

Noise-Resilient Symbolic Regression with Dynamic Gating Reinforcement Learning

Chenglu Sun¹, Shuo Shen¹, Wenzhi Tao¹, Deyi Xue¹, Zixia Zhou^{2*}

¹Cooperation Product Department, Interactive Entertainment Group, Tencent

²Stanford University

{clivesun, svenshen, simontao, davexue}@tencent.com, zixia@stanford.edu

Abstract

Symbolic regression (SR) has emerged as a pivotal technique for uncovering the intrinsic information within data and enhancing the interpretability of AI models. However, current state-of-the-art (sota) SR methods struggle to perform correct recovery of symbolic expressions from high-noise data. To address this issue, we introduce a novel noise-resilient SR (NRSR) method capable of recovering expressions from high-noise data. Our method leverages a novel reinforcement learning (RL) approach in conjunction with a designed noise-resilient gating module (NGM) to learn symbolic selection policies. The gating module can dynamically filter the meaningless information from high-noise data, thereby demonstrating a high noise-resilient capability for the SR process. And we also design a mixed path entropy (MPE) bonus term in the RL process to increase the exploration capabilities of the policy. Experimental results demonstrate that our method significantly outperforms several popular baselines on benchmarks with high-noise data. Furthermore, our method also can achieve sota performance on benchmarks with clean data, showcasing its robustness and efficacy in SR tasks.

Introduction

With the rise of electronic information technology, we have easy access to abundant data for acquisition, processing, and analysis. Extracting meaningful relationships from data is crucial for AI design, scientific discovery, and identifying core factors, et al. Deep learning (DL) has emerged as a powerful tool for data mining (Shu and Ye 2023; Sorscher et al. 2022), enabling neural networks to tackle a wide array of scientific tasks, such as regression and classification problems (Muthukumar et al. 2021). For instance, given a dataset (X, y) , where each point $X_i \in \mathbb{R}^n$ and $y_i \in \mathbb{R}$, DL methods can train a network to approximate $y_i \approx F(X_i)$, with F representing the learned network. However, as a black-box system (Buhmester, Münch, and Arens 2021), the relationships established by the network between data and targets are often opaque, making fine-grained control of the system and understand the information contained in the data challenging. Therefore, improving model interpretability and controllability has gradually become a key research direction in the current AI field.

Symbolic regression (SR) stands out as a promising technology, aiming to discover the correct symbolic expression for an unknown function f that best fits the dataset with $y_i = f(X_i)$. SR has the potential to explain black-box systems with simple, accurate, intelligible expressions, enhancing the development of data-driven AI systems (Udrescu and Tegmark 2020; Oliveira et al. 2018; Jiang and Xue 2023; Udrescu et al. 2020; Kim et al. 2020; Petersen et al. 2021; Kamienny et al. 2022, 2023). Despite its promise, SR is difficult due to the exponentially large combinatorial space of symbolic expressions. Traditional SR methods are usually based on genetic programming (GP) (Schmidt and Lipson 2009; Searson, Leahy, and Willis 2010; De Melo 2014). While GP-based SR method achieve good performance, they are slow, and struggle to scale to larger problems and are highly sensitive to hyperparameter settings (Mundhenk et al. 2021b). Recent advances have seen a pivot towards DL approaches for SR, leveraging networks to represent and learn the semantics of symbols (Kamienny et al. 2022; Kim et al. 2020; d’Ascoli et al. 2022). However, DL-based methods often encounter difficulties in achieving satisfactory regression performance or need specific revisions on the network and search space. Some DL-based methods pre-train an encoder-decoder network to learn the expression representations within a given dataset (Valipour et al. 2021; Holt, Qian, and van der Schaar 2023). These methods sample function f using the pre-trained network, thereby achieving low complexity at inference process. But the pre-trained model may lead to sub-optimal solutions and fail to discover highly complex equations (Holt, Qian, and van der Schaar 2023).

Reinforcement learning (RL) (Wang et al. 2022; Nguyen, Nguyen, and Nahavandi 2020; Crochepierre, Boudjeloud-Assala, and Barbesant 2022; Landajuela et al. 2021a) applied to SR combines the benefits of GP with environmental feedback and the representation power of neural networks. This combination has contributed to a growing trend in research adopting RL for SR tasks (Mundhenk et al. 2021a; Petersen et al. 2021; Landajuela et al. 2022; Zhang and Zhou 2021). However, a significant portion of real-world data is characterized by high-noisy¹, often containing abundant noisy information or irrelevant data. Performing SR on

*Corresponding Author: Zixia Zhou
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹In this paper, “noises” refers to irrelevant variables within a dataset.

such high-noise data is exceedingly complex due to the dramatic increase in the search space caused by the noisy input variables (Reinbold et al. 2021). Yet, most benchmarks used in SR research are typically clean, posing a significant challenge for the application of these algorithms in real-world scenarios. To address this issue, we introduce a novel end-to-end Noise-Resilient SR method, termed as NRSR, which can learn expressions from high-noise data via RL with a Noise-Resilient Gating Module (NGM). The NGM, designed to dynamically filter out noisy input variables during the RL training process, significantly boosts the efficiency of exploration. Additionally, we propose a Mixed Path Entropy (MPE) regularizer to bolster the exploration capabilities for searching expressions and prevent overfitting of the RL model. Hence, NRSR not only accurately recovers expressions from high-noise data but also enhances the precision of RL-based SR approaches. To evaluate the effectiveness of NRSR, we employed a suite of benchmarks containing twelve representative expressions and selected five popular SR approaches as baselines for a comprehensive test. The results demonstrate that NRSR significantly outperforms on benchmarks with high-noise data and can surpass all baselines on benchmarks with clean data. The main contributions of this study are threefold: (1) We introduce NRSR, a novel SR method that exhibits state-of-the-art (sota) performance on both high-noise data and clean data. (2) We design a dynamic NGM that effectively filters noisy variables, and the proposed MPE enhances the model’s exploration ability in expression generation. (3) Through experiments, we analyze the performance of the NGM and the MPE regularizer. These components can be decoupled from our method, allowing for integration with other scenarios.

Related Works

Reinforcement learning for symbolic regression Several recent approaches leverage RL-based method for SR, and represents decent performance. Petersen et al. (Petersen et al. 2021) introduced an SR framework utilizing RL, where a RNN generates mathematical expressions optimized through a risk-seeking policy gradient (PG) algorithm, outperforming established baselines and commercial software on twelve benchmarks. Crochepierre et al. (Crochepierre, Boudjeloud-Assala, and Barbesant 2022) proposed an interactive web-based platform that enhances grammar-guided SR by incorporating user preferences through a RL framework. Zhang et al. (Zhang and Zhou 2021) presented an SR method that combines genetic algorithms and RL to solve SR problems, addressing the challenge of balancing exploration and exploitation. The experimental results, based on ten benchmarks, demonstrate that this hybrid approach achieves competitive performance. And other existing RL-based SR methods also have shown promising results (Landajuela et al. 2021a; Mundhenk et al. 2021a; Landajuela et al. 2022). RL possesses robust decision-making capabilities and an exceptional aptitude for target fitting, making it an ideal foundational path for SR. Despite the advancements in RL-based SR methods, their performance on high-noise data has not been satisfactory. This area, therefore, warrants further exploration to im-

prove the robustness and applicability of these techniques.

L0 Regularization L1 and L2 regularization are commonly used regularization methods in neural networks (Ma et al. 2019; Cortes, Mohri, and Rostamizadeh 2012), which can enhance the generalization of the model and prevent overfitting (Ying 2019). L1 and L2 regularization respectively limit the absolute value and squared magnitude of the weights in the network. L0 regularization is to encourage the network to have a small number of non-zero parameters (Louizos, Welling, and Kingma 2018; Wei et al. 2022), which can lead to sparser models. However, L0 regularization is non-differentiable and thus harder to implement in practice. Louizos et al. (Louizos, Welling, and Kingma 2018) introduced an L0-norm regularization technique for neural networks that employs stochastic units to induce sparsity, enabling differentiable pruning during training, which enhances generalization. Wei et al. (Wei et al. 2022) present a projected neural network method, using differential equations, to tackle a range of sparse optimization problems. This method combines a non-smooth convex loss with L0-norm regularization and proved its global existence, uniqueness, and convergence properties. Some studies employ either L1-norm, L2-norm, or a combination of both for feature selection (Ng 2004). Indeed, the properties of L0 regularization make it particularly apt for feature selection. Consequently, our method leverages the principles of L0 regularization to design a gating module. This module is utilized to select input variables during the SR training process, as opposed to network features, which were the focus in the aforementioned studies.

Entropy regularization in reinforcement learning Enhancing the exploration capabilities during the RL training process can significantly boost the performance of learned policy (Ecoffet et al. 2019). Mnih et al. (Mnih et al. 2016) found that adding the policy entropy to the objective function improved exploration by discouraging premature convergence to suboptimal deterministic policies. Haarnoja et al. (Haarnoja et al. 2018) presented a maximum entropy RL algorithm, which added the policy entropy to the expectation of reward to addresses the challenges of high sample complexity and hyperparameter sensitivity common in RL algorithms. The incorporation of policy entropy as a regularization within the objective function has progressively been established as a standard operational procedure for the RL algorithms (Mysore et al. 2022; Vinyals et al. 2019). To address the early commitment phenomenon and from initialization bias in sequence search of SR, Landajuela et al. (Landajuela et al. 2021b) introduced a hierarchical entropy regularizer to enhance the entropy of early actions in sequences. To enhance the exploration of entire sequences rather than individual actions, we introduce the MPE approach, which increases the diversity of generated expressions by amplifying the uncertainty in the generation of complete sequences.

Method

In this section, we first introduce the designed NGM and how it can filter the noisy input variables. Subsequently, we present the expression generation process during the training

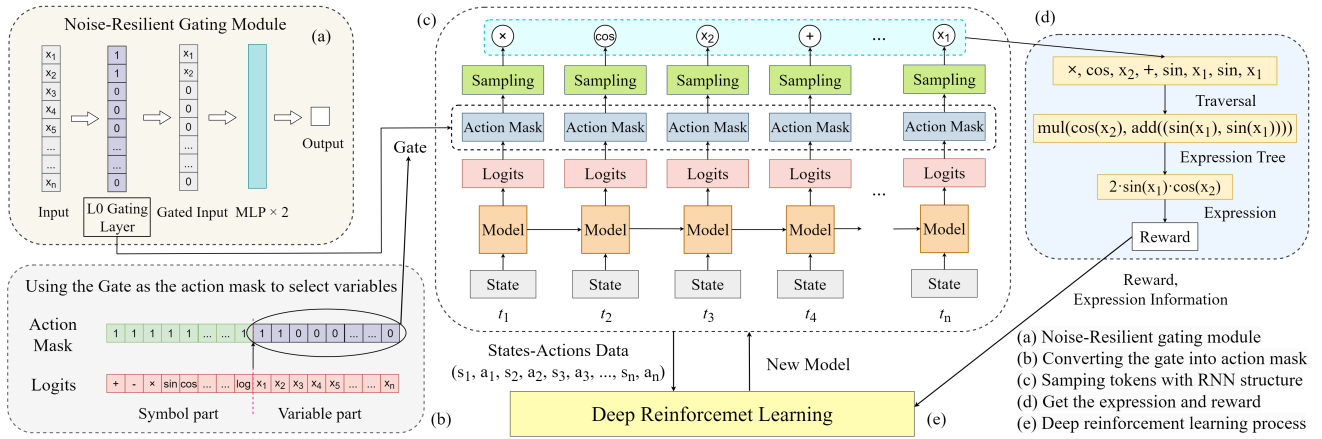


Figure 1: Overview of the NRSR training process. (a) Prior to the SR and RL training process, the NGM is trained with a sample network structure. (b) The obtained L0 gates are then combined with the original action mask to select the input variables. (c) During the SR process, the RNN model, serving as the policy, generates output logits which are processed with the new action mask. These processed logits are subsequently converted into action probabilities, which are used to sample symbolic tokens. The policy generates actions in a step-by-step manner (from t_0 to t_n) following a time sequence structure, thereby forming a trajectory. (d) Each trajectory, consisting of sequential tokens, represents a traversal that can form an expression using the expression tree approach. The resulting expressions are used to calculate the fitness reward. The training process concludes once the optimal expression is found. (e) The samples, comprising states, actions, and rewards, are used to train a new RL policy. After each training iteration, the updated model is used to generate new expressions. The procedures in (c), (d), and (e) constitute an iterative process in NRSR training. This iteration continues until the optimal expression is found or the limit of consumed expressions is reached, thereby ensuring a comprehensive exploration of the solution space.

process. We then describe the expression generation policy with RL approach, and finally, we introduce the MPE regularization with the RL training process.

Noise-Resilient Gating Module

The NGM is designed to induce sparsity in the input variables, effectively performing variable selection by deactivating noisy input as shown in Fig.1 (a). Given an original input $X \in \mathbb{R}^{m \times n}$, a L0 gating layer is utilized to selectively filter the input variables to achieve a noise-resilient function. The gating layer $G \in \{0, 1\}^{m \times n}$ is composed of binary values, where 0 indicates the suppression of the corresponding input and 1 indicates its retention. The processed input X' is computed as the Hadamard product of X and G :

$$X' = X \odot G \quad (1)$$

where \odot denotes the element-wise multiplication. The network output Y' in the gating module is then computed as:

$$Y' = W \cdot X' \quad (2)$$

where W representing the weights of the network.

The gating layer can be trained by achieving the objective, that is to minimize the mean squared error (MSE) between the network output Y' and the true labels y , subject to the L0 regularization constraint on G . The optimization problem can be formulated as:

$$J(W, G) = \min_{W, G} \frac{1}{m} \sum_{i=1}^m (y_i - W X'_i)^2 + \lambda \|G\|_0 \quad (3)$$

where $\|G\|_0$ denotes the L0 norm of G , which counts the number of non-zero parameters in the gating layer, and λ is a regularization parameter that controls the trade-off between the MSE and the sparsity of G . The goal of the L0 norm is to maximize the number of zeros in G , effectively deactivating the corresponding input variables and thus performing variable selection. The optimization of $\|G\|_0$ is computationally intractable due to the non-differentiability and huge possible combinatorial states of G . Hence, we employ an approximation approach (Louizos, Welling, and Kingma 2018) that allows for gradient-based optimization, after that, we can get the trained G as the gate to filter the input variables during the expression generation process.

Integration of Gating Layer with Action Mask

Action mask can prevent the selection of invalid or undesirable actions, thus accelerate the RL training process (Tang et al. 2020; Hou et al. 2023). In our scenario, NRSR trains the noise-resilient gating layer as an action mask (Tang et al. 2020) to filter out noisy input variables, thereby significantly enhancing exploration efficiency. Action mask is realized by modifying the output logits of policy network, where make the probabilities of masked actions to zero. Hence, it doesn't impact the network's parameters or the backpropagation process, thus preserving the model's learning capability. Upon the convergence of the training phase of NGM, the learned gating layer G is integrated with the original action mask used during the sampling stage, as depicted in Fig. 1 (b). The original action mask was designed to constrain the illogical sample of symbols within an expression to expedite

the exploration process. For instance, it prevents the selection of operations that are inverses or descendants of their predecessors, such as avoiding the sequence $\log(\exp(x))$ or $\sin(\cos(x))$ in traversals due to its meaningless (Petersen et al. 2021). The improved action mask inherits the capability to filter out noisy variables, thereby significantly reducing the complexity of the search space. This is achieved by applying the trained gate G to the action mask, which can be represented as:

$$A_{new} = A_{original} \odot G \quad (4)$$

where $A_{original}$ denotes the original action mask and A_{new} represents the new action mask that incorporates the learned gates. The element-wise multiplication ensures that only the permissible operators and selected variables are activated during the sampling process.

The improved action mask not only maintains the original constraints but also introduces an additional selection based on the relevance of the input variables. This dual functionality facilitates a more focused and efficient search by eliminating unnecessary exploratory steps and concentrating on the most promising regions of the search space.

Generating Expressions as Training Samples

Mathematical expressions can be effectively represented by expression trees, a specific type of binary tree where internal nodes correspond to mathematical operators and terminal nodes represent input variables or constants (Petersen et al. 2021). An expression tree can be sequentially described using pre-order traversal, thereby enabling the representation of an expression through the pre-order traversal of its corresponding expression tree. The tokens in these pre-order traversals are constituted by mathematical operators, input variables, or constants. These tokens can be selected from a pre-established token library, denoted as L . This library encompasses a range of commonly used operators and given variables, such as $\{+, -, \times, \div; \sin; \cos; \log; \exp; x_1; x_2; \dots\}$. Consequently, the generation of an expression can be achieved by sequentially producing tokens along the pre-order traversal as show in Fig. 1 (d). One traversal can be represented as τ , with τ_i denoting the i_{th} token of τ . The length of the traversal is symbolized by $|\tau| = T$. Consequently, the process of generating expressions can be viewed as a sequence generation process, which can be optimized using a recurrent neural network (RNN). The RNN serves as a potent tool for generating sequences that fit expressions, as it can encapsulate all previous information in the generating traversal. The process of generating expressions via the RNN constitutes the policy π , which can be optimized by the RL algorithm.

We sample tokens from L with probabilities derived from the output of the last layer of RNN with parameters θ . The sampled tokens, arranged in sequence as the traversal $(\tau_1, \tau_2, \dots, \tau_T)$, can form the mathematical expression as shown Fig. 1 (c). Specifically, the i_{th} output of the RNN passes through a Softmax layer to produce the probability distribution for selecting the i_{th} token τ_i , conditioned on the previously selected tokens $\tau_{1:(i-1)}$. This process is

noted as $\pi_\theta(\tau_i | \tau_{1:(i-1)})$. The likelihood of sampling the entire expression is simply the product of the likelihoods of its tokens: $\pi_\theta(\tau) = \pi_\theta(\tau_1) \prod_{i=2}^T \pi_\theta(\tau_i | \tau_{1:(i-1)})$. Hence, after getting the traversal $(\tau_1, \tau_2, \dots, \tau_T)$, the expression f is generated as a training sample.

Except for generating expressions as samples, this stage also necessitates the estimation of effects corresponding to these generated expressions. These effects will serve as reward feedback for the RL training process. Given the test dataset $(X; y)$ with size n and learned expression f , the normalized root-mean-square error (NRMSE) can be used as an indicator of the fitness of the learned expression f , as calculated by $\frac{1}{\sigma_y} \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - f(X_i))^2}$, where σ_y is the standard deviation of the target values. A smaller NRMSE signifies a better performance of the learned expression. In the ideal case where NRMSE equals zero, it implies that the underlying patterns within the data have been perfectly discovered, and the corresponding expression is correct.

Reinforcement Learning with Mixed Path Entropy (MPE) Regularization

The traversal generation process can be modeled as a standard Markov Decision Process (MDP), where the policy can be learned using RL approaches. We assume that the preceding operators in the traversals, along with their properties, can be treated as *observations* during the generation process. The selection of tokens in the traversals constitutes the *actions* of the policy based on the corresponding observations. As described in the previous section, the *reward*, which serves as the policy objective for SR, is defined by minimizing the NRMSE. Consequently, the token sequences in the generated traversals represent *trajectories*, and each generation process is considered as an *episode* in the MDP.

The Proximal Policy Optimization (PPO) (Schulman et al. 2017) as a popular policy gradient algorithm (Schulman et al. 2015; Song et al. 2020) can be used to achieve policy objectives. The reward obtained by a trajectory τ , denoted as $R(\tau)$, is calculated as $1 / (1 + \text{NRMSE})$, which is the inverse of the NRMSE (INRMSE). The reward $R(\tau)$ is assigned upon the completion of an episode and is attributed to each action within that episode. Then the empirical $(1 - \eta)$ -quantile of $R(\tau)$ in batch data, represented as R_η , is used to filter the batch data to only include the top ε fraction of samples. This approach make the policy focus on learning from best-case expressions, which can increase the training efficiency (Petersen et al. 2021). The term R_η can be regarded as the baseline within the advantage function of PPO, thereby serving as a substitute for the value function in PPO.

Entropy regularization is a popular approach in policy gradient methods to prevent premature policy convergence and to encourage exploration. We employ a hierarchical entropy term to increase the randomness of the policy at each individual step (Landajuela et al. 2021b), which is described as:

$$H(\pi_\theta) = - \sum_{\tau \in \Gamma} \sum_{t=1}^T \gamma^{t-1} \sum_{\tau_t \in a} \pi_\theta(\tau_t | s_t) \log \pi_\theta(\tau_t | s_t) \quad (5)$$

where $\gamma < 1$ is an exponential decay factor for the weights. τ_t is the action sampled by the policy at the state s_t , and a is all possible actions. $H(\pi_\theta)$ can encourage the exploration of policy, especially in the earliest tokens of each trajectory to alleviate the early commitment phenomenon and initialization bias in symbolic spaces.

To further promote exploration across sequences, we introduce path entropy regularization, $H_\tau(\pi_\theta)$, which is a measure of the uncertainty of an entire sequence of actions in all sequences taken by the policy. Given that the objective of RL is to maximize the expectation of reward, during the training of RL-based SR, it is possible to prematurely converge to a local-optimum state that closely approximates the target expression, but with tokens that are entirely absent from the target expression. This could lead to continuous exploration around the local-optimum state. In this case, the likelihood of the policy exploring the correct tokens present in the target expression is significantly low. The aim of path entropy regularization is to facilitate the discovery of the perfect expression that necessitates more exploration of completely distinct sequence paths. Hence, for a given trajectory set Γ , the path entropy can be defined as:

$$H_\tau(\pi_\theta) = - \sum_{\tau \in \Gamma} \pi_\theta(\tau) \log \pi_\theta(\tau) \quad (6)$$

where $\pi_\theta(\tau)$ is the joint probability of the entire action sequence $(\tau_1, \tau_2, \dots, \tau_T)$ of trajectory τ , which is computed as:

$$\pi_\theta(\tau) = \pi_\theta(\tau_1) \prod_{t=2}^T \pi_\theta(\tau_t | \tau_{t-1}, \dots, \tau_1) \quad (7)$$

Finally, the policy objective that combines both single-step entropy term and path entropy term is formulated as:

$$\mathcal{L}(\theta) = \mathcal{L}_p(\theta) + \alpha H_\tau(\pi_\theta) + \beta H(\pi_\theta) \quad (8)$$

where α and β are hyperparameters that determines the significance of the path entropy term and single-step entropy term, respectively. This combined entropy term aims to balance the immediate exploration benefits of single-step entropy with the long-term diversity encouraged by path entropy, which can be called mixed path entropy (MPE). The pseudocode of NRSR is shown in Appendix.

Experiments

In this section, we first delineate the experimental configurations. Following this, we present the results and conduct a comprehensive analysis to validate the efficacy of our method, as well as the individual modules encompassed within it.

Experimental Configurations

Benchmark In this study, we employed the Nguyen SR benchmark suite (Uy et al. 2011) to assess our proposed method. This suite, widely used in SR research, comprises twelve representative expressions. Each benchmark is defined by a ground truth expression, an operator library, and an input variable range, all of which are detailed in Appendix. Datasets are generated using the ground truth

and the input range, and are subsequently divided into three segments: one for training the NGM, one for calculating the fitness reward $R(\tau)$ of the expressions generated during the training process, and one for evaluating the best fit expression after each training iteration. The sample sizes for these three subsets are 20,000, 20, and 20, respectively. The operator library, which restricts the operators available for use during training, is denoted by $\{+; -; \times; \div; \sin; \cos; \log; \exp; x_i\}$ in this study, with the i_{th} input variables represented by x_i .

Baselines We compared NRSR against five sota SR baselines, providing a comprehensive comparison across two RL-based methods, a GP-based method, a pre-trained method and a commercial software. The first baseline, **DSR** (Petersen et al. 2021), is a RL-based SR framework that employs an RNN with a risk-seeking PG to generate and optimize mathematical expressions, demonstrating superior performance. The second baseline (Landajuela et al. 2021b) builds upon DSR by incorporating a hierarchical-entropy regularizer and a soft-length prior, which is noted by **HESL** in this paper. This enhancement mitigates early commitment and initialization bias, thereby improving exploration and performance. The third, **GP-Meld** (Landajuela et al. 2022), is an SR method that combines GP with DSR. It uses GP as an inner optimization loop, augmenting the exploration of the search space while addressing the non-parametric limitations of GP with DSR’s neural network. The fourth baseline, **DGSR** (Holt, Qian, and van der Schaar 2023), utilizes pre-trained deep generative models to exploit the inherent regularities of equations, enhancing the effectiveness of SR. It has excellent performance in terms of recovering true equations and computational efficiency, particularly dealing with a large number of input variables. The final baseline is **Eureqa** (White 2012), a widely-used commercial software based on a GP-based approach (Schmidt and Lipson 2009), serving as the gold standard for SR.

Training Process In the SR process for each benchmark, the NGM is initially trained to obtain the noise-resilient gating layer, a process that constitutes a regression task. During the RL training phase, the acquired gates G are employed to filter out noisy input variables. In the expression sampling phase, the RL policy sequentially generates tokens to produce a batch of trajectories. These trajectories are subsequently transformed into expressions to compute the fitness reward $R(\tau)$. For each iteration of policy training, the trajectory data, reward $R(\tau)$, and information about the generated expressions are utilized to train new policies. NRSR and the other comparative methods were implemented within a unified SR framework (Landajuela et al. 2022), with the exception of Eureqa, which was executed using the API interface of the DataRobot platform¹. Detailed specifications of the training settings can be found in Appendix.

Results

SR Performance on high-noise data The performance of NRSR and five comparative baselines is evaluated us-

¹[Online]. Available: <https://docs.datarobot.com>

Methods	RR \uparrow	EEN \downarrow	NMSE \downarrow
(a) with 5-noise data			
DSP	61.3%	955K	0.0352
HESL	63.9%	895K	0.0300
GP-Meld	51.0%	1.63M	0.0491
DGSR	72.5%	712K	0.0101
Eureqa	35.0%	NaN	0.176
NRSR	89.1%	425K	7.73e-3
(b) with 10-noise data			
DSP	23.2%	1.63M	0.138
HESL	26.4%	1.58M	0.123
GP-Meld	35.1%	1.80M	0.0718
DGSR	68.5%	864K	0.0121
Eureqa	34.3%	NaN	0.282
NRSR	89.1%	423K	8.52e-3

Table 1: Comparison of average RR, EEN, and NMSE between NRSR and five baseline methods across all benchmarks in high-noise data scenarios.

ing three metrics: recovery rate (RR), explored expression number (EEN), and normalized mean-square error (NMSE). RR quantifies the likelihood of identifying perfect expressions across all replicated tests under varying random seeds. EEN represents the average number of expressions examined across all replicated tests. A lower EEN indicates greater efficiency of a method in discovering the correct expression with fewer training resources. EEN is particularly critical for SR tasks, as they are conjectured to be NP-hard. NMSE measures the average fitness discrepancy between the ground-truth expression and the best-found expression across all replicated tests. Table 1 presents the average performance of six SR methods across all benchmarks when applied to high-noise data. Our proposed method, NRSR, significantly outperforms the five baseline methods in terms of RR, EEN, and NMSE. The results highlight a substantial degradation in the performance of the baseline methods when five noisy inputs are introduced. This performance decline becomes even more noticeable as the number of noisy inputs escalates to ten. Contrarily, NRSR maintains high performance levels even in the presence of high-noise data, thereby demonstrating its robustness against high-noise interference. Unless otherwise specified, all results reported in this study are the average of 100 replicated tests, each with different random seeds, for each benchmark expression.

Ablation Studies Ablation studies were conducted to estimate the individual contributions of the critical components within NRSR. As expounded in Section 3, the NGM and the MPE emerged as pivotal constituents. Furthermore, the efficacy of the PPO algorithm was also appraised against the traditional PG approach. To ensure a precise evaluation of each component’s influence, the NGM was specifically tested under noisy data conditions, whereas the performance metrics for the remaining modules were obtained using clean data. These ablation tests were carried out utilizing the Nguyen benchmark suite, employing RR, EEN, and NMSE as the metrics, with the results presented in Table 2. The outcomes demonstrate that the NGM significantly enhances SR performance on high-noise data scenarios. The

Ablations	RR \uparrow	EEN \downarrow	NMSE \downarrow
(a) on high-noise data			
NRSR	89.1%	423K	8.52e-3
No NGM	32.9%	1.45M	0.093
No NGM / MPE / PPO	28.3%	1.62M	0.112
(b) on clean data			
NRSR	89.7%	408K	7.66e-3
No MPE	85.3%	431K	5.98e-3
No PPO	86.7%	462K	8.92e-3
No MPE / PPO	84.1%	480K	9.60e-3

Table 2: Comparison of average RR, EEN, and NMSE across all benchmarks for various ablations of NRSR.

omission of MPE and PPO results in diminished performance across both high-noise and clean data contexts, representing the advantageous role of MPE and PPO in augmenting the RL training process for SR tasks. In the ablation tests, the high-noise data have 10 noisy input variables.

Analysis of noise-resilient gating module The NGM’s performance is crucial for SR in scenarios with high-noise data. Hence, a series of experiments were devised to assess the reliability of the gating mechanism under a variety of parameter configurations. Our evaluation are conducted on 93 expressions, which are detailed in the Supplementary File. We introduced four levels of noisy inputs across a comprehensive set of all benchmarks. The gating layers were then employed to filter the input variables, with the perfect filter rate being the primary metric of interest. The results, presented in Table 3, reveal that the gating layer adeptly eliminates noisy input variables with a 98.92% accuracy when the number of noise variables does not exceed ten. This mechanism demonstrates robust and stable performance, retaining high accuracy even with the introduction of up to twenty noisy input variables. Further investigation was conducted to compare the efficacy of employing the gating layer extracted from the final training epoch against an averaged gating layer computed across all epochs. The empirical evidence suggests that utilizing the averaged gating layer in conjunction with the Otsu threshold (Otsu 1979) significantly surpasses the approach using the gates on the final epoch. This enhancement is likely due to the averaged approach’s stability, which mitigates the variability inherent in the results from the final epoch. Moreover, we examined the impact of adjusting the Otsu threshold scale and the L0 gating loss coefficient λ . The optimal performance was achieved by scaling the Otsu threshold to 1.05 times its original value, facilitating more effective noise filtration. However, further scaling diminishes the benefits, as an excessively high threshold risks discarding true input variables. It is noteworthy that, with the exception of the first test set, the number of noisy input variables was fixed at ten for these evaluations.

Analysis of mixed path entropy In the context of SR, the MPE is instrumental in promoting the exploration of diverse complete sequences. To augment our understanding of MPE’s efficacy in SR, we introduce a metric named as the effective exploration ratio (EER). EER is calculated as the ratio of the unique expression number (UEN) to the EEN,

Settings	Accuracy
with 3 noisy input variables	98.92%
with 5 noisy input variables	98.92%
with 10 noisy input variables	98.92%
with 20 noisy input variables	96.77%
using the gating layer on the final epoch	60.21%
Otsu threshold scale 1.0	95.70%
Otsu threshold scale 1.05	98.92%
Otsu threshold scale 1.1	94.62%
gating loss λ 0.1	96.77%
gating loss λ 0.25	98.92%
gating loss λ 0.5	96.77%

Table 3: Comparison of accuracy results under different training settings for NGM.

wherein UEN is the amount of unique expressions generated through successful explorations. If the exploration fails to accurately reconstruct the target expression, the corresponding UEN is recorded as zero. EEN represents the amount of expression generated per task. Hence, a higher EER value indicates a more efficient exploration process. The evaluation of EER across NRSR and four comparative baselines, conducted on the Nguyen-5 benchmark, is presented in Table 4. The results indicate that NRSR has the highest EER, suggesting its superior proficiency in exploration efficiency within the SR framework.

To analyze the impact of the MPE on SR, we conducted a series of tests adjusting the parameter β within the MPE on the Nguyen-7 benchmark. The outcomes are represented in Fig. 2(a). The β is incremented from 0 to 0.06, there is a notable inflection in the performance. Specifically, the RR exhibits an initial ascent, followed by a descent. Concurrently, the EEN manifests a converse trend, initially presenting a decline, which then transitions into an ascent. This pattern implies a trade-off inherent in the MPE: it has the potential to improve exploratory behavior, yet an overly large β may heighten the computational expenditure of the algorithm. Further insights are provided in Fig. 2(b), which illustrates the dynamics of total entropy during the training phase on the Nguyen-12 benchmark. It demonstrates that the employment of MPE will result in a higher total entropy compared to excluding path entropy throughout the training duration, while simultaneously maintaining superior SR performance. These results were derived from ten independent tests, each conducted with a random seed.

SR Performance on clean data To thoroughly assess the performance of our proposed NRSR, we conducted the benchmark test against established baseline methods using

Matrix	NRSR	DSR	HESL	GP-Meld	DGSR
UEN	524.3K	718.4K	651.7K	104.7K	107.5K
EEN	735.2K	1132.4K	988.1K	1853.7K	195.6K
EER	0.713	0.634	0.660	0.056	0.550

Table 4: Comparison of UEN, EEN, and EER results between our proposed method and four baselines on the Nguyen-5 benchmark.

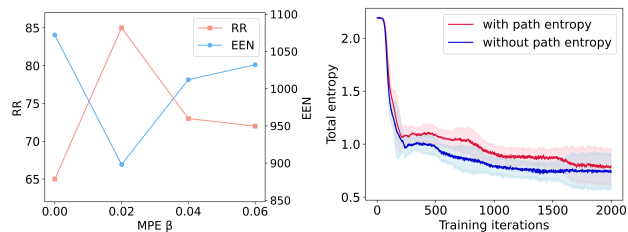


Figure 2: (a) Variation of RR and EEN metrics performance with respect to β in MPE. (b) Dynamics of total entropy during the training phase.

Methods	RR \uparrow	EEN \downarrow	NMSE \downarrow
DSP	83.2%	540K	8.75e-3
HESL	88.3%	441K	7.94e-3
GP-Meld	80.2%	879K	5.75e-3
DGSR	77.7%	507K	1.19e-2
Eureqa	67.4%	NaN	1.97e-3
NRSR	89.7%	408K	7.66e-3

Table 5: The comparison of average RR, EEN and NMSE for our proposed method and five baselines on all benchmarks with clean data.

a dataset without noises. The results, detailed in Table 5, show the comparative effectiveness of various SR methods on clean data. Notably, NRSR outperforms the five baseline methods on average RR and EEN. This demonstrates NRSR’s robustness and its ability to deliver superior performance as an independent SR method on clean data. Ablation studies partly illuminate the sources of this improved performance of our proposed method. Specifically, the adoption of the PPO algorithm and the MPE bonus are identified as significant contributors to the method’s success. While Eureqa exhibits superior performance in terms of NMSE, it registers the lowest RR. This discrepancy may stem from the employed algorithms in Eureqa.

Conclusion

We present an innovative symbolic regression (SR) method that demonstrates excellent capability to precisely recover expressions from data with high-noise, outperforming existing state-of-the-art (sota) methods on a comprehensive set of benchmark tasks. Our method utilizes a noise-resilient gating module (NGM) to filter out noisy information and employs reinforcement learning approach to develop policies for expression recovery. We also introduce the mixed path entropy (MPE), a novel policy bonus designed to enhance the exploration of expressions during training. Our experimental results demonstrate that our method not only handles high-noise data with superior performance but also achieves sota results on benchmarks with clean data. This shows the robustness of our method in various data quality scenarios. Importantly, the NGM and MPE are designed as modular elements, making them suitable for integration into other SR frameworks, thus expanding their potential utility and contributing to the advancement of SR methodologies.

Acknowledgments

We would like to express our sincere gratitude to the reviewers for their insightful comments and valuable suggestions. We also thank our colleagues at the Cooperation Product Department, Interactive Entertainment Group, Tencent for their helpful discussions and support throughout this work.

References

- Buhrmester, V.; Münch, D.; and Arens, M. 2021. Analysis of explainers of black box deep neural networks for computer vision: A survey. *Machine Learning and Knowledge Extraction*, 3(4): 966–989.
- Cortes, C.; Mohri, M.; and Rostamizadeh, A. 2012. L2 regularization for learning kernels. *arXiv preprint arXiv:1205.2653*.
- Crochepierre, L.; Boudjeloud-Assala, L.; and Barbesant, V. 2022. Interactive Reinforcement Learning for Symbolic Regression from Multi-Format Human-Preference Feedbacks. In *IJCAI*, 5900–5903.
- De Melo, V. V. 2014. Kaizen programming. In *Proceedings of the 2014 annual conference on genetic and evolutionary computation*, 895–902.
- d’Ascoli, S.; Kamienny, P.-A.; Lample, G.; and Charton, F. 2022. Deep symbolic regression for recurrence prediction. In *International Conference on Machine Learning*, 4520–4536. PMLR.
- Ecoffet, A.; Huizinga, J.; Lehman, J.; Stanley, K. O.; and Clune, J. 2019. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.
- Holt, S.; Qian, Z.; and van der Schaar, M. 2023. Deep Generative Symbolic Regression. In *The Eleventh International Conference on Learning Representations*.
- Hou, Y.; Liang, X.; Zhang, J.; Yang, Q.; Yang, A.; and Wang, N. 2023. Exploring the use of invalid action masking in reinforcement learning: A comparative study of on-policy and off-policy algorithms in real-time strategy games. *Applied Sciences*, 13(14): 8283.
- Jiang, N.; and Xue, Y. 2023. Symbolic regression via control variable genetic programming. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 178–195. Springer.
- Kamienny, P.-A.; d’Ascoli, S.; Lample, G.; and Charton, F. 2022. End-to-end symbolic regression with transformers. *Advances in Neural Information Processing Systems*, 35: 10269–10281.
- Kamienny, P.-A.; Lample, G.; Lamprier, S.; and Virgolin, M. 2023. Deep generative symbolic regression with Monte-Carlo-tree-search. In *International Conference on Machine Learning*, 15655–15668. PMLR.
- Kim, S.; Lu, P. Y.; Mukherjee, S.; Gilbert, M.; Jing, L.; Čeperić, V.; and Soljačić, M. 2020. 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.
- Landajuela, M.; Lee, C.; Yang, J.; Glatt, R.; Santiago, C. P.; Aravena, I.; et al. 2022. A Unified Framework for Deep Symbolic Regression. In *Advances in Neural Information Processing Systems*.
- Landajuela, M.; Petersen, B. K.; Kim, S.; Santiago, C. P.; Glatt, R.; Mundhenk, N.; Pettit, J. F.; and Faissol, D. 2021a. Discovering symbolic policies with deep reinforcement learning. In *International Conference on Machine Learning*, 5979–5989. PMLR.
- Landajuela, M.; Petersen, B. K.; Kim, S. K.; Santiago, C. P.; Glatt, R.; Mundhenk, T. N.; et al. 2021b. Improving exploration in policy gradient search: Application to symbolic optimization. In *1st Mathematical Reasoning in General Artificial Intelligence Workshop, ICLR*.
- Louizos, C.; Welling, M.; and Kingma, D. P. 2018. Learning Sparse Neural Networks through L0 Regularization. In *International Conference on Learning Representations*.
- Ma, R.; Miao, J.; Niu, L.; and Zhang, P. 2019. Transformed L1 regularization for learning sparse deep neural networks. *Neural Networks*, 119: 286–298.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; et al. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937. PMLR.
- Mundhenk, T.; Landajuela, M.; Glatt, R.; Santiago, C. P.; Petersen, B. K.; et al. 2021a. Symbolic regression via deep reinforcement learning enhanced genetic programming seeding. *Advances in Neural Information Processing Systems*, 34: 24912–24923.
- Mundhenk, T. N.; Landajuela, M.; Glatt, R.; Santiago, C. P.; Faissol, D. M.; and Petersen, B. K. 2021b. Symbolic Regression via Neural-Guided Genetic Programming Population Seeding. In *Advances in Neural Information Processing Systems*.
- Muthukumar, V.; Narang, A.; Subramanian, V.; Belkin, M.; Hsu, D.; and Sahai, A. 2021. Classification vs regression in overparameterized regimes: Does the loss function matter? *Journal of Machine Learning Research*, 22(222): 1–69.
- Mysore, S.; Cheng, G.; Zhao, Y.; Saenko, K.; and Wu, M. 2022. Multi-Critic Actor Learning: Teaching RL Policies to Act with Style. In *International Conference on Learning Representations*.
- Ng, A. Y. 2004. Feature selection, L1 vs. L2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, 78.
- Nguyen, T. T.; Nguyen, N. D.; and Nahavandi, S. 2020. Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications. *IEEE transactions on cybernetics*, 50(9): 3826–3839.
- Oliveira, L. O. V.; Martins, J. F. B.; Miranda, L. F.; and Pappa, G. L. 2018. Analysing symbolic regression benchmarks under a meta-learning approach. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1342–1349.

- Otsu, N. 1979. A Threshold Selection Method from Gray-Level Histograms. *IEEE Trans. Syst. Man Cybern.*, 9: 62–66.
- Petersen, B. K.; Landajuela, M.; Mundhenk, T. N.; Santiago, C. P.; Kim, S. K.; et al. 2021. Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients. In *Proc. of the International Conference on Learning Representations*.
- Reinbold, P. A.; Kageorge, L. M.; Schatz, M. F.; and Grigoriou, R. O. 2021. Robust learning from noisy, incomplete, high-dimensional experimental data via physically constrained symbolic regression. *Nature communications*, 12(1): 3219.
- Schmidt, M.; and Lipson, H. 2009. Distilling free-form natural laws from experimental data. *science*, 324(5923): 81–85.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Searson, D. P.; Leahy, D. E.; and Willis, M. J. 2010. GP-TIPS: an open source genetic programming toolbox for multigene symbolic regression. In *Proceedings of the International multiconference of engineers and computer scientists*, volume 1, 77–80. Citeseer.
- Shu, X.; and Ye, Y. 2023. Knowledge Discovery: Methods from data mining and machine learning. *Social Science Research*, 110: 102817.
- Song, H. F.; Abdolmaleki, A.; Springenberg, J. T.; Clark, A.; Soyer, H.; Rae, J. W.; et al. 2020. V-MPO: On-Policy Maximum a Posteriori Policy Optimization for Discrete and Continuous Control. In *International Conference on Learning Representations*.
- Sorscher, B.; Geirhos, R.; Shekhar, S.; Ganguli, S.; and Morcos, A. 2022. Beyond neural scaling laws: beating power law scaling via data pruning. *Advances in Neural Information Processing Systems*, 35: 19523–19536.
- Tang, C.-Y.; Liu, C.-H.; Chen, W.-K.; and You, S. D. 2020. Implementing action mask in proximal policy optimization (PPO) algorithm. *ICT Express*, 6(3): 200–203.
- 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. *Advances in Neural Information Processing Systems*, 33: 4860–4871.
- Udrescu, S.-M.; and Tegmark, M. 2020. AI Feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16): eaay2631.
- Uy, N. Q.; Hoai, N. X.; O’Neill, M.; McKay, R. I.; and Galván-López, E. 2011. Semantically-based crossover in genetic programming: application to real-valued symbolic regression. *Genetic Programming and Evolvable Machines*, 12: 91–119.
- Valipour, M.; You, B.; Panju, M.; and Ghodsi, A. 2021. Symbolicgpt: A generative transformer model for symbolic regression. *arXiv preprint arXiv:2106.14131*.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.
- Wang, X.; Wang, S.; Liang, X.; Zhao, D.; Huang, J.; Xu, X.; Dai, B.; and Miao, Q. 2022. Deep reinforcement learning: a survey. *IEEE Transactions on Neural Networks and Learning Systems*.
- Wei, Z.; Li, Q.; Wei, J.; and Bian, W. 2022. Neural network for a class of sparse optimization with L0-regularization. *Neural Networks*, 151: 211–221.
- White, D. R. 2012. Software review: the ECJ toolkit.
- Ying, X. 2019. An overview of overfitting and its solutions. In *Journal of physics: Conference series*, volume 1168, 022022. IOP Publishing.
- Zhang, H.; and Zhou, A. 2021. RL-GEP: symbolic regression via gene expression programming and reinforcement learning. In *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.