Introduction

In regression, we are given n observations

$$(x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)})$$

of d input variables $x \in \mathbb{R}^d$ and one output variable $y \in \mathbb{R}$. The goal is to use the observations to find a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$ that generalizes well beyond the given observations. Symbolic regression adds another goal, namely, to find an interpretable function f . In symbolic regression, f is taken from a space of symbolic expressions that are composed of the input variables x_1, \dots, x_d and a finite set of function symbols, such as, for instance,

$$\{+, -, \cdot, /, \sqrt{\cdot}, \log, \exp, \sin, \cos, \text{const_expr}\}.$$

Washburn’s equation

$$L = \sqrt{\frac{\gamma r t \cos(\phi)}{2\eta}}$$

is an illustrative example out of a set of 880 formulas that were extracted from Wikipedia articles by Guimerà

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

et al. (2020). Washburn’s equation describes the penetration length (L) of a liquid into a capillary pore as a function of time (t), properties of the liquid, namely, its dynamic viscosity (η) and its surface tension (γ), and the geometry of the set-up, that is, the pore’s radius (r) and the contact angle (ϕ) between the pore and the liquid. A key insight from Washburn’s equation is that the penetration length grows proportionally to the square root of the time variable, and the factor of proportionality is a function of the liquid’s properties and the geometry of the set-up. The right-hand side of Washburn’s equation, short Washburn’s formula, is a valid symbolic expression that is composed of five input variables

$$x_1 = t, x_2 = \eta, x_3 = \gamma, x_4 = r, \text{ and } x_5 = \phi,$$

and the function symbols \cdot , $/$, $\sqrt{\cdot}$, \cos , and `const_expr`, where the multiplication symbol is used three times and `const_expr` is instantiated by the constant 2. Washburn’s formula is already quite complex. Its expression tree has five input variable nodes and eight operator nodes that are labeled by function symbols, that is, 13 nodes in total, a number that serves as a measure of complexity.

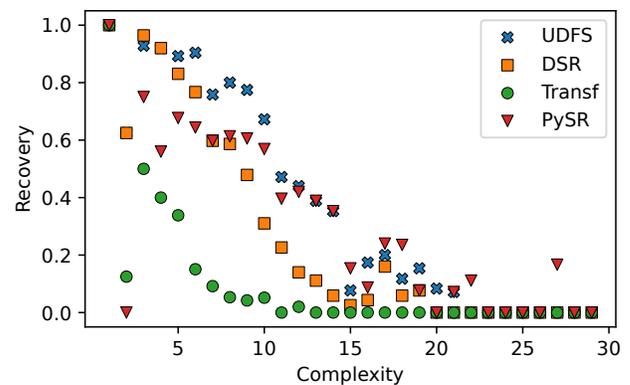


Figure 1: Fraction of successfully recovered expressions from a set of 880 formulas from Wikipedia’s list of eponymous equations (Guimerà et al. 2020) for the state-of-the-art symbolic regressors UDFS: search over expression DAGs, DSR: reinforcement learning, TRANSF: seq2seq learning, and PYSR: genetic programming.

Figure 1 shows, that the ability of state-of-the-art re-

gressors to recover formulas from sampled data decreases when the complexity of the formulas grows. As suggested by La Cava et al. (2021), the complexity of a formula is measured by the size of a corresponding minimal expression tree. It turns out that none of the different, state-of-the-art symbolic regressor algorithms from Figure 1 can recover Washburn’s formula from observational data.

We can simplify Washburn’s formula, for instance, by representing the geometric set-up that is given by r and ϕ in one variable $g(r, \phi) = r \cos(\phi)$. The size of the expression tree for the simplified Washburn formula, where the input variable nodes for r and ϕ are substituted by an input variable node for $g(r, \phi)$, is reduced from twelve to nine, rendering the symbolic regression more tractable. After the substitution, the UDFS regressor from Figure 1 can recover Washburn’s formula from the observational data. Obviously, more aggressive simplifications are possible. However, since we are only given observations for the input variables and the corresponding output variables, valid substitutions are not obvious and finding them becomes a task of its own.

Udrescu et al. (2020) were first to suggest predicates for valid substitutions in their symbolic regression framework AIFeynman. The predicates, however, work only for a limited set of substitutions $g(x_i, x_j)$, namely differences $x_i - x_j$, sums $x_i + x_j$, products $x_i \cdot x_j$, and quotients x_i/x_j of input variables x_i and x_j . The predicate for the substitution $g(x_i, x_j) = x_i - x_j$ works by comparing

$$f(\dots, x_i, \dots, x_j, \dots) \text{ with } f(\dots, x_i + a, \dots, x_j + a, \dots)$$

for different values for $x_i, x_j \in \mathbb{R}$ and offsets $a \in \mathbb{R}$. If $g(x_i, x_j)$ is a valid substitution, then the difference of the two function values should be very small. The predicates for the remaining three substitutions work similarly.

Here, we generalize the ideas of Udrescu et al. (2020) to much more general substitutions. We use the framework of Kahlmeyer et al. (2024) to enumerate expression DAGs up to a certain complexity. Every DAG represents a possible substitution. For checking if a possible substitution is a valid substitution, we build on recent work by Chatterjee (2021) and generalizations by Azadkia and Chatterjee (2021) and Deb, Ghosal, and Sen (2020) for certifying functional dependence from a sample. The fundamental assumption behind a regression problem is that the observations represent a functional dependence between the input variables x_1, \dots, x_d and the output variable y . Let x_I be a subset of the input variables and $g(x_I)$ a possible substitution for x_I . If g is valid, then the observations should still represent a functional dependence if we substitute $g(x_I)$ for x_I in every data point.

Our substitution finding approach can be used together with any symbolic regression algorithm. In fact, as we show in the experimental section of this paper, it reduces the number of variables in the formulas of an established benchmark data set (La Cava et al. 2021) by $\sim 50\%$ and thereby significantly boosts the recovery performance of many conceptually fairly different, state-of-the-art symbolic regression algorithms

Outline. In Section 2 we formally define substitutions, and review the recent progress in functional dependence

measures in Section 3. Then, in Section 4, we use the framework of Kahlmeyer et al. (2024) for enumerating small expression DAGs together with the dependence measures from the previous section to find valid substitutions that we iteratively use for reducing the dimension of symbolic regression problems. We evaluate our approach in Section 5 on the Wikipedia eponymous equations data set (Guimerà et al. 2020) and on the Feynman symbolic regression data set¹. Finally, we draw some conclusions in Section 6.

2 Substitutions

Here, we formally introduce the concept of substitutions that we are going to use for reducing the complexity of symbolic regression problems. Based on the example in the introduction, the concept should be intuitively clear, we simply replace some of the variables with a function thereof. This intuition is formalized in the definition of *input substitutions*.

Input substitutions. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a function, let $I \subseteq [d] = \{1, \dots, d\}$ with $|I| > 1$ be an index set, and let $g: \mathbb{R}^{|I|} \rightarrow \mathbb{R}^k$ be a function with $k < |I|$. Then g is an *input substitution* if there exists a function $f_g: \mathbb{R}^{d-|I|+k} \rightarrow \mathbb{R}$ such that

$$f(x) = f_g(g(x_I), x_{\setminus I})$$

for all $x \in \mathbb{R}^d$. Here, x_I denotes the projection of x onto the indices in I and $x_{\setminus I}$ denotes the projection of x onto the indices in $[d] \setminus I$.

In the context of regression problems, where we want to find the functional dependence f between the input variables x and the output variable y , we can replace the regression problem for x and y with a regression problem for input variables $g(x_I), x_{\setminus I}$ and output variable y . That is, for a given input substitution g , we are looking for a function f_g such that

$$f(x) = f_g(g(x_I), x_{\setminus I}).$$

Thus, the function f can be reconstructed symbolically from g and f_g . The four substitutions that are implemented in AIFeynman (Udrescu et al. 2020) are special cases of bivariate input substitutions, namely,

$$g(x_i, x_j) = x_i \text{ op } x_j \text{ with } \text{op} \in \{+, -, \cdot, /\} \text{ and } i, j \in [d].$$

When searching for input substitutions, another, less intuitive type of substitution makes sense, namely, substitutions that involve not only the input variables but also the output variable. We call this type of substitution an *out-input substitution*.

Out-input substitutions. Let $f: \mathbb{R}^d \rightarrow \mathbb{R}$ be a function, $I \subseteq [d]$ with $1 \leq |I| < d$ an index set, and $h: \mathbb{R}^{|I|+1} \rightarrow \mathbb{R}$. Then h is an *out-input substitution* if there exist functions

$$g: \mathbb{R}^{d-|I|} \rightarrow \mathbb{R} \quad \text{and} \quad f_g: \mathbb{R}^{|I|+1} \rightarrow \mathbb{R}$$

such that

$$h(x_I, f(x)) = g(x_{\setminus I}) \quad \text{and} \quad f(x) = f_g(g(x_{\setminus I}), x_I)$$

¹<https://space.mit.edu/home/tegmark/aifeynman.html>

for all $x \in \mathbb{R}^d$.

Conceptually, a *simple* out-input substitution h is used in an auxiliary symbolic regression problem to find a *complex* input substitution g . The auxiliary regression problem has input variables $x_{\setminus I}$ and output variable $h(x_I, y)$, that is, we are looking for a functional dependence of the form,

$$h(x_I, y) = g(x_{\setminus I}).$$

Given h and a solution g of the auxiliary problem in symbolic form, we can solve the equation $h(x_I, y) = g(x_{\setminus I})$ symbolically for y . That is, y becomes a symbolic expression

$$y = f_g(g(x_{\setminus I}), x_I),$$

and the function f that is sought in the original regression problem can be reconstructed from g and f_g as

$$f(x) = f_g(g(x_{\setminus I}), x_I).$$

The following example demonstrates that it can be beneficial to also consider out-input substitutions. The function

$$f(x_1, x_2, x_3) = x_1 x_2 x_3 + \frac{x_1(x_2 + \log(x_2))}{x_3}$$

does not have a *simple* input substitution, but the simple out-input substitution $h(x_1, y) = y/x_1$ with $I = \{1\}$. As one can easily check, the more *complex* (input substitution) function

$$g(x_2, x_3) = x_2 x_3 + \frac{x_2 + \log(x_2)}{x_3}$$

and the function $f_g(z, x_1) = x_1 z$ satisfy

$$\begin{aligned} h(x_1, f(x_1, x_2, x_3)) &= g(x_2, x_3) \text{ and} \\ f(x_1, x_2, x_3) &= x_1 g(x_2, x_3) = f_g(g(x_2, x_3), x_1). \end{aligned}$$

For both substitution types, that is, input and out-input substitutions, the substituted regression problems, including the auxiliary problems, have fewer variables than the original problems, and thus should be easier to solve for symbolic regression algorithms. In Section 5 we show that this is indeed the case for many state-of-the-art symbolic regression algorithms.

3 Functional Dependence Measures

The fundamental assumption behind regression problems is that there exists a functional dependence between input and output variables. This assumption can be tested with functional dependence tests, or more generally by functional dependence measures. Functional dependence measures use random samples $(X_i, Y_i)_{i \in [n]}$ to quantify the degree to which a function $f(X) = Y$ between two random variables X and Y exists. In general, it is necessary to introduce additional constraints on the function class from which f is taken. Otherwise, for pairwise distinct samples $(X_i, Y_i)_{i \in [n]}$, one always has the discrete function with $f(X_i) = Y_i$. Classical and widely used dependence measures are the Pearson correlation coefficient (Pearson 1920)

for linear functions or Spearman's rank correlation (Spearman 1904) for monotone functions. Recently, Chatterjee (2021) introduced a new rank-based dependence measure, which works for general measurable univariate functions. Follow-up work by Deb, Ghosal, and Sen (2020) and Azadkia and Chatterjee (2021) extend this measure to multivariate functions. In the following, we describe the latter three dependence measures.

We start with a description of Chatterjee (2021)'s correlation coefficient. Given n pairs of random variables (X_i, Y_i) that have been drawn i.i.d. from some probability distribution and the variables X_i have been sorted in ascending order. The Chatterjee coefficient is then defined as

$$\xi_n(X, Y) = 1 - \frac{n \sum_{i=1}^{n-1} |r_{i+1} - r_i|}{2 \sum_{i=1}^n l_i(n - l_i)},$$

where r_i denotes the rank of Y_i , that is, the number of indices j such that $Y_j \leq Y_i$. Analogously, l_i is defined as the number of indices j such that $Y_j \geq Y_i$. For $n \rightarrow \infty$, this coefficient $\xi_n(X, Y)$ converges to 0 if X_i and Y_i are independent variables. It converges to 1 if there exists a measurable function $f: \mathbb{R} \rightarrow \mathbb{R}$ such that $Y = f(X)$.

This can be explained intuitively as follows. If all X_i and all Y_i are distinct, then, after a small calculation, the coefficient simplifies to

$$\xi_n(X, Y) = 1 - \frac{3 \sum_{i=1}^{n-1} |r_{i+1} - r_i|}{n^2 - 1}.$$

Moreover, if we assume further that $Y_i = f(X_i)$ for a continuous function f and that the X_i are sorted and densely sampled, then not only

$$|Y_{i+1} - Y_i| = |f(X_{i+1}) - f(X_i)|$$

should be small, but also the rank difference $|r_{i+1} - r_i|$. In this case, we have,

$$\sum_{i=1}^{n-1} |r_{i+1} - r_i| \in O(n),$$

which means that $\xi_n(X, Y)$ converges to 1.

The Chatterjee coefficient is defined only for univariate functions $f: \mathbb{R} \rightarrow \mathbb{R}$ and thus is not suitable for our intended application in symbolic regression where we need to decide functional dependence of multivariate functions. Azadkia and Chatterjee (2021) have generalized the Chatterjee coefficient to the multivariate case. Their coefficient, called CODEC, is designed for detecting a functional dependence between random variables, conditioned on another set of variables. In our application, we can simply condition on the empty set. CODEC is defined for random variables X that take values in \mathbb{R}^d and Y that take values in \mathbb{R} . Let $(X_i, Y_i)_{i \in [n]}$ be a sample that is drawn i.i.d from some probability distribution on $\mathbb{R}^d \times \mathbb{R}$. The values r_i and l_i are the upper respectively lower ranks as for the Chatterjee coefficient. Let $\nu(i)$ denote the index such that, $X_{\nu(i)}$ is the nearest neighbor of X_i , where ties are broken arbitrarily. CODEC is now defined as

$$\text{CODEC}_n(X, Y) = \frac{\sum_{i=1}^n (n \min\{r_i, r_{\nu(i)}\} - l_i^2)}{\sum_{i=1}^n l_i(n - l_i)}.$$

As for the Chatterjee coefficient, if X and Y are independent, then $\text{CODEC}_n(X, Y)$ converges to 0 for $n \rightarrow \infty$. It converges to 1 if Y is a measurable function of X . The intuition behind CODEC is similar to the intuition behind the Chatterjee coefficient as well. Using the well known identity

$$\min\{a, b\} = \frac{1}{2}(a + b - |a - b|),$$

we can rewrite the numerator of the CODEC coefficient as

$$\begin{aligned} & \sum_{i=1}^n \left(n \min\{r_i, r_{\nu(i)}\} - l_i^2 \right) \\ &= \frac{n}{2} \left(R + S - \sum_{i=1}^n |r_i - r_{\nu(i)}| \right) - L, \end{aligned}$$

where

$$L = \sum_{i=1}^n l_i^2, \quad R = \sum_{i=1}^n r_i, \quad \text{and} \quad S = \sum_{i=1}^n r_{\nu(i)}.$$

The CODEC coefficient now becomes,

$$\text{CODEC}_n(X, Y) = \frac{\frac{n}{2} \left(R + S - \sum_{i=1}^n |r_i - r_{\nu(i)}| \right) - L}{\sum_{i=1}^n l_i(n - l_i)},$$

which becomes large, when $\sum_{i=1}^n |r_i - r_{\nu(i)}|$ is small. This happens exactly when the Y_i values, and thus the ranks, of sample neighboring points are close.

Both, the Chatterjee and the CODEC coefficient are special cases of a more general class of dependence measures, called KMAC, short for kernel measure of association (Deb, Ghosal, and Sen 2020). KMAC measures also quantify how close the Y_i values are for sample points X_i that are close to each other. To do so, KMAC measures use geometric graphs, such as, k -nearest-neighbor graphs or a Euclidean minimum spanning trees, on the sample points X_i to encode which sample points should be considered close to each other. The distance between the values Y_i is measured through a mapping into a reproducing kernel Hilbert space. In theory, this allows us to assess functional dependence not only on \mathbb{R}^d but on arbitrary topological spaces, as long as we have a notion of closeness for the sample points.

4 Finding and Using Substitutions

Our goal is to find input substitutions and out-input substitutions to transform the observations

$$(x^{(i)}, y^{(i)}) \in \mathbb{R}^d \times \mathbb{R} \quad \text{for } i \in [n],$$

that define a symbolic regression problem into a simpler regression problem with fewer variables.

Input substitutions. Let $I \subseteq [d]$ with $|I| > 1$ be an index set, and let $g: \mathbb{R}^{|I|} \rightarrow \mathbb{R}^k$ be a function with $k < |I|$ that transforms the observations into

$$\left((g(x_I^{(i)}), x_{\setminus I}^{(i)}), y^{(i)} \right) \in \mathbb{R}^{d-|I|+k} \times \mathbb{R}.$$

If there is a functional dependence between the transformed input variables $(g(x_I^{(i)}), x_{\setminus I}^{(i)})$ and the output variables $y^{(i)}$,

then the transformed observations amount to an *auxiliary* symbolic regression problem with only $d - |I| + k$ input variables. From g and a symbolic solution f_g of the auxiliary regression problem, a solution f of the original problem can be reconstructed.

Out-input substitutions. Let $I \subseteq [d]$ with $1 \leq |I| < d$ be an index set, and let $h: \mathbb{R}^{|I|+1} \rightarrow \mathbb{R}$ be a function that transforms the observations into

$$\left(x_{\setminus I}^{(i)}, h(x_I^{(i)}, y^{(i)}) \right) \in \mathbb{R}^{d-|I|} \times \mathbb{R}.$$

If there is a functional dependence between the transformed input variables $x_{\setminus I}^{(i)}$ and the output variable $h(x_I^{(i)}, y^{(i)})$, then the transformed observations amount to a symbolic regression problem with only $d - |I|$ input variables. Assume that there is such a dependence and that the entailed *auxiliary* symbolic regression problem is solved by a function $g: \mathbb{R}^{d-|I|} \rightarrow \mathbb{R}$. We can use the symbolically given functions h and g to derive a symbolic function f_g by solving the equation $h(x_I, y) = g(x_{\setminus I})$ for y . Examples can be found in Section 2 and in the full version of the paper. The symbolic solution f of the original symbolic regression problem can be reconstructed from g and f_g . As mentioned before, g is actually an input substitution that might be too complex to be found directly.

Given the observations that define a symbolic regression problem, we can use both input substitutions g and in-output substitutions h to derive a simpler, auxiliary symbolic regression problem. Therefore, the task becomes to find valid substitutions g and h . We address this task by systematically searching for g and h in a space of functions with succinct symbolic representations. In their search-based regressor, Kahlmeyer et al. (2024) enumerate expression DAGs up to a given size. For a fixed budget of operator nodes, the representation by DAGs allows covering more functions than the commonly used expression tree representation, because in DAGs common subexpressions need to be represented only once. Here, we use the enumeration of small expression DAGs to systematically enumerate *simple* candidate substitutions g and h . By a functional dependence test on transformed observations, we can decide if a candidate is indeed a valid substitution. Substitutions can be used not only on the observations for the original symbolic regression problem, but also for auxiliary problems that are given by transformed observations. That is, we can continue and iteratively simplify the problems further.

Beam search for symbolic regression. The iterative search for substitutions is naturally organized in a *search-tree*. At the root of the search tree is the original regression problem, which is given by the observations. The children of any node in the tree are auxiliary regression problems that are derived from valid substitutions, and the leafs of the tree are regression problems for which no substitution can be found. By the definition of substitutions, regression problems at child nodes always have fewer variable nodes than the regression problems at their parent nodes. The regression

problems at all nodes of the search tree can be addressed by any symbolic regression algorithm.

However, the functional dependence measures from Section 3 do not provide a definite answer, but only a functional dependence score between 0 and 1, where scores close to 1 indicate functional dependence and scores close to 0 its absence. The scores can be used to selectively explore the search tree by a *beam search*. That is, every candidate substitution becomes a node in the search tree, but now the nodes are weighted by a score. As usual in beam search, a beam size hyperparameter controls the size of the search space by limiting the number of nodes on each level to the beam-size-many highest scoring substitutions. The search tree is then explored in a breadth-first manner, where the beam size parameter restricts the number of paths that are followed. After the search has ended, the regression problems on the path that is leading to the highest scoring regression problem are passed on to an established symbolic regression algorithm. The beam-search algorithm for a symbolic regression is illustrated in Figure 2.

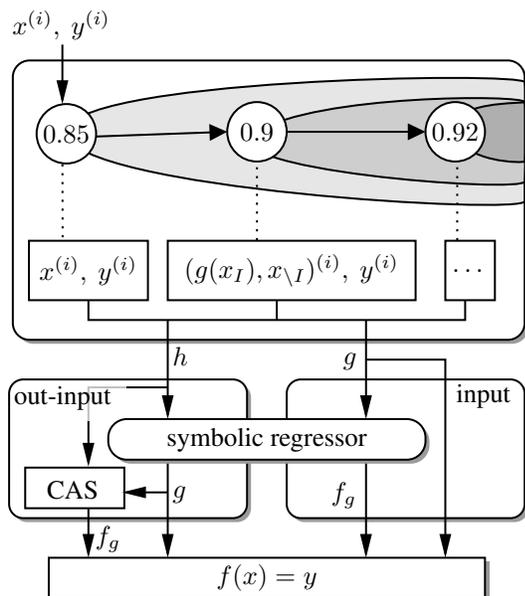


Figure 2: Illustration of the beam search approach for symbolic regression. Each node of the beam search tree has an associated regression data set and a functional dependence score for this data set. The edges are labeled by candidate substitutions. The highest scoring regression problems, over all levels, are passed to a symbolic regression algorithm. For input substitutions g , a symbolic regression algorithm computes a function f_g . For out-input substitutions h , the symbolic regression algorithm computes g , which is passed together with h to a computer algebra system (CAS) to find a function f_g . A solution f of the original symbolic regression problem is reconstructed from all functions g and f_g along the path from the root to the overall highest scoring node in the beam search tree.

5 Experiments

Since our work is motivated by the observation that symbolic regression becomes easier for smaller expressions, we evaluate the ability of our beam-search approach to reduce the complexity of symbolic regression problems in the following experiments. We assess the dimension-reduction ability by the *reduction rate* for formulas with known symbolic expressions. Our beam search is parameterized by a choice of functional dependence measure, by the beam size, by the choice of substitution type, by the symbolic regression algorithm, and by the expression DAG search space. To gauge the effect of the first three parameters, we evaluate the reduction rate for different dependence measures, different beam sizes, and three choices for the substitution types. In a second experiment, we evaluate the effectiveness of combining our beam search with different state-of-the-art symbolic regression algorithms. In this experiment, we compare the *recovery rate* of the symbolic regression algorithms with and without our beam search. For efficiency reasons, we keep the DAG search space small and fixed in all experiments.

Datasets. For our experiments, we have used two sets of regression problems, namely, Wikipedia’s list of 880 eponymous equations (Guimerà et al. 2020) and 114 formulas that were extracted from the Feynman lecture notes of physics (Udrescu and Tegmark 2020). We provide more details about both data sets in the full version of the paper.

Implementation Details. At each node of the beam search tree, we search for substitutions in a space of expression DAGs. To keep the search space small, we consider only expression DAGs with at most one intermediary node and one output node, that is, scalar functions with $k = 1$.

The KMAC measures are parameterized by a geometric graph and a kernel function. Here, we follow the original paper by Deb, Ghosal, and Sen (2020), and use a one-nearest neighbor graph and a Gaussian RBF kernel.

For out-input substitutions, we derive f_g from h and g by using the computer algebra system SymPy (Meurer et al. 2017).

Hardware. All experiments were run on a computer with an Intel Xeon Gold 6226R 64-core processor, 128 GB of RAM, running Python 3.10.

Parameter Study

For evaluating the effect of the choice of dependence measure, the beam size parameter, and the substitution type, we first need to define the reduction rate for formulas with known symbolic expressions, before we compute these rates in different experimental settings.

Reduction rate. For computing the reduction rate in a given parameter setting, we run our beam search on formulas with known symbolic expressions as described in Section 4, except for the last step. That is, we do not pass the reduced regression problems to a symbolic regression algorithm. Instead, we check the correctness of the selected substitutions.

The *reduction rate* is the best dimension reduction that is achieved by the beam search, that is,

$$\text{reduction rate} = 1 - \frac{\min_{\text{nodes}} \# \text{vars of reduced problem}}{\# \text{vars of original problem}}.$$

We use SymPy for checking the validity of a substitution. Let g be a candidate input substitution that replaces the input variables x_I with $g(x_I)$. Since the candidate substitutions are given symbolically, we can introduce a new variable symbol γ and solve the equation $\gamma = g(x_I)$ for some input variable x_i with $i \in I$. Then, we plug the expression for x_i , which depends on γ , into the known ground truth formula for f . This gives us a symbolic expression \hat{f} . If g is a valid substitution, then all input variables $x_j, j \in I \setminus \{i\}$ can be eliminated from \hat{f} . That is, \hat{f} has an equivalent symbolic representation that only depends on γ and the input variables $x_{\setminus I}$. The validity of a candidate out-input substitution h that replaces the input variables x_I and the output variable y with $h(x_I, y)$ can be checked similarly. Given the ground truth formula f in symbolic form, we can substitute $f(x)$ for y in $h(x_I, y)$. If h is valid, then the symbolic expression for $h(x_I, f(x))$ does not depend on x_I . We provide examples in the full version of the paper.

Choice of dependence measure. We measure the reduction rate of our beam search when used with the CODEC and KMAC dependence measures, and a simple baseline measure. The baseline measure is a simple volume heuristic. The observations $(x^{(i)}, y^{(i)})_{i \in [n]}$ are contained in a submanifold of \mathbb{R}^{d+1} if they are sampled from a well-behaved function f with $y = f(x)$. In our volume heuristics, we compute for every point $x^{(i)}$ the volume of the parallelepiped that is spanned by the difference vectors of $x^{(i)}$ to its $d + 1$ nearest neighbors. These volumes should be close to zero if the observations are sampled densely from a well-behaved function. The volume is given by the determinant of the matrix of difference vectors. For the dependence measure, we take the average over the n volumes,

For the experiment, we follow the experimental protocol by La Cava et al. (2021) and sample the ground-truth functions in the eponymous equations and the Feynman equations data sets, for which the true functional dependencies $y = f(x)$ are known. The samples for Feynman equations data set are directly given by La Cava et al. (2021), who also add different levels of Gaussian noise to the sample points. Sampling from the functions in the eponymous equations data set is described in the full version of the paper.

We report average reduction rates for the different dependence measures and beam size 1 in Figure 3. In the noise-free case, all three dependence measures perform similarly. The number of input variables is reduced by $\sim 50\%$ for the Feynman equations data set and by $\sim 35\%$ for the eponymous equations data set. At non-zero noise levels, however, the baseline measure breaks down, whereas CODEC and KMAC still achieve substantial average reduction rates.

While CODEC and KMAC perform similarly in terms of average reduction rate, they differ significantly, almost an

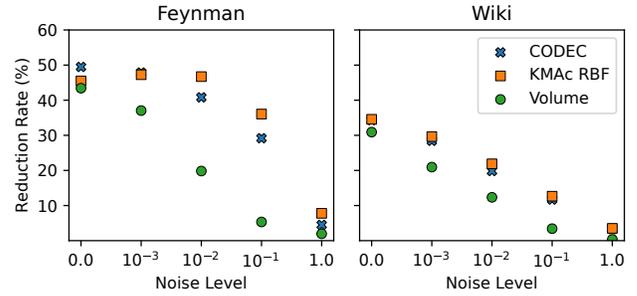


Figure 3: Average reduction rates of our beam search, with beam size 1, on the eponymous equations (Wiki) and the Feynman equations (Feynman) data sets for the functional dependence measures CODEC and KMAC, and a volume-based baseline measure.

order of magnitude, in running times, as can be seen in Table 1. Therefore, we use our beam search with CODEC in the following.

	CODEC	KMAC	Volume
Feynman	3.32 ± 0.08	24.29 ± 0.29	21.58 ± 0.44
Wiki	2.88 ± 0.02	23.24 ± 0.14	20.92 ± 0.15

Table 1: Average running times (in milliseconds) for the identification of a single substitution on the eponymous equations and the Feynman equations data sets. The average and standard deviations have been computed over 50 runs.

Beam size. We measure the reduction rate of our beam search for five values of the beam size parameter, while keeping the functional dependence measure CODEC fixed. A larger beam size broadens the view of the search process and potentially finds better substitutions, however, at the cost of increased running times. We report results for beam sizes 1 to 5 in Table 2. Since the reduction rate improves only marginally with increasing beam size, we use our beam search with beam size 1 in the following.

beam size	1	2	3	4	5
Feynman	0.489	0.504	0.491	0.491	0.491
Wiki	0.346	0.354	0.354	0.357	0.358

Table 2: Average reduction rates of for beam sizes 1 to 5.

Substitution type. So far, we have used input substitutions as well as out-input substitutions in the experiment. Here, we assess how much out-input substitutions add to the reduction-rate performance of the beam search. In Table 3 we compare the reduction rate of our beam search for out-input and input substitutions, only input substitutions, and the four substitutions from the AIFeynman project (Udrescu et al. 2020). Remember that, conceptually, out-input substitutions are used for indirectly finding input substitutions that are more complex than the functions in the direct search

space. Here, the search space is not very large, since we only consider expression DAGs with one input and one output node. However, as can be seen on the eponymous equations data set (Wiki), it covers more input substitutions than AIFeynman. Adding out-input substitutions increases the reduction rates significantly on both data sets.

	out+input	input	AIFeynman
Feynman	0.49	0.39	0.39
Wiki	0.34	0.19	0.08

Table 3: Comparison of the average reduction rate for out-input and input substitutions (out+input), only input substitutions, and the AIFeynman substitutions.

Recovery Rate

As described in Section 4, our beam search can be turned into a symbolic regressor by passing reduced problem instances that have been discovered in the search process to a symbolic regression algorithm. An established measure for assessing the quality of a symbolic regression algorithm is the *recovery rate*, that is, the fraction of formulas with known symbolic expressions that are recovered by a symbolic regression algorithm up to symbolic equivalence (La Cava et al. 2021). Here, we use again SymPy to automatically decide the symbolic equivalence of expressions. Following the symbolic regression benchmark paper by La Cava et al. (2021), we also measure model fit and model complexity, besides the recovery rate. A good symbolic regression algorithm computes models that provide high recovery, have good model fits, and are of low complexity. Here, model fit is measured by the normalized root-mean-square error (NRMSE) on hold out data, and the complexity of an expression is measured by the number of nodes in a corresponding expression tree that has been simplified by SymPy.

Symbolic regression algorithms. In the experiment, we compare the recovery rate of several state-of-the-art symbolic regression algorithms in combination with our beam search, or without. The beam search in the experiment uses beam size 1 and the CODEC functional dependence measure. One can roughly distinguish three types of symbolic regression algorithms, namely, genetic programming algorithms, sampling algorithms, and search algorithms. We provide a review of the different types of symbolic regression algorithms, together with examples in the full version of paper. For our experiment, we used our beam search with representatives for all three types: OPERON (Burlacu, Kronberger, and Kommenda 2020) and PYSR (Cranmer 2023) for genetic programming, the sequence-to-sequence learning approach TRANF by Kamienny et al. (2022) and deep symbolic regression (DSR) by Mundhenk et al. (2021) as two examples for sampling algorithms, and the unbiased DAG frame search (UDFS) by Kahlmeyer et al. (2024) as an example for a search algorithm. Moreover, we use the beam search also with a standard regression model, namely, polynomial regression with polynomials of degree at most two (POLY).

If applicable, we set the hyperparameters of the symbolic regression algorithms so that our experiments run in reasonable time. An overview of the respective hyperparameters can be found in the full version of the paper.

Results. Results on the Feynman equations data set are shown in Table 4. More details and results for the eponymous equations data set are provided in the full version of the paper. Notably, the recovery rate improves significantly for all regression algorithms included in the experiment. As can be seen, this typically also affects the average model fit, because correctly recovered expressions have RMSEs of zero, or numerically very close to zero, on test data. The complexity of the expressions that are returned by the symbolic regression algorithms is not affected much by the beam search. The only exception is polynomial regression, where the expression trees shrink on average by more than 20 nodes when using the beam search.

	recovery		NRMSE		complexity	
	base	beam	base	beam	base	beam
UDFS	0.58	0.69	0.08	0.03	10.26	12.71
DSR	0.37	0.66	0.14	0.03	19.09	21.17
TRANSF	0.01	0.11	0.09	0.03	33.53	28.97
PYSR	0.51	0.62	0.09	0.03	11.74	14.86
OPERON	0.41	0.52	0.05	0.02	34.81	40.93
POLY	0.06	0.34	0.18	0.08	54.2	32.09

Table 4: Performance measures of different symbolic regression algorithms in combination with our search, or without (**bold** is better). Recovery rate, average normalized root-mean-square error on test data, and average model complexity on the Feynman equations data set.

It turns out, that even for problems that are not recovered by the algorithms, the returned expressions have more subexpressions in common with the ground-truth formula when used together with our beam search. See the full version of the paper for details.

6 Conclusions

Motivated by the straightforward, yet seemingly so far mostly overlooked observation that symbolic regression becomes more difficult for more complex formulas, we have designed and implemented a dimension reduction method for symbolic regression. Given a symbolic regression problem in the form of observations, our method systematically searches for small substitutions that transform the observed data into a lower dimensional data set, that is, a data set where each observation is replaced with a data point with fewer variables. The method can be applied iteratively to the dimension-reduced data sets. We have implemented the idea of an iterative dimension reduction in a beam search. Combining the beam search with state-of-the-art symbolic regression algorithms significantly boosts the symbolic recovery performance of these algorithms.

Acknowledgements

This work was supported by the Carl Zeiss Stiftung within the project "Interactive Inference".

References

- Azadkia, M.; and Chatterjee, S. 2021. A simple measure of conditional dependence. *The Annals of Statistics*, 49(6): 3070 – 3102.
- Burlacu, B.; Kronberger, G.; and Kommenda, M. 2020. Operon C++: An Efficient Genetic Programming Framework for Symbolic Regression. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion, GECCO '20*, 1562–1570.
- Chatterjee, S. 2021. A New Coefficient of Correlation. *Journal of the American Statistical Association*, 116(536): 2009–2022.
- Cranmer, M. 2023. Interpretable Machine Learning for Science with PySR and SymbolicRegression.jl.
- Deb, N.; Ghosal, P.; and Sen, B. 2020. Measuring Association on Topological Spaces Using Kernels and Geometric Graphs.
- Guimerà, R.; Reichardt, I.; Aguilar-Mogas, A.; Massucci, F. A.; Miranda, M.; Pallarès, J.; and Sales-Pardo, M. 2020. A Bayesian machine scientist to aid in the solution of challenging scientific problems. *Science Advances*, 6(5): eaav6971.
- Kahlmeyer, P.; Giesen, J.; Habeck, M.; and Voigt, H. 2024. Scaling Up Unbiased Search-based Symbolic Regression. In Larson, K., ed., *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 4264–4272. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Kamienny, P.-A.; d'Ascoli, S.; Lample, G.; and Charton, F. 2022. End-to-end Symbolic Regression with Transformers. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- La Cava, W.; Orzechowski, P.; Burlacu, B.; de Franca, F.; Virgolin, M.; Jin, Y.; Kommenda, M.; and Moore, J. 2021. Contemporary Symbolic Regression Methods and their Relative Performance. In Vanschoren, J.; and Yeung, S., eds., *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1. Curran.
- Meurer, A.; Smith, C. P.; Paprocki, M.; Čertík, O.; Kirpichev, S. B.; Rocklin, M.; Kumar, A.; Ivanov, S.; Moore, J. K.; Singh, S.; Rathnayake, T.; Vig, S.; Granger, B. E.; Muller, R. P.; Bonazzi, F.; Gupta, H.; Vats, S.; Johansson, F.; Pedregosa, F.; Curry, M. J.; Terrel, A. R.; Roučka, v.; Saboo, A.; Fernando, I.; Kulal, S.; Cimrman, R.; and Scopatz, A. 2017. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3: e103.
- Mundhenk, T.; Landajuela, M.; Glatt, R.; Santiago, C.; Faisol, D.; and Petersen, B. 2021. Symbolic Regression via Neural-Guided Genetic Programming Population Seeding.
- Pearson, K. 1920. Notes on the history of correlation. *Biometrika*, 13(1): 25–45.
- Spearman, C. 1904. The proof and measurement of association between two things. *The American journal of psychology*, 15(1): 72–101.
- Udrescu, S.-M.; Tan, A.; Feng, J.; Neto, O.; Wu, T.; and Tegmark, M. 2020. AI Feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 4860–4871. Curran Associates, Inc.
- Udrescu, S.-M.; and Tegmark, M. 2020. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16): eaay2631.