

GINN-LP: A Growing Interpretable Neural Network for Discovering Multivariate Laurent Polynomial Equations

Nisal Ranasinghe¹, Damith Senanayake¹, Sachith Seneviratne^{1,2}, Malin Premaratne³, Saman Halgamuge¹

¹AI, Optimization and Pattern Recognition Research Group, Dept. of Mechanical Eng., University of Melbourne, Australia

²Melbourne School of Design, University of Melbourne, Australia

³Department of Electrical and Computer Systems Engineering, Monash University, Australia
nsranasinghe@student.unimelb.edu.au, {damith.senanayake, sachith.seneviratne}@unimelb.edu.au,
malin.premaratne@monash.edu, saman@unimelb.edu.au

Abstract

Traditional machine learning is generally treated as a black-box optimization problem and does not typically produce interpretable functions that connect inputs and outputs. However, the ability to discover such interpretable functions is desirable. In this work, we propose GINN-LP, an interpretable neural network to discover the form and coefficients of the underlying equation of a dataset, when the equation is assumed to take the form of a multivariate Laurent Polynomial. This is facilitated by a new type of interpretable neural network block, named the “power-term approximator block”, consisting of logarithmic and exponential activation functions. GINN-LP is end-to-end differentiable, making it possible to use backpropagation for training. We propose a neural network growth strategy that will enable finding the suitable number of terms in the Laurent polynomial that represents the data, along with sparsity regularization to promote the discovery of concise equations. To the best of our knowledge, this is the first model that can discover arbitrary multivariate Laurent polynomial terms without any prior information on the order. Our approach is first evaluated on a subset of data used in SRBench, a benchmark for symbolic regression. We first show that GINN-LP outperforms the state-of-the-art symbolic regression methods on datasets generated using 48 real-world equations in the form of multivariate Laurent polynomials. Next, we propose an ensemble method that combines our method with a high-performing symbolic regression method, enabling us to discover non-Laurent polynomial equations. We achieve state-of-the-art results in equation discovery, showing an absolute improvement of 7.1% over the best contender, by applying this ensemble method to 113 datasets within SRBench with known ground-truth equations.

Introduction

Although supervised machine learning (ML) can create arbitrary mappings between a large number of input variables and target variables, typically, there is no ability to discover an interpretable mathematical expression between the two variable spaces. However, in cases where the recovery of an underlying governing equation is required, this lack of interpretability is an obstacle. Interpretable machine learning has the potential to harness data to discover underlying equations that govern them.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Phenomenon	Equation
Coulomb’s Law	$E_f = \frac{1}{4\pi\epsilon} \frac{q_1 q_2}{r^2}$
Kinetic Energy	$K = \frac{m}{2} (u^2 + v^2 + w^2)$
Potential Energy (Gravity)	$U = Gm_1 m_2 \left(\frac{1}{r_2} - \frac{1}{r_1} \right)$

Table 1: Three example equations as multivariate LPs and their relevant phenomena

In this work, we focus on the data-driven discovery of equations that take the form of multivariate Laurent polynomials (LP). LPs produce equations important in physics and real-world systems, with some examples shown in Table 1.

More formally, we aim to discover a multivariate LP that can accurately map features $\mathbf{X} \in \mathbb{R}^d$ to a target $y \in \mathbb{R}$ using a dataset of paired samples for \mathbf{X} and y . An exhaustive search is tractable to solve this problem when the search space is small but quickly becomes intractable as the number of variables and possible terms increases, as demonstrated in the appendix¹.

In this work, we present GINN-LP, a Growing Interpretable Neural Network for discovering multivariate Laurent Polynomials, which can efficiently discover the form and coefficients of a multivariate LP equation that describes a given dataset, without explicitly searching through the equation space.

The main contributions of this work are as follows.

- A new type of interpretable neural network (NN) block named the “power-term approximator” (PTA).
- GINN-LP: An interpretable neural network architecture that uses multiple PTA blocks to discover multivariate LPs using observed data.
- A neural network growth strategy allowing the automatic discovery of the number of terms in the underlying polynomial while reducing overfitting and training time.
- A model training strategy that can effectively train the interpretable neural network and decide the best trade-off between accuracy and model simplicity.

¹Appendix available in <https://arxiv.org/abs/2312.10913>

Preliminaries

Here, we describe the preliminaries in formulating our rationale.

Definition: A multivariate LP has the form

$$P = \sum_{i=1}^n c_i \prod_{j=1}^k V_j^{p_i(j)} \quad (1)$$

where $c_i \in \mathbb{R}$ is the i -th term of the polynomial and $p_i(j) \in \mathbb{Z}$ is the power of the j -th variable in that term. The set of variables $V_j \in \{V_1, \dots, V_k\}$ is the finite set of variables considered in the LP.

To the best of our knowledge, GINN-LP is the first method which can discover governing equations with arbitrary multivariate LP terms. The NN is end-to-end differentiable and therefore, can be trained using backpropagation. A neural network growth strategy is implemented to automatically discover the correct number of terms in the underlying equation. The model is trained on a dataset consisting of (\mathbf{x}_i, y_i) pairs (input-output) where $0 \leq i \leq N$ and N is the number of data points in the dataset.

In the following section, we review the existing literature on related methods. Then, we introduce GINN-LP, our interpretable NN that can discover multivariate LPs describing datasets. The results of the experiments conducted are then presented, and the final section provides a summary of this work, its limitations and some future research directions.

Related Work

Data-driven discovery of equations has been mostly dominated by genetic programming (GP) methods in the early literature, exemplified by the fact that symbolic regression (SR) was first introduced as a GP problem in (John R. Koza 1992). However, with the increased interest in explainable AI, a variety of methods have been developed in the recent past (La Cava et al. 2021).

Methods for symbolic regression: Symbolic regression is an area of research focused on discovering equations that describe given datasets, by efficiently searching through the space of possible equations.

A well-known method named AIFeynman discovers simplifying properties (e.g. symmetry, separability) of real-world equations using a NN to greatly reduce the search space till the search space becomes tractable (Udrescu and Tegmark 2020; Udrescu et al. 2020). Another class of SR methods use sparse regression such as least absolute shrinkage and selection operator (LASSO) (Tibshirani 1996) to recover a sparse coefficient matrix of a pre-defined equation structure, making the overall search more efficient (Brunton, Proctor, and Kutz 2015; Rudy et al. 2017; Brunton, Proctor, and Kutz 2016; Kaheman, Kutz, and Brunton 2020; Quade et al. 2018; Messenger and Bortz 2021). Sparse regression-based methods for equation recovery have been further improved using Bayesian methods (Zhang and Lin 2018; Jin et al. 2019), allowing the quantification of uncertainty.

Genetic programming (GP): Most early SR methods use genetic programming (Koza 1994; Dubčáková 2011; Schmidt and Lipson 2009). In these algorithms, a population

of equations are evolved along multiple iterations till an acceptably accurate equation is found. GP based methods have been successfully used in areas such as microbial modelling (Vidanaarachchi et al. 2020). GP algorithms suffer from the combinatorial nature of the problems, i.e., the search space is vast and a large number of equally suitable equations may be recovered. Additionally, GP algorithms tend to be slower in convergence compared to other optimization methods such as gradient-based methods and therefore do not scale well.

Shallow Neural Networks: A class of equation learning neural networks, called equation learners (EQL) is presented in (Martius and Lampert 2016; Sahoo, Lampert, and Martius 2018). In EQL, the conventional activation functions are replaced by operators such as summation and trigonometric functions which are commonly found in real-world equations. This EQL network is integrated with deep learning architectures for end-to-end training in (Kim et al. 2021). PDE-NET 2.0 (Long, Lu, and Dong 2018) uses a symbolic neural network inspired by (Martius and Lampert 2016) to approximate the non-linear response function of the underlying partial differential equation (PDE) governing a system.

Deep Neural Networks: Some recent works in SR circumvent the uninterpretable nature of deep neural networks (DNN) by using a large model to search the smaller space of symbolic models. Deep symbolic regression (Petersen et al. 2021) proposes the use of a recurrent neural network (RNN) to output a distribution over symbolic expressions. The network is trained using a reinforcement learning method to learn the structure of the expression and the constants are learned using a non-linear optimization method such as the Broyden–Fletcher–Goldfarb–Shanno algorithm (BFGS) (Fletcher 2000).

In some recent deep learning methods, the model is pre-trained on a large collection of symbolic regression datasets with known equations (Valipour et al. 2021; Biggio et al. 2021; Kamienny et al. 2022). This contrasts with the majority of other methods, where the algorithms have to be trained from scratch for each SR dataset. Though they have been shown to be faster in equation discovery, the pre-training time could be quite large.

Performance estimation and benchmarking: The field of symbolic regression lacks a widely agreed-upon benchmarking platform or dataset. A recent work has introduced SRBench, an open-source, reproducible platform for benchmarking symbolic regression methods (La Cava et al. 2021), aiming to establish a standardized framework for evaluating SR methods. SRBench has incorporated over 130 datasets with ground truth equation forms to the Penn Machine Learning Benchmark (PMLB) (Olson et al. 2017), which can be used to evaluate SR methods. The main source of symbolic regression datasets used in SRBench is the Feynman dataset, which originates from the Feynman lecture series on physics (Feynman, Leighton, and Sands 2011). SRBench compares 14 different symbolic regression algorithms using these SR datasets as well as black-box datasets.

Method

In this work, we present GINN-LP, a novel growing interpretable neural network that can discover multivariate

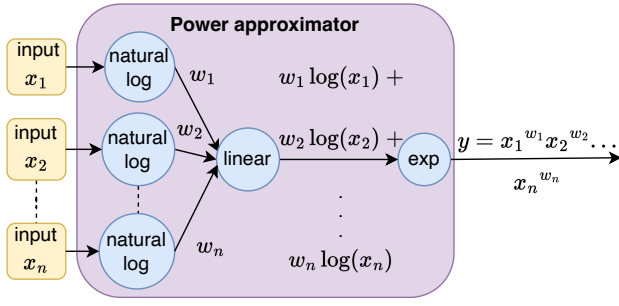


Figure 1: The architecture of the proposed interpretable NN block named the “PTA” block. This block can discover a single term in a multivariate LP. x_1, x_2, \dots, x_n are the inputs to the block and w_1, w_2, \dots, w_n are weights of the linear activated neuron.

LP equations which describe a dataset². Once trained on a dataset of input-output pairs, the network can produce a LP that fits the data, while accurately predicting whether the LP assumption is valid for the given dataset.

The main component of the proposed architecture is the “power-term approximator” (PTA) block, which is an interpretable NN block consisting of a logarithmic unit and a single linear activated neuron followed by an exponential activation. We illustrate the architecture of a PTA in Figure 1. If x_1, x_2, \dots, x_n denote the inputs to the block, the output y_d of the linear activated neuron and the output of the network y are given by,

$$y_d = w_1 \log(x_1) + w_2 \log(x_2) + \dots + w_n \log(x_n) \quad (2)$$

$$= \log(x_1^{w_1} x_2^{w_2} \dots x_n^{w_n}) \quad (3)$$

$$y = e^{y_d} = x_1^{w_1} x_2^{w_2} \dots x_n^{w_n} \quad (4)$$

where w_1, w_2, \dots, w_n are the weights of the linear activated neuron.

Since the PTA can be generalized to support any arbitrary number of inputs, a single block can in theory approximate an arbitrary multivariate LP with a single term.

To enable the network to discover equations with multiple terms, a set of PTA blocks are stacked in parallel, followed by a single neuron with a linear activation as shown in Figure 2. The output is a linear combination of the outputs of all PTA blocks. Each PTA is expected to discover a single term in a polynomial equation, including linear terms. More generally, they can recover equations with multiple additive terms each having an arbitrary number of variables raised to arbitrary powers. To promote the learning of exact equations, we perform a rounding of all coefficients and exponents to the nearest ϵ after training. ϵ was empirically set to 0.001 in our experiments.

Training Strategies

In our implementation, GINN-LP is trained to minimize the mean squared error (MSE) using the Adam optimizer and an

²Source code available in <https://github.com/nisalt/ginn-lp>

Algorithm 1: GINN-LP training strategy

Input: Dataset of input \mathbf{x} , output \mathbf{y} pairs

Parameters: Training instance count ($N = 4$), model complexity weight ($\alpha = 10^{-6}$), L_1 regularization factor ($\lambda_1 = 10^{-4}$), L_2 regularization factor ($\lambda_2 = 10^{-4}$), maximum number of PTA blocks ($P = 4$), number of epochs per network growth stage (E), rounding precision (ϵ)

Output: An equation that describes the relationship between inputs and outputs

```

1: for  $k_{inst} \leftarrow 1$  to  $N$  do
2:   model  $\leftarrow$  InitNetwork(num_PTA=0)
3:   MSE  $\leftarrow$   $\infty$ 
4:   max_SEM  $\leftarrow$   $\infty$ 
5:   for  $k_{PTA} \leftarrow 1$  to  $P$  do
6:     GrowModel(model)
7:     TrainModel(model,  $L_1=\lambda_1$ ,  $L_2=\lambda_2$ , epochs= $E/2$ )
8:     TrainModel(model,  $L_1=0$ ,  $L_2=0$ , epochs= $E/2$ )
9:     RoundExponents(model, round_prec= $\epsilon$ )
10:    new_MSE = CalcMSE(model, validation_data)
11:    if new_MSE > MSE * 0.8 then
12:      break
13:    end if
14:    MSE = new_MSE
15:  end for
16:  eq  $\leftarrow$  GetEQ(model)
17:   $C_M \leftarrow$  CalcComplexity(eq)
18:   $SE_M \leftarrow$  MSE +  $\alpha * C_M$ 
19:  if  $SE_M > max\_SEM$  then
20:    max_SEM  $\leftarrow$   $SE_M$ 
21:    best_eq  $\leftarrow$  eq
22:  end if
23: end for
24: return best_eq

```

exponentially decaying learning rate starting from 0.01. Algorithm 1 outlines the training strategies used to train GINN-LP efficiently, along with the empirically determined hyperparameters. We further discuss these strategies below.

Regularization to enforce equation conciseness. In practice, simpler equations are preferred over more complex ones to describe relationships between variables since they are easier to understand and interpret. To enforce this property, we impose sparsity-promoting regularization on our interpretable neural network. Since the values of the weight matrix correspond to the coefficients and power terms of the discovered equation, a sparser weight matrix would result in a simpler equation.

However, using regularization throughout training from beginning to end causes some discovered equations to be slightly different to the ground-truth equation. To avoid this, we introduce an unregularized training phase after regularized training. This allows the equation coefficients to converge towards more accurate values.

During the regularized training phase, a linear combination of L_1 and L_2 regularization is applied. Hence, the objective function would be,

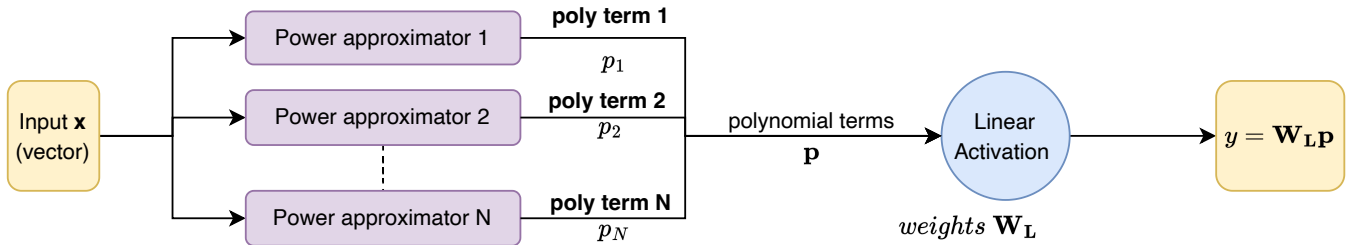


Figure 2: The architecture of the proposed interpretable neural network, GINN-LP. This consists of multiple PTA blocks in parallel, each discovering a single term in the underlying multivariate LP.

$$L = \frac{1}{N} \sum_{i=1}^N (f(\mathbf{x}) - y_i)^2 + \lambda_1 \sum_{i=1}^n |W_i| + \lambda_2 \sum_{i=1}^n W_i^2 \quad (5)$$

Here, N is the number of samples in the dataset, n is the number of weights in the weight matrix and λ_1, λ_2 are the regularization constants. When the unregularized phase was not included, we observed that the equation coefficients would sometimes deviate from their true values by small amounts (e.g. 0.001). The point at which the regularization changes is expressed as a fraction of the total number of training epochs. This was empirically set to 0.5.

Neural network growth. Having a higher number of PTA blocks than required could lead to the model overfitting. Therefore, starting from a single PTA block, the number of PTA blocks is increased, thereby growing the GINN-LP architecture. When the network is grown, a new, randomly initialized PTA block is added in parallel without altering the weights of the already trained PTA blocks. Then, the grown network is trained for a fixed number of epochs on the dataset. This is performed iteratively till the network is grown to a pre-defined maximum size, or till an early stopping condition is reached.

Training is stopped if the validation MSE does not decrease by a certain percentage after each growth iteration. This percentage is a hyperparameter of the model and is empirically set to 20% in our experiments. We expect the training to stop when the number of PTA blocks is equal to the number of terms in the equation.

Training multiple instances. When training GINN-LP, the network sometimes gets stuck in local optimums. To ensure a well-fitting model, we train the network multiple times with different random initializations, and the best model is selected as outlined in the following sub-section.

Model selection. When there are multiple models (equations) with similar performances on unseen data, we prefer the simpler model. To select the best model from all training instances, we propose a new performance metric named the symbolic error (SE) using a linear combination of the model complexity, and the validation set MSE. We use the same complexity metric as SRBench (La Cava et al. 2021), where the complexity (C_M) of a given model M is defined simply as the number of mathematical operators (op_M), constants ($const_M$) and features (n_m) in the equation discovered by

GINN-LP Output	LP/Non-LP?	Secondary SR output	Ensemble output
$x_1^{2.342}x_2^2$	Non-LP	$\sin(x_1)/x_2$	$\sin(x_1)/x_2$
$x_1^2x_2^{-1}$	LP	N/A	$x_1^2x_2^{-1}$

Table 2: Hypothetical examples of the ensemble method output, showing the intermediate outputs.

the model. The best model \hat{M} is determined as shown below,

$$C_M = op_M + const_M + n_M \quad (6)$$

$$SE_M = MSE(y, \hat{y}|\theta_M) + \alpha * C_M \quad (7)$$

$$\hat{M} = \underset{M}{\operatorname{argmin}}[SE_M] \quad (8)$$

Here, θ_M are the parameters of model M , y is the target, \hat{y} is the predicted value and α is the weighting factor of the model complexity. $\alpha = 10^{-6}$ is determined empirically.

Ensemble Method for Discovering Non-Laurent Polynomial Equations

To demonstrate the broader applicability of GINN-LP, we propose an ensemble equation discovery pipeline that combines GINN-LP with high-performing SR methods, to enable the discovery of both LP and non-LP equations, thereby achieving state-of-the-art performance in equation discovery. In this ensemble method, we first fit GINN-LP, and if at least one of the exponents of the discovered equation is not an integer, we determine that the assumption of the ground-truth equation being an LP is invalid. In this case, we fit a secondary SR method to discover an accurate equation. The ensemble pipeline is illustrated in Figure 3 and two hypothetical examples showing the GINN-LP and secondary SR method outputs are given in Table 2.

We perform multiple experiments with different options for the secondary SR method, and show that ensembling with GINN-LP improves the performance of all other SR methods used for evaluation.

Results

We evaluate our approach on the Feynman symbolic regression benchmark dataset, using SRBench (La Cava et al.

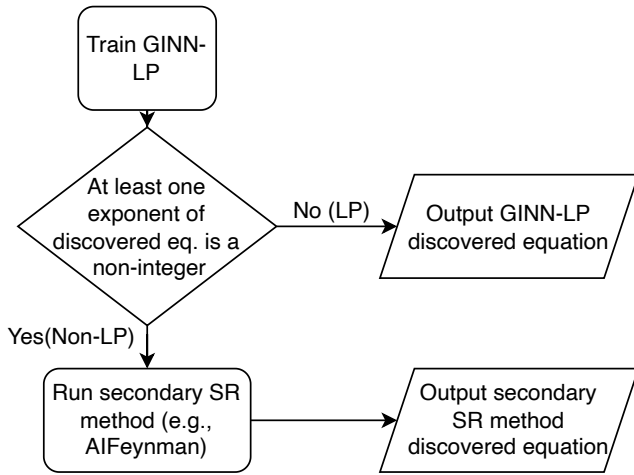


Figure 3: The ensemble pipeline combines GINN-LP with another high-performing SR method, enabling it to discover both LP and non-LP equations

2021). The Feynman dataset consists of multiple SR datasets generated from real-world equations, along with their corresponding ground-truth equations. We note that our method may not qualify as a symbolic regression method, since we are not explicitly searching through a space of equations. However, due to the similarities in the aim, we evaluate our method against SR methods.

GINN-LP was compared against 14 other popular SR methods in terms of how well it recovers the ground-truth equation. Each SR method was run for five trials per equation dataset to investigate whether they can consistently recover the correct ground truth equation with different random initializations.

All experiments were conducted on a computing cluster where each experiment used a 32-core, 2.90GHz Intel Xeon Gold 6326 CPU and a single NVIDIA A-100 GPU. We note that our approach does not see a major performance improvement when training on a GPU, since the network architecture is small. Moreover, we did not re-run experiments for other methods where results were already available within SRBench³.

Performance Metric

We use the same definition of exact recovery as SRBench (La Cava et al. 2021). Each recovered equation is simplified using SymPy (Meurer et al. 2017) and compared with the corresponding ground-truth equation to determine whether it has been recovered exactly. For each SR method tested, the performance is reported using the symbolic solution recovered percentage, also called the solution rate. This is calculated as,

$$solution\ rate = \frac{correctly\ recovered\ equations}{total\ number\ of\ equations} \times 100\% \tag{9}$$

³SRBench code available in <https://github.com/cavalab/srbench>

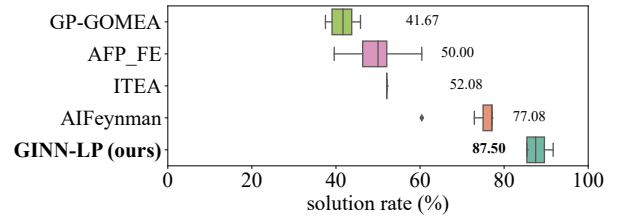


Figure 4: Solution rate of the top five algorithms, for all datasets with LP ground-truths. The median solution rates are shown on the side of each plot.

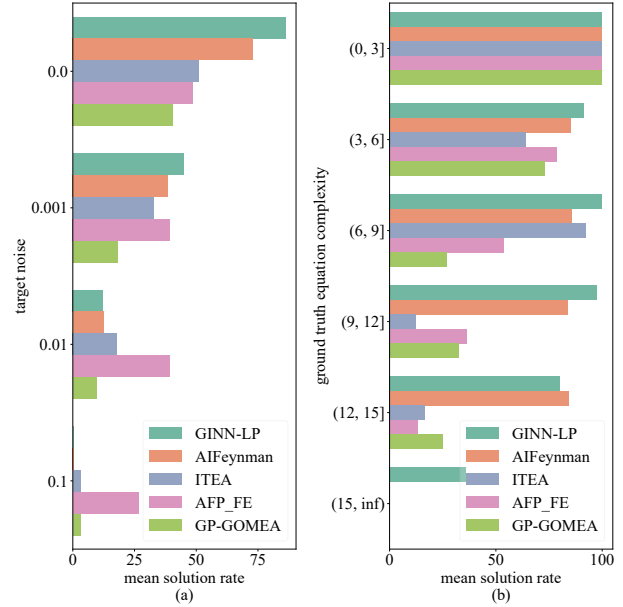


Figure 5: (a) Comparison of SR methods, with the presence of target noise. The mean solution rates are reported. (b) Comparison of SR methods with respect to the ground truth equation complexity. The mean solution rate is reported

Since multiple trials (with different random initializations) are performed for each dataset and algorithm, we report a range of the solution rate or the mean/median value. In each trial, the network was trained on 10,000 data points, generated from the ground-truth equation of each of the datasets.

Standalone GINN-LP Performance Results – Recovering LP Equations

We first selected the subset of 48 multivariate LPs within the Feynman dataset to evaluate our method, as GINN-LP is designed to retrieve such equations. The ground-truth equations of these datasets are listed in the appendix.

Comparison of models without added noise. We compare GINN-LP with 14 popular symbolic regression algorithms using SRBench. The results of the experiments without adding any noise to the datasets are shown in Figure 4

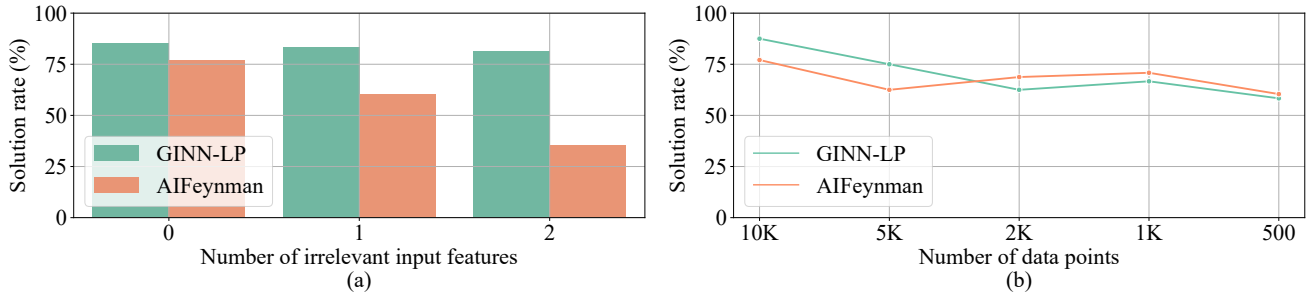


Figure 6: (a) Model performance (solution rate) with respect to the number of irrelevant inputs features (b) Model performance (solution rate) with a varying number of data points used for training

for the five top-performing algorithms. Our method outperforms all the 14 algorithms it was compared against in terms of minimum, maximum and median solution rates. To ensure fairer comparison, experiments were performed by reducing the search space of AIFeynman only to include operators present in LPs, to incorporate the same prior that we are incorporating into our method. However, the performance of AIFeynman did not increase noticeably with this reduced search space.

Impact of adding Gaussian noise. To evaluate the noise-resilience of GINN-LP, experiments were conducted after adding Gaussian white noise to the target as a fraction of the target root mean square value. Experiments were repeated with noise fractions of 0.001, 0.01 and 0.1, on the top five algorithms identified in the previous section (without adding noise) and the results are shown in Figure 5(a). Our method outperforms the other top methods with low noise (0.001) but does not manage to beat the AFP_FE method with high noise values.

Equations that were recovered by our method in the high noise cases, seemed to have recovered the structure with some slight changes to the constants. We hypothesize that the network overfits due to attempting to model the noise.

However, we note GINN-LP is still able to model noisy data quite accurately, even if it’s unable to recover the exact groundtruth equations, as shown in the appendix.

Solution rate vs ground truth equation complexity. Further experiments were conducted to observe how the GINN-LP performance changes with the complexity of the equation we attempt to recover. Here, complexity C_M is calculated using the metric defined in Equation 6. The equations in our dataset consisted of complexity values ranging from 3 to 27. To visualize our results, we grouped the complexity values into six bins. The intervals of these bins are $(0, 3]$, $(3, 6]$, $(6, 9]$, $(9, 12]$, $(12, 15]$ and $(15, +\infty)$.

We compare the performance of the five top-performing algorithms on each of the six bins. The mean solution rate (across all experiments) for each bin for a particular algorithm is visualized in Figure 5(b). Our method maintains good solution rates across most of the complexity bins, with a moderate solution rate of close to 40% even for the most complex bin, where all other methods fail.

Training with irrelevant inputs. In practice, we might not know which of our input variables are related to the output. We investigate how GINN-LP performs when there are such variables, by adding irrelevant inputs sampled from a uniform distribution within the same range as the other inputs. Experiments were conducted with up to 2 irrelevant inputs and the results are shown in Figure 6(a). We compare our method with AIFeynman, which was the top-performing method (among other methods) in earlier experiments. Our method does not see a significant drop in performance even with 2 irrelevant inputs, while the performance of AIFeynman drops considerably.

Effect of training data size. We compared our method with AIFeynman when trained with different sizes of training datasets and present the results in Figure 6(b). We note that GINN-LP can still recover a majority of the equations even with a small number of samples, and remains competitive with AIFeynman.

Effect of epoch count and number of training instances. Further experiments were done to investigate the effect of these hyperparameters on the performance of GINN-LP. Results of these experiments are provided in the appendix.

Ensemble Method Performance Results

We evaluate the performance of the proposed ensemble method using the entire Feynman symbolic regression benchmark dataset within SRBench. This consists of 113 datasets, including all 48 datasets with LP ground-truths. GINN-LP perfectly identifies when the LP assumption does not hold for a given dataset, as demonstrated by the confusion matrix included in the appendix. This indicates that GINN-LP can successfully identify the problems it cannot solve, enabling them to be solved using a secondary, high-performing SR method. We show the results after ensembling our method with AIFeynman and Multiple Regression Genetic Programming (MRGP) in the below sub-sections, since these two methods show the best performances in terms of their solution rate and predictive accuracy respectively. We further show in the appendix, that ensembling with GINN-LP enhances the performance of all 14 other SR methods.

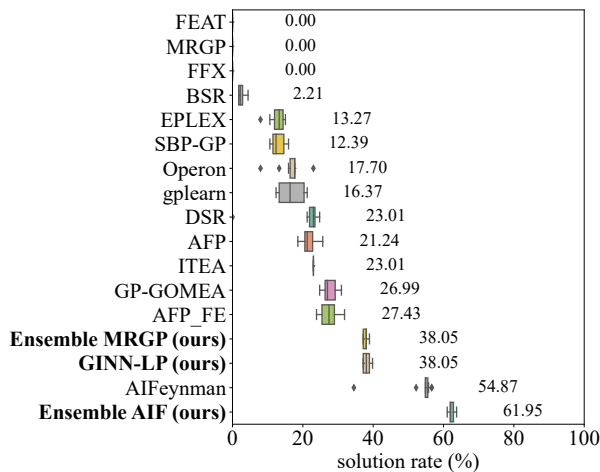


Figure 7: Solution rate of ensemble models (with MRGP, AIFeynman) and standalone GINN-LP, compared against other methods, for all 113 datasets in the Feynman symbolic regression benchmark dataset (including non-LP). The median solution rate is shown on the side of each plot.

Solution rate. The results of the proposed ensemble methods in terms of their solution rates are shown in Figure 7. We show that our method ensemble with AIFeynman outperforms state-of-the-art methods by 7.1% in terms of the median solution rate (absolute). Moreover, despite being formulated for recovering LP equations, the standalone GINN-LP still outperforms the majority of other methods on the complete Feynman dataset.

Predictive accuracy. Further experiments were conducted to evaluate how accurately our methods can predict outputs for new, previously unseen data. The prediction accuracy is calculated as the percentage of datasets with R^2 above 0.99 on the test data. A train-test split of 75-25 is used. We compare the prediction accuracy of all methods on all 113 datasets, including GINN-LP, and the proposed ensemble methods. As shown in Figure 8, all of our methods, including the standalone GINN-LP perform reasonably well while GINN-LP ensemble with MRGP outperforms all other methods by 0.9%, while achieving near perfect $R_2 > 0.99$ accuracy.

Conclusion

In this work, we present GINN-LP, an end-to-end differentiable interpretable neural network that can recover mathematical expressions that take the form of multivariate LPs. This is made possible by a new type of interpretable neural network block we introduce, named the PTA block, that can discover terms of multivariate LPs. Our method can automatically discover the number of terms in the underlying equation using a neural network growth strategy while promoting the discovery of simpler equations by applying sparsity regularization.

The results of the experiments conducted show that our

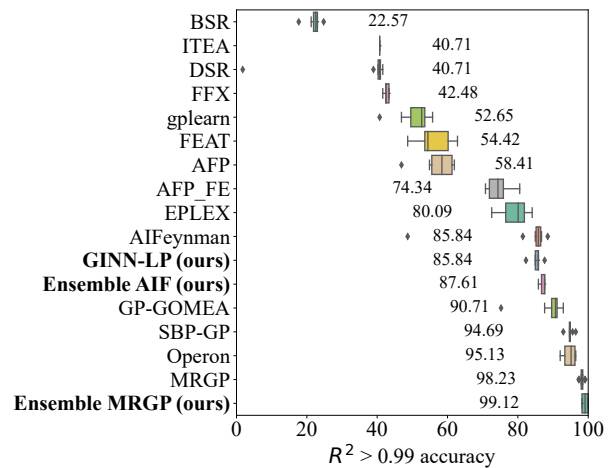


Figure 8: $R^2 > 0.99$ accuracy of ensemble models (with MRGP, AIFeynman) and standalone GINN-LP, compared against other methods, for all 113 datasets (including non-LP) in the Feynman symbolic regression benchmark dataset. The median accuracy is shown on the side of each plot.

approach achieves the state-of-the-art results, outperforming the best contender in symbolic regression algorithms on datasets with multivariate LP ground-truth equations in the Feynman symbolic regression benchmark dataset. We also show that GINN-LP can successfully identify when a given dataset cannot be accurately described using a multivariate LP. Building on these observations, we propose an ensemble method which uses a secondary SR method to discover equations when GINN-LP identifies that the LP assumption does not hold for a given dataset. Further experiments show that this ensemble method achieves state-of-the-art results in equation discovery, when evaluated on the entire Feynman dataset, including both LP and non-LP equations.

A limitation of the proposed architecture is that it assumes that all input values inside and outside of the training dataset are positive real numbers. This is because of the natural logarithmic activation used within the PTA block. This could be overcome by adding a constant to all input features with negative values. However, this will make it more difficult to discover the mathematical equation since adding a constant to the input could change an LP equation to a more complex equation (e.g. x_1/x_2 changes to $(x_1 + 2)/(x_2 + 2)$). Future work may investigate NN architectures that can perform power approximation in the presence of negative input values.

Another interesting future research direction is to extend this architecture to discover a wider variety of equations, without being limited to only LPs. One way to achieve this is to grow the network sequentially, thereby stacking PTA blocks in multiple layers one after the other. However, this would significantly increase the risk of overfitting. Therefore, greater attention will need to be paid to regularizing the weights, in addition to the growth strategy.

Acknowledgements

This study was supported by the Melbourne Graduate Research Scholarship for the first author, the Australian Research Council grant DP210101135 for the first author and the Australian Research Council grant DP220101035 for the second, third and the last authors. This study was supported partially for the third author by the Australian Government through the Australian Research Council's Centre of Excellence for Children and Families over the Life Course (Project ID CE200100025). This research was also supported by The University of Melbourne's Research Computing Services and the Petascale Campus Initiative. The authors thank Nadarasar Bahavan, Deshani Poddenige and Rashindrie Perera for proofreading.

References

- Biggio, L.; Bendinelli, T.; Neitz, A.; Lucchi, A.; and Parascandolo, G. 2021. Neural Symbolic Regression that Scales. In *International Conference on Machine Learning*.
- Brunton, S. L.; Proctor, J. L.; and Kutz, J. N. 2015. Discovering governing equations from data: Sparse identification of nonlinear dynamical systems.
- Brunton, S. L.; Proctor, J. L.; and Kutz, J. N. 2016. Sparse Identification of Nonlinear Dynamics with Control (SINDYc). In *IFAC-PapersOnLine*, volume 49, 710–715. Elsevier B.V.
- Dubčáková, R. 2011. Eureqa: software review. *Genetic Programming and Evolvable Machines*, 12(2): 173–178.
- Feynman, R. P.; Leighton, R. B.; and Sands, M. 2011. *The Feynman lectures on physics, Vol. I: The new millennium edition: mainly mechanics, radiation, and heat*. Basic Books.
- Fletcher, R. 2000. *Practical Methods of Optimization*. Chichester, West Sussex England: John Wiley & Sons, Ltd. ISBN 9781118723203.
- Jin, Y.; Fu, W.; Kang, J.; Guo, J.; and Guo, J. 2019. Bayesian Symbolic Regression.
- John R. Koza. 1992. *Genetic programming - on the programming of computers by means of natural selection*. MIT Press. ISBN 0-262-11170-5.
- Kaheman, K.; Kutz, J. N.; and Brunton, S. L. 2020. SINDY-PI: A robust algorithm for parallel implicit sparse identification of nonlinear dynamics: SINDY-PI. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2242).
- Kamienny, P.-A.; d'Ascoli, S.; Lample, G.; and Charton, F. 2022. End-to-end symbolic regression with transformers. In *Advances in Neural Information Processing Systems*.
- Kim, S.; Lu, P. Y.; Mukherjee, S.; Gilbert, M.; Jing, L.; Ceperic, V.; and Soljacic, M. 2021. Integration of Neural Network-Based Symbolic Regression in Deep Learning for Scientific Discovery. *IEEE Transactions on Neural Networks and Learning Systems*, 32(9): 4166–4177.
- Koza, J. 1994. Genetic programming as a means for programming computers by natural selection. *Statistics and Computing*, 4(2).
- La Cava, W.; Orzechowski, P.; Burlacu, B.; de França, F. O.; Virgolin, M.; Jin, Y.; Kommenda, M.; and Moore, J. H. 2021. Contemporary Symbolic Regression Methods and their Relative Performance. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Long, Z.; Lu, Y.; and Dong, B. 2018. PDE-Net 2.0: Learning PDEs from Data with A Numeric-Symbolic Hybrid Deep Network.
- Martius, G.; and Lampert, C. H. 2016. Extrapolation and learning equations.
- Messenger, D. A.; and Bortz, D. M. 2021. Learning Mean-Field Equations from Particle Data Using WSINDy.
- 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.; Roucka, S.; Saboo, A.; Fernando, I.; Kulal, S.; Cimrman, R.; and Scopatz, A. 2017. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3: e103.
- Olson, R. S.; La Cava, W.; Orzechowski, P.; Urbanowicz, R. J.; and Moore, J. H. 2017. PMLB: a large benchmark suite for machine learning evaluation and comparison. *Bio-Data Mining*, 10(1): 36.
- Petersen, B. K.; Larma, M. L.; Mundhenk, T. N.; Santiago, C. P.; Kim, S. K.; and Kim, J. T. 2021. Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients. In *International Conference on Learning Representations*.
- Quade, M.; Abel, M.; Nathan Kutz, J.; and Brunton, S. L. 2018. Sparse identification of nonlinear dynamics for rapid model recovery. *Chaos*, 28(6).
- Rudy, S. H.; Brunton, S. L.; Proctor, J. L.; and Kutz, J. N. 2017. Data-driven discovery of partial differential equations. *Science Advances*, 3(4).
- Sahoo, S. S.; Lampert, C. H.; and Martius, G. 2018. Learning Equations for Extrapolation and Control.
- Schmidt, M.; and Lipson, H. 2009. Distilling Free-Form Natural Laws from Experimental Data. Technical Report 5923.
- Tibshirani, R. 1996. Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1): 267–288.
- 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.
- Udrescu, S.-M.; and Tegmark, M. 2020. AI Feynman: A physics-inspired method for symbolic regression. Technical report.
- Valipour, M.; You, B.; Panju, M.; and Ghodsi, A. 2021. SymbolicGPT: A Generative Transformer Model for Symbolic Regression.
- Vidanaarachchi, R.; Shaw, M.; Tang, S.-L.; and Halgamuge, S. 2020. IMPARO: inferring microbial interactions through

parameter optimisation. *BMC Molecular and Cell Biology*, 21(S1): 34.

Zhang, S.; and Lin, G. 2018. Robust data-driven discovery of governing physical laws with error bars. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2217).