

AI Feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity

Silviu-Marian Udrescu, Andrew Tan, Jiahai Feng, Orisvaldo Neto, Tailin Wu & Max Tegmark
 MIT Dept. of Physics & Center for Brains, Minds & Machines
 Cambridge, MA 02139
 sudrescu@mit.edu

Abstract

We present an improved method for symbolic regression that seeks to fit data to formulas that are Pareto-optimal, in the sense of having the best accuracy for a given complexity. It improves on the previous state-of-the-art by typically being orders of magnitude more robust toward noise and bad data, and also by discovering many formulas that stumped previous methods. We develop a method for discovering generalized symmetries (arbitrary modularity in the computational graph of a formula) from gradient properties of a neural network fit. We use normalizing flows to generalize our symbolic regression method to probability distributions from which we only have samples, and employ statistical hypothesis testing to accelerate robust brute-force search.

1 Introduction




Figure 1: Our symbolic regression of data on how kinetic energy depends on mass, velocity and the speed of light discovers a Pareto-frontier of four formulas that are each the most accurate given their complexity. Convex corners reveal particularly useful formulas, in this case Einstein’s formula and the classical approximation $mv^2/2$.

A central challenge in science is *symbolic regression*: discovering a symbolic expression that provides a simple yet accurate fit to a given data set. More specifically, we are given a table of numbers, whose rows are of the form $\{x_1, \dots, x_n, y\}$ where $y = f(x_1, \dots, x_n)$, and our task is to discover the correct symbolic expression (composing mathematical functions from a user-provided set) for the unknown mystery function f , optionally including the complication of noise and outliers. Science aside, symbolic regression has the potential to replace some inscrutable black-box neural networks

by simple yet accurate symbolic approximations, helping with the timely goal of making high-impact AI systems more interpretable and reliable [1–5].

Symbolic regression is difficult because of the exponentially large combinatorial space of symbolic expressions. Traditionally, symbolic regression has relied on human intuition, leading to the discovery of some of the most famous formulas in science. More recently, there has been great progress toward fully automating the process [6–26], and open-source software now exists that can discover quite complex physics equations by combining neural networks with techniques inspired by physics and information theory [25]. The goal of this paper is to further improve this state-of-the-art, by making three main contributions:

1. We use learned neural-network gradients to discover and exploit modularity in the function’s computational graph.
2. We use statistical hypothesis testing and recursive composition of description-length-based Pareto-frontiers to accelerate and robustify the symbolic regression.
3. We use normalizing flows to enable regression of probability distributions from samples.

We describe our symbolic regression algorithm in Section 2 and test it with numerical experiments in Section 3.

2 Method

Our symbolic regression algorithm uses a divide-and-conquer approach as in [25]. We directly solve a mystery in two base cases: if the mystery function $f(x_1, \dots, x_n)$ is a low-order polynomial or if it is simple enough to be discovered by brute-force search. Otherwise, we recursively try the strategies that we will now describe for replacing it by one or more simpler mysteries, ideally with fewer input variables.

2.1 Leveraging graph modularity against the curse of dimensionality

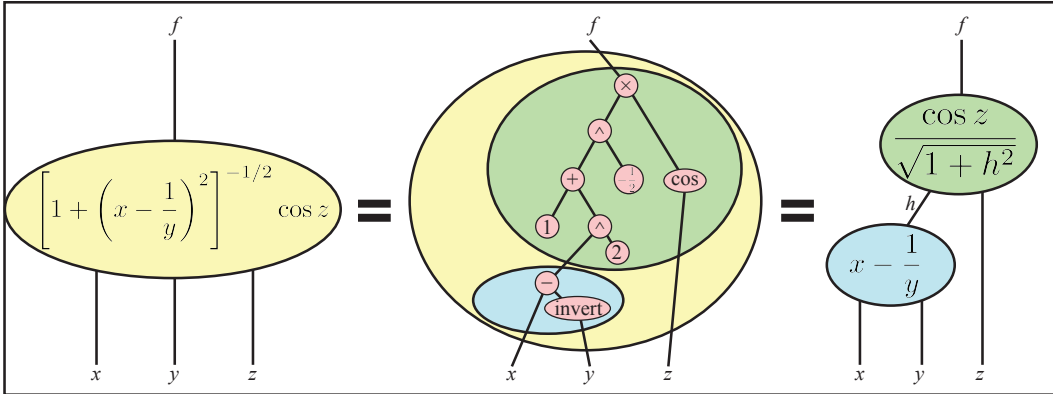


Figure 2: All functions can be represented as tree graphs whose nodes represent a set of basic functions (middle panel). Using a neural network trained to fit a mystery function (left panel), our algorithm seeks a decomposition of this function into others with fewer input variables (right panel), in this case of the form $f(x, y, z) = g[h(x, y), z]$.

When we define and evaluate a mathematical function, we typically represent it as composed of some basis set S of simpler functions. As illustrated in Figure 2 (middle panel), this representation can be specified as a graph whose nodes contain elements of S . The most popular basis functions in the scientific literature tend to be functions of two variables (such as $+$ or \times), one variable (such as \sin or \log) or no variables (constants such as 2 or π). For many functions of scientific interest, this graph is *modular* in the sense that it can be partitioned in terms of functions with fewer input variables, as in Figure 2 (right panel).

A key strategy of our symbolic regression algorithm is to recursively discover such modularity, thereby reverse-engineering the computational graph of a mystery function, starting with no information about

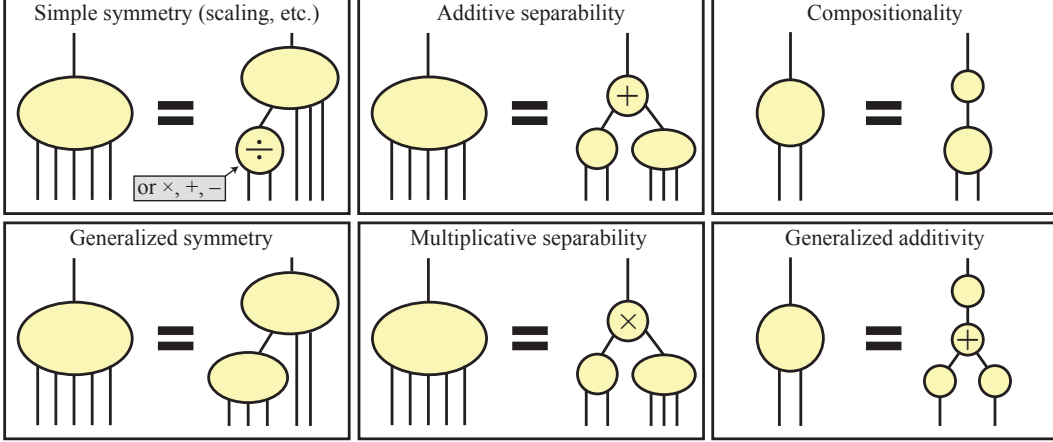


Figure 3: Examples of graph modularity that our algorithm can auto-discover. Lines denote real-valued variables and ovals denote functions, with larger ones being more complex.

it other than an input-output data table. This is useful because there are exponentially many ways to combine n basis functions into a module, making it extremely slow and difficult for brute-force or genetic algorithms to discover the correct function when n is large. Our divide-and-conquer approach of first breaking the function into smaller modules with smaller n that can be solved separately thus greatly accelerates the solution. We implement this modularity discovery algorithm in two steps:

1. Use the user-provided data table to train a neural network $f_{\text{NN}}(\mathbf{x})$ that accurately approximates the mystery function $f(\mathbf{x})$.
2. Perform numerical experiments on $f_{\text{NN}}(\mathbf{x})$ to discover graph modularity.

Specifically, we test for the six types of graph modularity illustrated in Figure 3 and listed in Table 1, and choose between the discovered candidates as described in Section 2.2. Our method for discovering separability is described in [25]. As we will see below, all our other types of graph modularity (compositionality, symmetry and generalized additivity) can be revealed by ∇f , the gradient of our mystery function f .

Compositionality Let us first consider the case of *compositionality* (Figure 3, top right), where $f(\mathbf{x}) = g(h(\mathbf{x}))$ and h is a scalar function simpler than f in the sense of being expressible with a smaller graph as in Figure 3. By the chain rule, we have

$$\nabla f(\mathbf{x}) = g'(h(\mathbf{x}))\nabla h(\mathbf{x}), \quad \text{so} \quad \widehat{\nabla} f = \pm \widehat{\nabla} h, \quad (1)$$

where hats denote unit vectors: $\widehat{\nabla} f \equiv \nabla f / |\nabla f|$, etc. This means that if we can discover a function h whose gradient is proportional to that of f (which we will describe a process for in Section 2.2), then we can simply replace the variables \mathbf{x} in the original mystery data table by the single variable $h(\mathbf{x})$ and recursively apply our AI Feynman algorithm to the new one-dimensional symbolic regression problem of discovering $g(h)$.

Generalized symmetry Let us now turn to *generalized symmetry* (Figure 3, bottom left), where k of the n arguments enter only via some scalar function h of them. Specifically, we say that an f has generalized symmetry if the n components of the vector $\mathbf{x} \in \mathbb{R}^n$ can be split into groups of k and $n - k$ components (which we denote by the vectors $\mathbf{x}' \in \mathbb{R}^k$ and $\mathbf{x}'' \in \mathbb{R}^{n-k}$) such that $f(\mathbf{x}) = f(\mathbf{x}', \mathbf{x}'') = g[h(\mathbf{x}', \mathbf{x}'')] for some function g . By the chain rule, we have$

$$\nabla_{\mathbf{x}'} f(\mathbf{x}', \mathbf{x}'') = g_1[h(\mathbf{x}', \mathbf{x}'')] \nabla h(\mathbf{x}'), \quad \text{so} \quad \widehat{\nabla}_{\mathbf{x}'} f = \pm \widehat{\nabla} h, \quad (2)$$

where g_1 denotes the derivative of g with respect to its first argument. This means that $\widehat{\nabla}_{\mathbf{x}'} f(\mathbf{x}', \mathbf{x}'')$ is independent of \mathbf{x}'' , which it would not be for a generic function f . \mathbf{x}'' -independence of the normalized gradients $\widehat{\mathbf{v}}(\mathbf{x}', \mathbf{x}'') \equiv \widehat{\nabla}_{\mathbf{x}'} f(\mathbf{x}', \mathbf{x}'')$ thus provides a smoking gun signature of generalized symmetry. Whereas our compositionality discovery above requires discovering an explicit function h , we can

Table 1: Simplification strategies

Name	Property	Action
Negativity	$f(x_1, x_2, \dots) < 0$	Solve for $g \equiv -f$
Positivity	$f(x_1, x_2, \dots) > 0$	Solve for $g \equiv \ln f$
Additive separability	$f(x_1, \dots, x_k, x_{k+1}, \dots, x_n) = g(x_1, \dots, x_k) + h(x_{k+1}, \dots, x_n)$	Solve for g & h
Multiplicative separability	$f(x_1, \dots, x_k, x_{k+1}, \dots, x_n) = g(x_1, \dots, x_k)h(x_{k+1}, \dots, x_n)$	Solve for g & h
Simple symmetry	$f(x_1, x_2, \dots) = g(x_1 \odot x_2, \dots)$, $\odot \in \{+, -, \times, /\}$	Solve for g
Compositionality	$f(x_1, \dots, x_n) = g(h(x_1, \dots, x_n))$, h simpler than f	Find h with $\nabla h \propto \nabla f$
Generalized symmetry	$f(x_1, \dots, x_k, x_{k+1}, \dots, x_n) = g[h(x_1, \dots, x_k), x_{k+1}, \dots, x_n]$	Find h satisfying $h_{,i} \propto f_{,i}$, $i = 1, \dots, k$
Generalized additivity	$f(x_1, x_2) = F[g(x_1) + h(x_2)]$	Solve for F , g & h
Zero-snap	\tilde{f} has numerical parameters \mathbf{p}	Replace p_i by 0
Integer snap	\tilde{f} has numerical parameters \mathbf{p}	Round p_i to integer
Rational snap	\tilde{f} has numerical parameters \mathbf{p}	Round p_i to fraction
Reoptimize	\tilde{f} has numerical parameters \mathbf{p}	Reoptimize \mathbf{p} to minimize inaccuracy

discover generalized symmetry without knowing h , thus only performing the time-consuming task of searching for an h satisfying equation (2) after determining that a solution exists. The Supplementary Material details how we numerically test for \mathbf{x}'' -independence of $\hat{\mathbf{v}}(\mathbf{x}', \mathbf{x}'')$.

Generalized additivity If f is a function of two variables, then we also test for *generalized additivity* (Figure 3, bottom right), where $f(x_1, x_2) = F[g(x_1) + h(x_2)]$. If we define the function

$$s(x_1, x_2) \equiv \frac{\partial f / \partial x_1}{\partial f / \partial x_2}, \quad \text{then} \quad s(x_1, x_2) = \frac{g'(x_1)}{h'(x_2)} \quad (3)$$

if f satisfies the generalized additivity property. In other words, we simply need to test if s is of the multiplicatively separable form $s(x_1, x_2) = a(x_1)b(x_2)$, and we do this using a variant of the separability test described in [25]. The Supplementary Material details how we perform this separability test numerically.

2.2 Robustness through recursive Pareto-optimality

As illustrated in Figure 1, the goal of our symbolic regression of a data set is to approximate $f(\mathbf{x})$ by functions $\tilde{f}(\mathbf{x})$ that are not only accurate, but also simple, in the spirit of Occam’s razor. As in [10], we seek functions that are *Pareto-optimal* in the sense of there being no other function that is both simpler and more accurate. We will adopt an information-theoretical approach and use bits of information to measure lack of both accuracy and simplicity.

For *accuracy*, we wish the vector $\boldsymbol{\varepsilon}$ of prediction errors $\varepsilon_i \equiv y_i - \tilde{f}(\mathbf{x}_i)$ to be small. We quantify this not by the mean-squared error $\langle \varepsilon_i^2 \rangle$ or max-error $\max |\varepsilon_i|$ as in [10, 25], but by the MEDL, the *mean error-description-length* $\langle L_d(\varepsilon_i) \rangle$ defined in Table 2. As argued in [27] and illustrated in Figure 4, this improves robustness to outliers. We analogously quantify *complexity* by the description length L_d defined as in [27], summarized in Table 2.

L_d can be viewed as a crude but computationally convenient approximation of the number of bits needed to describe each object, made differentiable where possible. We choose the precision floor $\epsilon \equiv 2^{-30} \sim 10^{-9}$. For function complexity, both input variables and mathematical functions (e.g., \cos and $+$) count toward n and k . For example, the classical kinetic energy formula has

Table 2: Complexity definitions

Object	Symbol	Description length L_d
Natural number	n	$\log_2 n$
Integer	m	$\log_2(1 + m)$
Rational number	m/n	$L_d(m) + L_d(n) = \log_2[(1 + m)n]$
Real number	r	$\log_+(\frac{r}{\epsilon}), \log_+(x) \equiv \frac{1}{2} \log_2(1 + x^2)$
Parameter vector	\mathbf{p}	$\sum_i L_d(p_i)$
Parametrized function	$f(\mathbf{x}; \mathbf{p})$	$L_d(\mathbf{p}) + k \log_2 n; n$ basis functions appear k times




Figure 4: When fitting a function (the right panel shows the example $\frac{ax^3}{e^{x/b}-1}$) to data with outliers, minimizing mean-squared-error (MSE) biases the curve toward the outliers (here finding $a \approx 0.89$, $b \approx 1.056$), whereas minimizing mean error description length (MEDL) ignores the outliers and recovers the correct answer $a = b = 1$. Left panel compares MSE and MEDL loss functions for $b = 1$.

$L_d("m \times v \times v/2^n") = L_d(2) + k \log_2 n = \log_2 3 + 6 \log_2 4 \approx 13.6$ bits, since the formula contains $n = 4$ basis functions (m, v, \times and $/$) used $k = 6$ times.

We wish to make the symbolic regression implementation of [25] more robust; it sometimes fails to discover the correct expression because of noise in the data or inaccuracies introduced by the neural network fitting. The neural network accuracy may vary strongly with \mathbf{x} , becoming quite poor in domains with little training data or when the network is forced to extrapolate rather than interpolate, and we desire a regression method robust to such outliers. We expect our insistence on Pareto-optimal functions in the information plane of Figure 1 to increase robustness, both because $\langle L_d(\epsilon_i) \rangle$ is robust (Figure 4) and because noise and systematic errors are unlikely to be predictable by a simple mathematical formula with small L_d . More broadly, minimization of total exact description length (which L_d crudely approximates) probably avoids the overfitting problem that plagues many alternative machine-learning strategies [28–30].

Speedup by recursive Pareto frontier composition When recursively symbolically regressing various modules (see Figure 2), we end up with a Pareto frontier of candidate functions for each one. If there are n_i functions on the i^{th} frontier, then combining them all would produce $\prod_i n_i$ candidates $\tilde{f}(\mathbf{x})$ for the original function $f(\mathbf{x})$. We speed up our algorithm by Pareto-pruning after each merge step: whenever two modules are combined (via composition or multiplication, say), the resulting $n_1 n_2$ functions are pruned by removing all functions that are Pareto-dominated by another function that is both simpler and more accurate. Pruning models on the Pareto frontier significantly reduces the number of models that need to be evaluated, since in typical scenarios, the number of Pareto-optimal points grows only logarithmically with the total number of points.

Robust speedup of brute-force graph search with hypothesis testing Our recursive reduction of regression mysteries into simpler ones terminates at the base case when the mystery function has only one variable and cannot be further modularized. As in [25], we subject these (and also all multivariate modules) to two solution strategies, polynomial fitting up to some degree (4 by default) and brute force search, and then add all candidates functions to the Pareto plane and prune

as above. The brute-force search would, if run forever, try all symbolic expressions by looping over ever-more-complex graphs (the middle panel of Figure 2 shows an example) and over function options for each node.

Our brute-force computation of the Pareto frontier simply tries all functions $f_k(\mathbf{x})$ ($k = 1, 2, \dots$) in order of increasing complexity $L_d(f_k)$ and keeps only those with lower mean error-description-length $d_k \equiv \frac{1}{N} \sum_{i=1}^N d_{ki}$ than the previous record holder, where $d_{ki} \equiv L_d[y_i - f_k(\mathbf{x}_i)]$. When instead fitting normalized gradient vectors $\widehat{\nabla} f$ as in Section 2.1, we define $d_{ki} \equiv L_d[1 - |\hat{\mathbf{y}}_i \cdot \widehat{\nabla} f_k(\mathbf{x}_i)|]$ to handle the sign ambiguity. The bad news is that computing d_k exactly is slow, requiring evaluation of $f_k(\mathbf{x}_i)$ for all N data points \mathbf{x}_i . The good news is that this is usually unnecessary, since for the vast majority of all candidate functions, it becomes obvious that they provide a poor fit after trying merely a handful of data points. We therefore accelerate the search via the following procedure. Before starting the loop over candidate functions, we sort the data points in random order to be able to interpret the numbers d_{ki} as random samples from a probability distribution whose mean is the sought-for d_k and whose standard deviation is σ_k . Let d_{k*} and σ_{k*} denote the corresponding quantities that were computed for the previous best-fit function we added to the Pareto frontier. We make the simplifying approximations that $\sigma_k = \sigma_{k*}$ and that all errors are uncorrelated, so that the loss estimate from the first m data points $\bar{d}_{km} \equiv \frac{1}{m} \sum_{i=1}^m d_{ki}$ has mean d_k and standard deviation σ_{k*}/\sqrt{m} . We now test our candidate function f_i on one data point at a time and reject it as soon as

$$z > \nu, \quad \text{where} \quad z \equiv \sqrt{m} \frac{\bar{d}_{km} - d_{k*}}{\sigma_{k*}}, \quad (4)$$

where ν is a hyperparameter that we can interpret as the “number of sigmas” we require to rule out a candidate function as viable when its average error exceeds the previous record holder. We find that $\nu = 10$ usually works well, generically requiring no more than a handful of evaluations m per candidate function asymptotically. We can further increase robustness by increasing ν at the price of longer runtime.

Speedup by greedy search of simplification options We do not *a priori* know which of the modular decompositions from Figure 3 are most promising, and recursively trying all combinations of them would involve trying exponentially many options. We therefore accelerate our algorithm with a greedy strategy where at each step we compare the decomposition in a unified way and try only the most accurate one — our runtime thus grows roughly linearly with n , the number of input variables. $f(\mathbf{x})$ stays constant along constant- h curves for generalized symmetry, simple symmetry (where $h(x, y) = x + y, x - y, xy$ or x/y) and generalized additivity (where $h(x, y) = a(x) + b(y)$). We thus test the accuracy of all such h -candidates by starting at a datapoint \mathbf{x}_i and computing an error $\epsilon_i \equiv f(\tilde{\mathbf{x}}_i) - f(\mathbf{x}_i)$ for some $\tilde{\mathbf{x}}_i$ satisfying $h(\tilde{\mathbf{x}}_i) = h(\mathbf{x}_i)$. For additive and multiplicative separability, we follow [25] by examining a rectangle in parameter space and predicting f at the fourth corner from the other three, defining ϵ_i as the mismatch. The supplementary material details how our test points are chosen.

After this greedy recursive process has terminated, we further improve the Pareto frontier in two ways. We first add models where rational numbers are replaced by reals and optimized by gradient descent to fit the data. We then add models with zero-snap, integer-snap and rational-snap from Table 1 applied to all real-valued parameters as described in [27], pruning all Pareto-dominated models after each step. For example, if there are 3 real-valued parameters, integer-snap generates 3 new models where the 1, 2 and 3 parameters closest to integers get rounded, respectively.

2.3 Leveraging normalizing flows to symbolic regress probability distributions

An important but more difficult symbolic regression problem is when the unknown function $f(\mathbf{x})$ is a probability distribution from which we have random samples \mathbf{x}_i rather than direct evaluations $y_i = f(\mathbf{x}_i)$. We tackle this by adding preceding the regression by a step that estimates $f(\mathbf{x})$. For this step, we use the popular *normalizing flow* technique [31–35], training an invertible neural network mapping $\mathbf{x} \mapsto \mathbf{x}' \equiv g(\mathbf{x})$ such that \mathbf{x}' has a multivariate normal distribution $n(\mathbf{x}')$ as illustrated in Figure 5. We then obtain our estimator $f_{\text{NN}}(\mathbf{x}) = n[g(\mathbf{x})]|J|$, where J is the Jacobian of g .

We find rational-quadratic neural spline flows (RQ-NSF) suitable for relatively low-dimensional applications due to their enhanced expressivity. Specifically, we used three steps of the RQ-NSF with RQ-NSF (C) coupling layers as described in [34], parametrized by three 16-neuron softplus layers,




Figure 5: A normalizing flow g maps samples from a probability distribution f (right) into a normal distribution (left), enabling an estimate (middle) of f , here illustrated for the $n = 2, l = 1, m = 0$ hydrogen orbital from Table 5.

Table 3: Robustness to noise

New robustness	10^{-5}	10^{-4}	10^{-3}	10^{-2}	10^{-1}	Total	
Old robustness	10^{-5}	0	1	2	2	0	5
	10^{-4}	0	1	3	5	12	21
	10^{-3}	0	0	5	6	24	35
	10^{-2}	0	0	0	2	37	39
Total	0	2	10	15	73	100	

trained for 50,000 epochs with the Adam optimizer. The learning rate was initialized to 3×10^{-4} and halved every time the test loss failed to improve for 2500 epochs.

3 Results

We now turn to quantifying the performance of our method with numerical experiments, comparing it with that of [25] which recently exceeded the previous state-of-the-art performance of [10]. To quantify robustness to noise, we add Gaussian noise of standard deviation 10^r to y_i and determine the largest integer $r < 0$ for which the method successfully discovers the correct mystery function $f(\mathbf{x})$. As seen in Table 3, our method solves 73 of 100 the baseline problems from the Feynman Symbolic Regression Database [25] with $r = -1$, and is typically 1-3 orders of magnitude more robust than that of [25]. Crudely speaking, we found that adding progressively more noise shifted the most accurate formula straight upward in the Pareto plane (Figure 1) until it no longer provided any accuracy gains compared with simpler approximations.

To quantify the ability of our method to discover more complex equations, we reran it on all 17 mysteries that [25] tackled and failed to solve. We also tested a dozen new mysteries exhibiting various forms of graph modularity (see Table 4) that were all chosen before any of them were tested. Allowing at most two hours of run-time, the method of [25] solved equations 5, 12, ..., 15, whereas our new method solved them all, as well as four of the outstanding mysteries from [25] (rows 1-4). For these first four, our method got the numerical parameters in the right ballpark with rational approximations, then discovered their exact values through gradient descent.

To quantify the ability of our method to discover probability distributions, we tested it on samples from the ten distributions in Table 5. As seen in the table, 80% were solved, requiring between 10^2 and 10^5 samples \mathbf{x}_i . The flows trained in about 20 minutes on one CPU, scaling roughly linearly with sample size and number of network weights. The Supplementary Material details failure modes.

4 Conclusions

We have presented a symbolic regression method that exploits neural networks, graph modularity, hypothesis testing and normalizing flows. It improves state-of-the-art performance both by being

Table 4: Test equations exhibiting translational symmetry $h = x \pm y$ (T), scaling symmetry $h = x/y$ (S), product symmetry $h = xy$ (P), generalized symmetry (G), multiplicative separability (M), compositionality (C) and generalized additivity (A).

Equation	Symmetries
1 $\delta = -5.41 + 4.9 \frac{\alpha - \beta + \gamma / \chi}{3\chi}$	TC
2 $\chi = 0.23 + 14.2 \frac{\alpha + \beta}{3\gamma}$	TS
3 $\beta = 213.80940889 (1 - e^{-0.54723748542\alpha})$	
4 $\delta = 6.87 + 11\sqrt{\alpha\beta\gamma}$	P
5 $V = [R_1^{-1} + R_2^{-1} + R_3^{-1} + R_4^{-1}]^{-1} I_0 \cos \omega t$ (Parallel resistors)	PGSM
6 $I_0 = \frac{V_0}{\sqrt{R^2 + (\omega L - \frac{1}{\omega C})^2}}$ (RLC circuit)	MG
7 $I = \frac{V_0 \cos \omega t}{\sqrt{R^2 + (\omega L - \frac{1}{\omega C})^2}}$ (RLC circuit)	MG
8 $V_2 = (\frac{R_2}{R_1 + R_2} - \frac{R_x}{R_x + R_3}) V_1$ (Wheatstone bridge)	SGMA
9 $v = c \frac{(v_1 + v_2 + v_3)/c + v_1 v_2 v_3 / c^3}{1 + (v_1 v_2 + v_1 v_3 + v_2 v_3)/c^2}$ (Velocity addition)	AG
10 $v = c \frac{(v_1 + v_2 + v_3 + v_4)/c + (v_2 v_3 v_4 + v_1 v_3 v_4 + v_1 v_2 v_4 + v_1 v_2 v_3)/c^3}{1 + (v_1 v_2 + v_1 v_3 + v_1 v_4 + v_2 v_3 + v_2 v_4 + v_3 v_4)/c^2 + v_1 v_2 v_3 v_4 / c^4}$ (Velocity addition)	GA
11 $z = (x^4 + y^4)^{1/4}$ (L_4 -norm)	AC
12 $w = xyz - z\sqrt{1-x^2}\sqrt{1-y^2} - y\sqrt{1-x^2}\sqrt{1-z^2} - x\sqrt{1-y^2}\sqrt{1-z^2}$	GA
13 $z = \frac{xy + \sqrt{1-x^2-y^2+x^2y^2}}{y\sqrt{1-x^2-x}\sqrt{1-y^2}}$	A
14 $z = y\sqrt{1-x^2} + x\sqrt{1-y^2}$	A
15 $z = xy - \sqrt{1-x^2}\sqrt{1-y^2}$	A
16 $r = \frac{a}{\cot(\alpha/2) + \cot(\beta/2)}$ (Incircle)	GMAC

Table 5: Probability distributions and number of samples N required to discover them

Distribution Name	Probability distribution	N
Laplace distribution	$\frac{1}{2} e^{- x }$	10^2
Beta distribution ($\alpha = 0.5, \beta = 0.5$)	$\frac{1}{\pi} \frac{1}{\sqrt{x(1-x)}}$	10^4
Beta distribution ($\alpha = 5, \beta = 2$)	$30x^4(1-x)$	10^4
Harmonic oscillator ($n = 2, \frac{m\omega}{\hbar} = 1$)	$\frac{2}{\sqrt{\pi}} x^2 e^{-x^2}$	10^5
Sinc diffraction pattern	$\frac{1}{\pi} \left(\frac{\sin x}{x}\right)^2$	10^4
2D normal distribution (correlated)	$\frac{1}{\sqrt{3\pi}} e^{-\frac{2}{3}(x^2 - xy + y^2)}$	10^3
2D harmonic oscillator ($n = 2, m = 1, \frac{m\omega}{\hbar} = 1$)	$\frac{2}{\pi} x^2 e^{-x^2 - y^2}$	10^5
Hydrogen orbital ($n = 1, l = 0, m = 0$)	$\frac{1}{4} e^{-2r}$	10^3
Hydrogen orbital ($n = 2, l = 1, m = 0$)	$\frac{1}{16} r^2 e^{-r} \cos^2 \theta$	-
Hydrogen orbital ($n = 3, l = 1, m = 0$)	$\frac{1}{729} r^2 \left(4 - \frac{2r}{3}\right)^2 e^{-\frac{2r}{3}} \cos^2 \theta$	-

more robust towards noise and by solving harder problems, including symbolic density estimation. These core ideas can enable further improvements. For example, gradients can reveal more types of graph modularity than the Figure 3 examples that we exploited; additional simplification strategies can be included in the Pareto-optimal recursion; and flow-based regression can be used for regularized density estimation from sparse high-dimensional data. Larger and more challenging collections of science-based equations are needed to benchmark and inspire improved algorithms.

Pareto-optimal symbolic regression has the power to not only discover exact formulas, but also approximate ones that are useful for being both accurate and simple. The mainstream view is that all known science formulas are such approximations. We live in a golden age of research with ever-larger datasets produced by both experiments and numerical computations, and we look forward to a future when symbolic regression is as ubiquitous as linear regression is today, helping us better understand the relations hidden in these datasets.

Acknowledgments and Disclosure of Funding

The authors wish to thank Philip Tegmark for helpful comments, and the Center for Brains, Minds, and Machines (CBMM) for hospitality.

Funding: This work was supported by The Casey and Family Foundation, the Ethics and Governance of AI Fund, the Foundational Questions Institute, the Rothberg Family Fund for Cognitive Science and the Templeton World Charity Foundation, Inc.

Competing interests: The authors declare that they have no competing interests.

Data and materials availability: Data and code will be publicly released when this paper is accepted for publication.

Broader Impact

Who may benefit from this research

Our research presumably has quite broad impact, since discovery of mathematical patterns in data is a central problem across the natural and social sciences. Given the ubiquity of *linear* regression in research, one might expect that there will be significant benefits to a broad range of researchers also from more general symbolic regression once freely available algorithms get sufficiently good.

Who may be put at disadvantage from this research

Although it is possible that some numerical modelers could get their jobs automated away by symbolic regression, we suspect that the main effect of our method, and future tools building on it, will instead be that these people will simply discover better models than today.

Risk of bias, failure and other negative outcomes

Pareto-optimal symbolic regression can be viewed as an extreme form of lossy data compression that uncovers the simplest possible model for any given accuracy. To the extent that overfitting can exacerbate bias, such model compression is expected to help. Moreover, since our method produces closed-form mathematical formulas that have excellent interpretability compared to black-box neural networks, they make it easier for humans to interpret the computation and pass judgement on whether it embodies unacceptable bias. This interpretability also reduces failure risk.

Another risk is automation bias, whereby people overly trust a formula from symbolic regression when they extrapolate it into an untested domain. This could be exacerbated if symbolic regression promotes scientific laziness and enfeeblement, where researchers fit phenomenological models instead of doing the work of building models based on first principles. Symbolic regression should inform but not replace traditional scientific discovery.

Although the choice of basis functions biases the discoverable function class, our method is agnostic to basis functions as long as they are mostly differentiable.

The potentially greatest risk associated with this work does not stem from it failing but from it succeeding: accelerated progress in symbolic regression, modularity discovery and its parent discipline program synthesis could hasten the arrival of artificial general intelligence, which some authors have argued that humanity still lacks the tools to manage safely [5]. On the other hand, our work may help accelerate research on intelligible intelligence more broadly, and powerful future artificial intelligence is probably safer if we understand aspects of how it works than if it is an inscrutable black box.

Supplementary material

Below we provide additional technical details about how we implement our method and numerical experiments.

A Testing for generalized symmetry

We showed that generalized symmetry can be revealed by $\widehat{\mathbf{v}}(\mathbf{x}', \mathbf{x}'')$ being independent of \mathbf{x}'' . We will now describe how we test for such \mathbf{x}'' -independence numerically. Given a point $\mathbf{x}_i \in \mathbb{R}^k$ from our data set, we compute a set of normalized gradients $\widehat{\mathbf{v}}_j \equiv \widehat{\mathbf{v}}_i(\mathbf{x}'_i, \mathbf{x}''_j)$, where $\mathbf{x}''_j \in \mathbb{R}^{n-k}$ correspond to a sample of m other data points, and quantify the variation between them by the quantity

$$V(\mathbf{x}) \equiv 1 - \max_{|\mu|=1} \frac{1}{m} \sum_{j=1}^m (\widehat{\mu} \cdot \widehat{\mathbf{v}}_j)^2 = 1 - \max_{|\mu|=1} \widehat{\mu}^t \mathbf{V} \widehat{\mu}, \quad \text{where} \quad \mathbf{V} \equiv \frac{1}{m} \sum_{j=1}^m \widehat{\mathbf{v}}_j \widehat{\mathbf{v}}_j^t. \quad (5)$$

We can intuitively interpret the optimal $\widehat{\mu}$ as maximally aligned with the vectors $\widehat{\mathbf{v}}_j$ up to a sign. Equation (5) implies that our variation measure V is simply one minus the smallest eigenvalue of \mathbf{V} , so V ranges from 0 when all $\widehat{\mathbf{v}}_j$ are identical to $1 - \frac{1}{m}$ when all eigenvalues are equal (equal to $1/m$, since $\text{tr} \mathbf{V} = 1$). As illustrated in Figure 6, we compute $V(\mathbf{x}_i)$ for each subset of up to n_g input variables, and select the subset with the smallest median $V(\mathbf{x}_i)$ as the most promising generalized symmetry candidate. In our numerical experiments, we set the hyperparameter $n_g = 3$ to save time, since we do not wish to consider all 2^n subsets for large n .




Figure 6: Distribution of $V(\mathbf{x}_i)$ for the function from Figure 2, revealing that evidence for the generalized symmetry $f(x, y, z) = g[h(x, y), z]$ (shaded distribution) is stronger than for $f(x, y, z) = g[h(x, z), y]$ (blue curve) or $f(x, y, z) = g[h(y, z), x]$ (red curve). The curves are shown slightly smoothed for clarity.

B Testing for generalized additivity

We showed that generalized additivity holds when the function $s(x_1, x_2)$ from Equation (3) is multiplicatively separable. We will now describe how we test for such separability numerically. If $s(x_1, x_2)$ being multiplicatively separable is equivalent to $f(x_1, x_2) \equiv \ln s(x_1, x_2)$ being additively separable. We numerically quantify this by the and we test the function $\ln s_{NN}(x_1, x_2)$ for additive separability using the normalized score S defining

$$S[f] = \frac{|f_{,xy}|^2}{|f_{,xx}f_{,yy}| + |f_{,xy}|^2}. \quad (6)$$

It is easy to see that $S[f] = 0$ if f is additively separable functions f , and $S[f] > 0$ otherwise. If the median value of S over all points \mathbf{x}_i in the dataset, we take this as evidence for generalized additivity in the dataset and proceed as below. It is important to use smooth (not, *e.g.*, ReLU) activation functions in the activation function for this derivative-based test to be useful.

If this property holds, then we recursively apply our algorithm to the two new 1-dimensional symbolic regression problems of discovering $a(x_1)$ and $b(x_2)$. If this succeeds and we are able to discover the

functions $g(x_1)$ and $h(x_2)$ by symbolically integrating our solutions $g' = a$ and $h' = 1/b$, then we have reduced the original problem to the same state as when we found compositionality above, now with $h(x_1, x_2) = g(x_1) + h(x_2)$. Just as in that case, we simply replace the variables \mathbf{x} in the original mystery data table by the single variable $h(\mathbf{x})$ and recursively apply our AI Feynman algorithm to the new 1-dimensional symbolic regression problem of discovering how f depends on h .

If we have determined that generalized additivity holds but the above aforementioned method for discovering $g(x_1) + h(x_2)$ fails, we make a second attempt by training a neural network of the modular form $f_{\text{NN}}(x_1, x_2) = F[g(x_1) + h(x_2)]$ to fit the data. If this succeeds, we then recursively apply our AI Feynman algorithm to the three new 1-dimensional symbolic regression problems of discovering F , g and h .

C Further details on success and failure modes

Our paper reported which symbolic regression problems our method succeeded and failed on. Here we add specifics on how these successes and failures occurred.

Success definition Given a data set $\{x_1, \dots, x_n, y\}$, we use 90% of the data to compute a Pareto-optimal set of candidate functions $\tilde{f}_i(\mathbf{x})$, then rank them based on their MEDL accuracy on the held-back 10% of the data. We count our method as successful only if the top-ranked function matches the true $f(\mathbf{x})$ exactly, or, if the definition of f involves irrational numerical parameters, if these parameters are recovered to better than 0.01% relative accuracy.

We considered an equation solved even if the top solution was not in the exact form presented in our tables, but Mathematically equivalent. For example, our method predicted that Equation (12) in Table 4 was $w = \cos[\arccos(x) + \arccos(y) + \arccos(z)]$, which is mathematically equivalent within the domain of our provided data set, where $x, y, z \in [-1, 1]$.

For the problem of density estimation from samples, our goal was to obtain the correct normalized probability distributions. The candidate functions on the Pareto-frontier were therefore discarded unless they were non-negative and normalizable. The surviving candidates then normalized to integrate to unity by symbolic/numerical integration to obtain the appropriate normalization constant, and quality-ranked by the surprisal loss function

$$L_i = - \sum \log \tilde{f}_i(\mathbf{x}_k)$$

evaluated on the held-back test data.

Success examples Tables 1 and 2 below show the highest noise level allowing allowing each of the 100 equations from the Feynman Database for Symbolic regression to be solved in the original paper analyzing it and in the present paper.

For many of the solved equations, the modularity discovery had to be used multiple times in order for the correct equation to be discovered, reflecting the power of the recursive algorithm. For example, for the quadruple velocity addition equation in Table 4, generalized symmetry was exploited twice. First, the code discovered that the first two velocities only enter in the combination $\frac{v_1+v_2}{1+v_1v_2}$, and these two variables were replaced by a new variable v_{12} . The same method then discovered that v_{12} and v_3 only enter in that same combination $\frac{v_{12}+v_3}{1+v_{12}v_3}$, and thus the initial 3 variables v_1, v_2 and v_3 were replaced by a single variable v_{123} . Now the remaining equation had only 2 variables left, and was solved by brute force. In principle, this recursive method can be used to discover relativistic addition of an arbitrary number of velocities, by reducing the number of variables by one at each step.

Failure examples As mentioned in the main text, numerous equations remained unsolved, motivating further work. In some cases, the main reason for this failure was that the form of the equation did not allow our method to break it into small enough pieces. For example, for the equation

$$\alpha^3 \exp(-\alpha) \cos(\alpha) \sin(\alpha) (\sin(\alpha)^2 \cos(\alpha) - 1) (\beta - 5),$$

our algorithm is able to discover the multiplicative separability into terms including only α and only β . However, the remaining α -term is too complicated to be solved in a reasonable amount of time by the brute force code, and none of the graph modularity methods apply because they only help for functions of more than one variable.

For other equations, our method fails but not irreparably. For example, for the function

$$22 - 4.2(\cos(\alpha) - \tan(\beta)) \tanh(\gamma) / \sin(\chi),$$

our code is able to discover that γ and χ can be separated from the rest of the equation. However, given that we allow the brute force code to run for only a minute at each iteration, the expression $\tanh(\gamma) / \sin(\chi)$ is not discovered, mainly because we did not include \tanh as one of the functions used, so the brute force would have to write that as $\frac{e^{-2x}-1}{e^{2x}+1}$. By allowing the code to run for longer and experimenting with additional basis functions (such as \tanh), it is likely that the code would solve this and several other mysteries that we reported as failures.

Fitting a normalizing flow for density estimation presents a ‘real-world’ test of the symbolic regression code, which needs to cope with flow errors. We demonstrated the robustness of the code to more challenging errors (i.e. not i.i.d. Gaussian); some of the failure modes are investigated and expanded on below.

It is worth noting that our definition of complexity is dependent on the chosen set of operations and does not always match our intuition. For example, in fitting the probability distribution

$$p(r, \theta) = \frac{1}{16} r^2 e^{-r} \cos^2 \theta$$

of electron positions in the $n = 2, l = 1, m = 0$ hydrogen orbital, solutions with θ -dependence $\cos(\cos(\theta))$ are preferred over $\cos^2 \theta$. This is because, up to additive and multiplicative prefactors, the two formulas differ by at most approximately 2×10^{-2} over our parameter range, but given a set of operations that includes only $\{\times, \cos\}$ denoted by "*" and "C" respectively in reverse Polish notation, $\cos(\cos(\theta))$ (encoded as "xCCC") is simpler than $\cos^2 \theta$ (encoded as "xCC*"). In the presence of the imprecisions introduced by the normalizing flow, we were unable to perform the density estimation a level at which the accuracy for the correct $\cos^2 \theta$ was preferred over the simpler alternative.

Furthermore, more interpretable approximations (e.g. Taylor expansions) are not always favored by our definition of complexity. For example, in Figure 1, the unfamiliar solution

$$mc^2 \left(\frac{1}{\cos v/c} - 1 \right)$$

intermediate to the more familiar $mv^2/2$ and $mc^2 \left(\frac{1}{\sqrt{1-v^2/c^2}} - 1 \right)$ can be understood as a fourth-order approximation about $v = 0$ of the exact formula. Specifically, $mc^2 \left(\frac{1}{\sqrt{1-v^2/c^2}} - 1 \right) = \frac{mv^2}{2} + \frac{3mv^4}{8c^2} + O(v^6)$, and $mc^2 \left(\frac{1}{\cos v/c} - 1 \right) = \frac{mv^2}{2} + \frac{5mv^4}{24c^2} + O(v^6)$. The Taylor expansions themselves are not preferred for reasons of complexity.

Feynman eq.	Equation	Old Noise tolerance	New Noise tolerance
I.6.20a	$f = e^{-\theta^2/2}/\sqrt{2\pi}$	10^{-2}	10^{-1}
I.6.20	$f = e^{-\frac{\theta^2}{2\sigma^2}}/\sqrt{2\pi\sigma^2}$	10^{-4}	10^{-2}
I.6.20b	$f = e^{-\frac{(\theta-\theta_1)^2}{2\sigma^2}}/\sqrt{2\pi\sigma^2}$	10^{-4}	10^{-2}
I.8.14	$d = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$	10^{-4}	10^{-1}
I.9.18	$F = \frac{Gm_1m_2}{(x_2-x_1)^2+(y_2-y_1)^2+(z_2-z_1)^2}$	10^{-5}	10^{-3}
I.10.7	$m = \frac{m_0}{\sqrt{1-\frac{v^2}{c^2}}}$	10^{-4}	10^{-2}
I.11.19	$A = x_1y_1 + x_2y_2 + x_3y_3$	10^{-3}	10^{-1}
I.12.1	$F = \mu N_n$	10^{-3}	10^{-1}
I.12.1a	$K = \frac{1}{2}m(v^2 + u^2 + w^2)$	10^{-4}	10^{-1}
I.12.2	$F = \frac{q_1q_2}{4\pi\epsilon r^2}$	10^{-2}	10^{-1}
I.12.4	$U = \frac{q_1}{4\pi\epsilon r^2}$	10^{-2}	10^{-1}
I.12.5	$F = q_2E_f$	10^{-2}	10^{-1}
I.12.11	$F = q(E_f + Bv \sin \theta)$	10^{-3}	10^{-1}
I.13.12	$U = Gm_1m_2(\frac{1}{r_2} - \frac{1}{r_1})$	10^{-4}	10^{-1}
I.14.3	$U = mgz$	10^{-2}	10^{-1}
I.14.4	$U = \frac{k_{spring}x^2}{2}$	10^{-2}	10^{-1}
I.15.3x	$x_1 = \frac{x-ut}{\sqrt{1-u^2/c^2}}$	10^{-3}	10^{-3}
I.15.3t	$t_1 = \frac{t-ux/c^2}{\sqrt{1-u^2/c^2}}$	10^{-4}	10^{-3}
I.15.1	$p = \frac{m_0v}{\sqrt{1-v^2/c^2}}$	10^{-4}	10^{-1}
I.16.6	$v_1 = \frac{u+v}{1+uv/c^2}$	10^{-3}	10^{-2}
I.18.4	$r = \frac{m_1r_1+m_2r_2}{m_1+m_2}$	10^{-2}	10^{-1}
I.18.12	$\tau = rF \sin \theta$	10^{-3}	10^{-1}
I.18.14	$L = mrv \sin \theta$	10^{-3}	10^{-1}
I.24.6	$E = \frac{1}{4}m(\omega^2 + \omega_0^2)x^2$	10^{-4}	10^{-1}
I.25.13	$V_e = \frac{q}{C}$	10^{-2}	10^{-1}
I.26.2	$\theta_1 = \arcsin(n \sin \theta_2)$	10^{-2}	10^{-1}
I.27.6	$f_f = \frac{1}{\frac{1}{d_1} + \frac{1}{d_2}}$	10^{-2}	10^{-1}
I.29.4	$k = \frac{\omega}{c}$	10^{-2}	10^{-1}
I.29.16	$x = \sqrt{x_1^2 + x_2^2 - 2x_1x_2 \cos(\theta_1 - \theta_2)}$	10^{-4}	10^{-3}
I.30.3	$I_* = I_{*0} \frac{\sin(n\theta/2)}{\sin(\theta/2)}$	10^{-3}	10^{-3}
I.30.5	$\theta = \arcsin(\frac{\lambda}{nd})$	10^{-3}	10^{-1}
I.32.5	$P = \frac{q^2a^2}{6\pi\epsilon^3}$	10^{-2}	10^{-1}
I.32.17	$P = (\frac{1}{2}\epsilon c E_f^2)(8\pi r^2/3)(\omega^4/(\omega^2 - \omega_0^2)^2)$	10^{-4}	10^{-3}
I.34.8	$\omega = \frac{qvB}{p}$	10^{-2}	10^{-1}
I.34.10	$\omega = \frac{1+v/c}{1-v/c} \omega_0$	10^{-3}	10^{-2}
I.34.14	$\omega = \frac{1+v/c}{\sqrt{1-v^2/c^2}} \omega_0$	10^{-3}	10^{-3}
I.34.27	$E = \hbar\omega$	10^{-2}	10^{-1}
I.37.4	$I_* = I_1 + I_2 + 2\sqrt{I_1I_2} \cos \delta$	10^{-3}	10^{-2}
I.38.12	$r = \frac{4\pi\epsilon\hbar^2}{mq^2}$	10^{-2}	10^{-1}
I.39.10	$E = \frac{3}{2}p_F V$	10^{-2}	10^{-1}
I.39.11	$E = \frac{1}{\gamma-1} p_F V$	10^{-3}	10^{-1}
I.39.22	$P_F = \frac{nk_b T}{V} \frac{mgx}{k_b T}$	10^{-4}	10^{-1}
I.40.1	$n = n_0 e^{-\frac{mgx}{k_b T}}$	10^{-2}	10^{-1}
I.41.16	$L_{rad} = \frac{\hbar\omega^3}{\pi^2 c^2 (e^{\frac{\hbar\omega}{k_b T}} - 1)}$	10^{-5}	10^{-4}
I.43.16	$v = \frac{\mu_{drift} q V_e}{d}$	10^{-2}	10^{-1}
I.43.31	$D = \mu_e k_b T$	10^{-2}	10^{-1}
I.43.43	$\kappa = \frac{1}{\gamma-1} \frac{k_b v}{A}$	10^{-3}	10^{-1}
I.44.4	$E = nk_b T \ln(\frac{V_2}{V_1})$	10^{-3}	10^{-1}
I.47.23	$c = \sqrt{\frac{\gamma p F}{\rho}}$	10^{-2}	10^{-1}
I.48.2	$E = \frac{mc^2}{\sqrt{1-v^2/c^2}}$	10^{-5}	10^{-3}
I.50.26	$x = x_1[\cos(\omega t) + \alpha \cos(\omega t)^2]$	10^{-2}	10^{-1}

Table 6: Tested Equations, part 1

Feynman eq.	Equation	Old Noise tolerance	New Noise tolerance
II.2.42	$P = \frac{\kappa(T_2 - T_1)A}{r^d}$	10^{-3}	10^{-1}
II.3.24	$F_E = \frac{P^d}{4\pi r^2}$	10^{-2}	10^{-1}
II.4.23	$V_e = \frac{q}{4\pi\epsilon r}$	10^{-2}	10^{-1}
II.6.11	$V_e = \frac{1}{4\pi\epsilon} \frac{p_d \cos(\theta)}{r^2}$	10^{-3}	10^{-1}
II.6.15a	$E_f = \frac{3}{4\pi\epsilon} \frac{p_d z}{r^5} \sqrt{x^2 + y^2}$	10^{-3}	10^{-2}
II.6.15b	$E_f = \frac{3}{4\pi\epsilon} \frac{p_d}{r^3} \cos\theta \sin\theta$	10^{-2}	10^{-2}
II.8.7	$E = \frac{3}{5} \frac{q^2}{4\pi\epsilon d}$	10^{-2}	10^{-1}
II.8.31	$E_{den} = \frac{\epsilon E_f^2}{2}$	10^{-2}	10^{-1}
II.10.9	$E_f = \frac{\sigma_{den}}{\epsilon} \frac{1}{1+\chi}$	10^{-2}	10^{-1}
II.11.3	$x = \frac{q E_f}{m(\omega_0^2 - \omega^2)}$	10^{-3}	10^{-2}
II.11.7	$n = n_0 \left(1 + \frac{p_d E_f \cos\theta}{k_b T}\right)$	10^{-2}	10^{-1}
II.11.20	$P_* = \frac{n_\rho p_d^2 E_f}{3k_b T}$	10^{-3}	10^{-1}
II.11.27	$P_* = \frac{n\alpha}{1-n\alpha/3} \epsilon E_f$	10^{-3}	10^{-1}
II.11.28	$\theta = 1 + \frac{n\alpha}{1-(n\alpha/3)}$	10^{-4}	10^{-2}
II.13.17	$B = \frac{1}{4\pi\epsilon c^2} \frac{2I}{r}$	10^{-2}	10^{-1}
II.13.23	$\rho_c = \frac{\rho_{c0}}{\sqrt{1-v^2/c^2}}$	10^{-4}	10^{-2}
II.13.24	$j = \frac{\rho_{c0} v}{\sqrt{1-v^2/c^2}}$	10^{-4}	10^{-1}
II.15.4	$E = -\mu_M B \cos\theta$	10^{-3}	10^{-1}
II.15.5	$E = -p_d E_f \cos\theta$	10^{-3}	10^{-1}
II.21.32	$V_e = \frac{q}{4\pi\epsilon r(1-v/c)}$	10^{-3}	10^{-1}
II.24.17	$k = \sqrt{\frac{\omega^2}{c^2} - \frac{\pi^2}{d^2}}$	10^{-5}	10^{-2}
II.27.16	$F_E = \epsilon c E_f^2$	10^{-2}	10^{-1}
II.27.18	$E_{den} = \epsilon E_f^2$	10^{-2}	10^{-1}
II.34.2a	$I = \frac{qv}{2\pi r}$	10^{-2}	10^{-1}
II.34.2	$\mu_M = \frac{qvr}{2}$	10^{-2}	10^{-1}
II.34.11	$\omega = \frac{g q \dot{B}}{2m}$	10^{-4}	10^{-1}
II.34.29a	$\mu_M = \frac{qh}{4\pi m}$	10^{-2}	10^{-1}
II.34.29b	$E = \frac{g \mu_M B J_z}{h}$	10^{-4}	10^{-1}
II.35.18	$n = \frac{n_0}{\exp(\mu_m B/(k_b T)) + \exp(-\mu_m B/(k_b T))}$	10^{-2}	10^{-2}
II.35.21	$M = n_\rho \mu_M \tanh\left(\frac{\mu_M B}{k_b T}\right)$	10^{-4}	10^{-4}
II.36.38	$f = \frac{\mu_m B}{k_b T} + \frac{\mu_m \alpha M}{\epsilon c^2 k_b T}$	10^{-2}	10^{-1}
II.37.1	$E = \mu_M (1 + \chi) B$	10^{-3}	10^{-1}
II.38.3	$F = \frac{Y A x}{d}$	10^{-3}	10^{-1}
II.38.14	$\mu_S = \frac{Y}{2(1+\sigma)}$	10^{-3}	10^{-1}
III.4.32	$n = \frac{1}{e^{\frac{\hbar\omega}{k_b T}} - 1}$	10^{-3}	10^{-2}
III.4.33	$E = \frac{\hbar\omega}{e^{\frac{\hbar\omega}{k_b T}} - 1}$	10^{-3}	10^{-3}
III.7.38	$\omega = \frac{2\mu_M B}{\hbar}$	10^{-2}	10^{-1}
III.8.54	$p_\gamma = \sin\left(\frac{Et}{\hbar}\right)^2$	10^{-3}	10^{-3}
III.9.52	$p_\gamma = \frac{\frac{p_d E_f t}{\hbar} \sin((\omega - \omega_0)t/2)^2}{((\omega - \omega_0)t/2)^2}$	10^{-3}	10^{-1}
III.10.19	$E = \mu_M \sqrt{B_x^2 + B_y^2 + B_z^2}$	10^{-4}	10^{-1}
III.12.43	$L = n\hbar$	10^{-3}	10^{-1}
III.13.18	$v = \frac{2Ed^2 k}{\hbar}$	10^{-4}	10^{-1}
III.14.14	$I = I_0 \left(e^{\frac{qV_e}{k_b T}} - 1\right)$	10^{-3}	10^{-1}
III.15.12	$E = 2U(1 - \cos(kd))$	10^{-4}	10^{-1}
III.15.14	$m = \frac{\hbar^2}{2Ed^2}$	10^{-2}	10^{-1}
III.15.27	$k = \frac{2\pi\alpha}{nd}$	10^{-3}	10^{-1}
III.17.37	$f = \beta(1 + \alpha \cos\theta)$	10^{-3}	10^{-1}
III.19.51	$E = \frac{-mq^4}{2(4\pi\epsilon)^2 \hbar^2} \frac{1}{n^2}$	10^{-5}	10^{-2}
III.21.20	$j = \frac{-\rho_{c0} q A_{vec}}{m}$	10^{-2}	10^{-1}

Table 7: Tested Equations, part 2.

References

- [1] Stuart Russell, Daniel Dewey, and Max Tegmark. Research priorities for robust and beneficial artificial intelligence. *Ai Magazine*, 36(4):105–114, 2015.
- [2] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- [3] Margaret Boden, Joanna Bryson, Darwin Caldwell, Kerstin Dautenhahn, Lilian Edwards, Sarah Kember, Paul Newman, Vivienne Parry, Geoff Pegman, Tom Rodden, et al. Principles of robotics: regulating robots in the real world. *Connection Science*, 29(2):124–129, 2017.
- [4] Viktoriya Krakovna and Finale Doshi-Velez. Increasing the interpretability of recurrent neural networks using hidden markov models. *arXiv preprint arXiv:1606.05320*, 2016.
- [5] Stuart Russell. *Human compatible: Artificial intelligence and the problem of control*. Penguin, 2019.
- [6] James P Crutchfield and Bruce S McNamara. Equation of motion from a data series. *Complex systems*, 1(417-452):121, 1987.
- [7] Saso Dzeroski and Ljupco Todorovski. Discovering dynamics: from inductive logic programming to machine discovery. *Journal of Intelligent Information Systems*, 4(1):89–108, 1995.
- [8] Elizabeth Bradley, Matthew Easley, and Reinhard Stolle. Reasoning about nonlinear system identification. *Artificial Intelligence*, 133(1-2):139–188, 2001.
- [9] Pat Langley, Dileep George, Stephen D Bay, and Kazumi Saito. Robust induction of process models from time-series data. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 432–439, 2003.
- [10] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *science*, 324(5923):81–85, 2009.
- [11] Randall K McRee. Symbolic regression using nearest neighbor indexing. In *Proceedings of the 12th annual conference companion on Genetic and evolutionary computation*, pages 1983–1990. ACM, 2010.
- [12] Dominic P Searson, David E Leahy, and Mark J Willis. Gptips: an open source genetic programming toolbox for multigene symbolic regression. In *Proceedings of the International multiconference of engineers and computer scientists*, volume 1, pages 77–80. IMECS Hong Kong, 2010.
- [13] Renáta Dubčáková. Eureqa: software review. *Genetic programming and evolvable machines*, 12(2):173–178, 2011.
- [14] Sean Stijven, Wouter Minnebo, and Katya Vladislavleva. Separating the wheat from the chaff: on feature selection and feature importance in regression random forests and symbolic regression. In *Proceedings of the 13th annual conference companion on Genetic and evolutionary computation*, pages 623–630. ACM, 2011.
- [15] Michael D Schmidt, Ravishankar R Vallabhajosyula, Jerry W Jenkins, Jonathan E Hood, Abhishek S Soni, John P Wikswo, and Hod Lipson. Automated refinement and inference of analytical models for metabolic networks. *Physical biology*, 8(5):055011, 2011.
- [16] Christopher Hillar and Friedrich Sommer. Comment on the article "distilling free-form natural laws from experimental data". *arXiv preprint arXiv:1210.7273*, 2012.
- [17] Bryan C Daniels and Ilya Nemenman. Automated adaptive inference of phenomenological dynamical models. *Nature communications*, 6:8133, 2015.
- [18] Pat Langley and Adam Arvay. Heuristic induction of rate-based process models. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [19] Ignacio Arnaldo, Una-May O’Reilly, and Kalyan Veeramachaneni. Building predictive models via feature synthesis. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, pages 983–990, 2015.
- [20] 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.

- [21] Matthew Guzdial, Boyang Li, and Mark O Riedl. Game engine learning from video. In *IJCAI*, pages 3707–3713, 2017.
- [22] Markus Quade, Markus Abel, J Nathan Kutz, and Steven L Brunton. Sparse identification of nonlinear dynamics for rapid model recovery. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28(6):063116, 2018.
- [23] Maciej Koch-Janusz and Zohar Ringel. Mutual information, neural networks and the renormalization group. *Nature Physics*, 14(6):578, 2018.
- [24] Weiwei Kong, Christopher Liaw, Aranyak Mehta, and D Sivakumar. A new dog learns old tricks: RL finds classic optimization algorithms. *ICLR*, 2018.
- [25] Silviu-Marian Udrescu and Max Tegmark. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16):eaay2631, 2020.
- [26] Jiechun Liang and Xi Zhu. Phillips-inspired machine learning for band gap and exciton binding energy prediction. *The journal of physical chemistry letters*, 10(18):5640–5646, 2019.
- [27] Tailin Wu and Max Tegmark. Toward an artificial intelligence physicist for unsupervised learning. *Physical Review E*, 100(3):033311, 2019.
- [28] J. Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465 – 471, 1978.
- [29] Marcus Hutter. A theory of universal artificial intelligence based on algorithmic complexity. *arXiv preprint cs/0004001*, 2000.
- [30] Peter D Grünwald, In Jae Myung, and Mark A Pitt. *Advances in minimum description length: Theory and applications*. MIT press, 2005.
- [31] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1530–1538, Lille, France, 07–09 Jul 2015. PMLR.
- [32] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. 2017.
- [33] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 10215–10224. Curran Associates, Inc., 2018.
- [34] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. In H. Wallach et al, editor, *Advances in Neural Information Processing Systems 32*, pages 7511–7522. Curran Associates, Inc., 2019.
- [35] I. Kobyzev, S. Prince, and M. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020.