



Made available through Montana State University's [ScholarWorks](#)

Univariate Skeleton Prediction in Multivariate Systems Using Transformers

Giorgio Morales & John W. Sheppard



Copyright Springer Nature [academic journals on nature.com]

Accessibility Disclaimer:

For a more accessible version of this document, please submit an accessibility request form through the Montana State University Library website.



Univariate Skeleton Prediction in Multivariate Systems Using Transformers

Giorgio Morales^(✉)  and John W. Sheppard 

Gianforte School of Computing, Montana State University, Bozeman, MT, USA
giorgiomorales@ieee.org, john.sheppard@montana.edu

Abstract. Symbolic regression (SR) methods attempt to learn mathematical expressions that approximate the behavior of an observed system. However, when dealing with multivariate systems, they often fail to identify the functional form that explains the relationship between each variable and the system's response. To begin to address this, we propose an explainable neural SR method that generates univariate symbolic skeletons that aim to explain how each variable influences the system's response. By analyzing multiple sets of data generated artificially, where one input variable varies while others are fixed, relationships are modeled separately for each input variable. The response of such artificial data sets is estimated using a regression neural network (NN). Finally, the multiple sets of input–response pairs are processed by a pre-trained Multi-Set Transformer that solves a problem we termed Multi-Set Skeleton Prediction and outputs a univariate symbolic skeleton. Thus, such skeletons represent explanations of the function approximated by the regression NN. Experimental results demonstrate that this method learns skeleton expressions matching the underlying functions and outperforms two GP-based and two neural SR methods.

Keywords: Symbolic regression · Transformer networks · Symbolic skeletons · Multivariate regression · Explainable artificial intelligence

1 Introduction

Symbolic regression (SR) aims to identify mathematical equations or symbolic expressions that capture the underlying relationships and dynamics of the studied phenomena [18]. One of the main advantages of the expressions learned by an SR model is that they are interpretable by humans and allow for the identification of cause–effect relationships between the inputs and outputs of a system [10]. These techniques not only capture the behavior of empirical data into analytical equations but also reduce the computational complexity during the inference phase, and have more powerful extrapolation ability than black-box models [22].

A limitation of existing SR approaches lies in their primary focus on minimizing prediction errors rather than distilling the underlying equations that govern the system dynamics [4]. Consequently, the generated equations may exhibit

high complexity, effectively approximating the observed data but failing to correspond to the underlying equations [18]. This limitation poses challenges when performing out-of-sample inference, where generalization may be ineffective.

SR methods are often based on population-based algorithms, especially genetic programming (GP) [24]. Nevertheless, a notable drawback of GP-based SR methods is that they suffer from slow computation. Their computational inefficiency is related to the inherent complexity of the search space and the expensive iterative calls to numerical optimization routines after each generation [24]. In addition, these methods do not consider past experiences, as they require learning each problem from scratch. As such, the obtained models do not benefit from additional data or insights from different equations, hindering their capacity for improvement and limiting their generalization capabilities [4, 17].

In response to these limitations, recent advancements have seen the emergence of neural SR methods as a promising alternative. These methods may utilize general-purpose pre-trained models that generate symbolic expressions through a straightforward forward pass and, possibly, a single call to a numerical optimization routine [4, 17, 29]. Thus, neural SR methods offer a substantial time speedup compared to GP-based approaches. However, there remains a gap in prediction accuracy between the two paradigms [17].

Although recent neural SR methods may achieve low prediction errors when dealing with multivariate systems, they struggle to represent the dependency of the system’s response and each independent variable correctly. As a first step in addressing this, our proposed method aims to generate univariate symbolic skeletons that estimate functional relationships for each variable with respect to the response. A symbolic skeleton expression is an abstract representation of a mathematical expression that captures its structural form without setting specific numerical values. In other words, we propose to generate univariate skeletons that constitute mathematical “explanations” describing the interaction between each independent variable and the system’s response.

Given a multivariate regression problem that can be expressed in terms of a mathematical equation, our method identifies univariate symbolic skeleton expressions for each variable. The skeleton prediction process begins by training a regression model, such as a neural network (NN), that approximates the system’s function. This regression model is used to estimate the response of multiple randomly generated sets of points where only the variable of interest is allowed to vary while the remaining variables are held fixed. We then define a novel Multi-Set Transformer model, which is pre-trained on a large dataset of synthetic symbolic expressions, that receives as inputs the sets generated by the NN and is used to find the univariate skeletons.

Contributions. We hypothesize that our method will generate univariate skeletons that are more similar to those corresponding to the underlying equations in comparison to other SR methods. The generated skeletons represent explanations of how each variable is related to the system’s response. Thus, from an explainability standpoint, producing more faithful univariate skeletons means

that we can offer better explanations about the system’s behavior. In addition, the generated skeletons may be taken as building blocks to be used to estimate the overall function of the system. Our main contributions are summarized as:

- We introduce a method that learns univariate symbolic skeletons that explain the functional form between each independent variable and the system’s response using regression NNs and pre-trained transformer models.
- We introduce a new SR problem called Multi-Set symbolic skeleton prediction (MSSP). It receives multiple sets of input–response pairs, where all sets correspond to the same functional form but use different equation constants, and outputs a common skeleton expression.
- We present a novel transformer model called “Multi-Set Transformer” to solve the MSSP problem. The model is pre-trained on a large dataset of synthetic symbolic expressions.

2 Related Work

SR is commonly tackled using GP-based methods. Such methods evolve a population of tree-like individuals using operations like selection, crossover, and mutation to improve their fitness over multiple generations. Each individual represents a symbolic expression that maps the inputs and the output, and its fitness function determines how well it fits the data set being modeled. Variations of this approach attempt to design improved operators and fitness functions to reduce the complexity of the search [12, 24].

Two significant challenges of GP for SR are code growth (i.e., bloat) and the huge search space [1]. Code growth leads to computationally expensive evolution of large programs and hinders their generalization ability [8]. The huge search space is attributed to the variability in program lengths permitted during the evolutionary process. This allows for the generation of multiple solution trees representing mathematically equivalent functions [9]. However, GP tends to generate a greater number of large solution trees compared to smaller ones [20]. The high complexity of large solutions entails poor generalization performance. In most GP-based methods, the program fitness is determined by its overall output and not by the intermediate outputs of its subexpressions. Thus, a program’s subexpressions are optimized indirectly. Arnaldo *et al.* [2] pointed out that the indirect optimization approach may result in suboptimal fitness optimization because it does not focus on evolving and learning suitable building blocks.

Furthermore, Martius and Lampert [22], Sahoo *et al.* [26], and Werner *et al.* [31] proposed the use of NN architectures whose units represent conventional mathematical operators. The objective is to reduce the prediction error and gradually prune irrelevant parts of the network until a simple equation can be extracted from the network. However, similar to the case of GP-based methods, the learned expressions tend to be large and, as a consequence, difficult to explain. Another limitation is that each problem is learned from scratch.

As a recent alternative, a few methods based on large language models have been proposed. Biggio *et al.* [4] introduced the use of pre-trained transformer models for SR. A dataset of multivariate equations is generated to pre-train a model based on a Set Transformer [21], and this transformer acts as a general-purpose model to predict the symbolic skeleton from a corresponding set of input-output pairs. The skeleton’s constants are then fit using a non-convex optimizer, such as the BFGS algorithm [11].

Kamienny *et al.* [17] pointed out that the loss function minimized by the BFGS algorithm can be highly non-convex and the correct constants of the skeletons are not guaranteed to be found. They avoided performing skeleton prediction as an intermediary step and proposed an end-to-end (E2E) transformer model that estimates the full mathematical expression directly. The learned constants can then be refined using a non-convex optimizer. Experiments showed that E2E performs better than previous NN-based methods. Since previous methods have shown scalability issues when dealing with equations having many variables, Chu *et al.* [6] proposed a method that decomposes multi-variable SR into a sequence of single-variable SR problems, combined in a bottom-up manner. The process involves learning a data generator using NNs from observed data, generating variable-specific samples with controlled input variables, applying single-variable SR, and iteratively adding variables until completion.

3 Multivariate Skeleton Prediction

Consider a system whose response $y \in \mathbb{R}$ is sensitive to variables $\mathbf{x} = \{x_1, \dots, x_t\}$ ($\mathbf{x} \in \mathbb{R}^t$). The underlying function that maps the feature value space and the response value space is denoted as $f : \mathcal{X}^t \rightarrow \mathcal{Y}$, such that $y = f(\mathbf{x}) = f(x_1, \dots, x_t)$. In addition, let $\kappa(\cdot)$ represent a skeleton function that replaces the numerical constants of a given symbolic expression by placeholders c_i ; e.g., $\kappa(3x^2 + e^{2x} - 4) = c_1 x^2 + e^{c_2 x} + c_3$. We assume f can be expressed as a mathematical expression with unary (e.g., \sin , \cos , and \log) and binary operators (e.g. $+$, $-$, $*$, and $/$). We aim to obtain the univariate skeletons $\hat{e}(x_1), \dots, \hat{e}(x_t)$ that describe the functional form between each variable and the system’s response. In the following, we lay out the steps proposed in this method.

3.1 Neural Network Training

The underlying function f can be approximated based on observed data using any regression model. Let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{N_R}\}$ be a data set with N_R samples, where each sample is denoted as $\mathbf{x}_j = \{x_{j,1}, \dots, x_{j,t}\}$, and $\mathbf{y} = \{y_1, \dots, y_{N_R}\}$ is the set of corresponding target observations. A NN regression model, whose compute function is denoted as $\hat{f}(\cdot; \boldsymbol{\theta}_{NN})$ ($\boldsymbol{\theta}_{NN}$ represents the weights of the network), is constructed to capture the association between \mathbf{X} and \mathbf{y} . Thus, a target estimate for a given input \mathbf{x}_j is computed as $\hat{y}_j = \hat{f}(\mathbf{x}_j, \boldsymbol{\theta}_{NN})$ or, simply, $\hat{y}_j = \hat{f}(\mathbf{x}_j)$. The parameters $\boldsymbol{\theta}_{NN}$ of the network \hat{y} are obtained by minimizing the mean squared error of the predictions; that is, $\boldsymbol{\theta}_{NN}^* =$




Fig. 1. x_1 vs. y curves when $x_2 = 4.45$, 0.2 , and 1.13 .

$\operatorname{argmin}_{\theta_{NN}} \frac{1}{N_R} \sum_{j=1}^{N_R} (\hat{y}_j - y_j)^2$. Note that we selected an NN to generate the function \hat{f} due to its ease of training and high accuracy; however, other regression methods could be applied.

3.2 Multi-set Symbolic Skeleton Prediction

We tackle the SR problem by decomposing it into single-variable sub-problems. To do so, we deviate from the Symbolic Skeleton Prediction (SSP) problem that has been explored previously by existing research [4, 25, 29]. We explain the rationale behind this decision through the use of an example: Consider the function $y = x_1 / \log(x_1^2 + x_2)$. We analyze the relationship between x_1 and y . Similar to the work presented by Chu *et al.* [6], the variables that are not being analyzed are held fixed. Figure 1 depicts that the function’s behavior differs when using three different x_2 values. As such, conventional SSP solvers may produce expressions with different functional forms depending on the fixed values of the variables not currently under analysis.

What is more, the fixed values of the remaining variables could project the function into a space where the functional form is not easily identifiable, which could be worsened due to the limited range of values that the analyzed variable can take. Therefore, the SSP problem would benefit from the injection of additional context data. When analyzing the variable x_v , we could use multiple sets of input–response pairs, each of which is constructed using different fixed values for the variables $\mathbf{x} \setminus \{x_v\}$. The key idea is to process the information from these sets simultaneously to produce a skeleton that is common to all input sets. We refer to this problem as Multi-Set Symbolic Skeleton Prediction (MSSP).

More formally, assume we are given a data set of N_R input–response pairs (\mathbf{X}, \mathbf{y}) where $\mathbf{X} \in \mathbb{R}^{N_R \times t}$ and $\mathbf{y} \in \mathbb{R}^{N_R \times 1}$. Suppose we analyze the relationship between the v -th variable of the system, x_v (i.e., $v \in [1, \dots, t]$), and the response variable y . We construct a collection of N_S sets, denoted as $\mathbf{D} = \{\mathbf{D}^{(1)}, \dots, \mathbf{D}^{(N_S)}\}$. Each set $\mathbf{D}^{(s)}$ comprises n input–response pairs such that $\mathbf{D}^{(s)} = (\mathbf{X}_v^{(s)}, f(\mathbf{X}^{(s)})) = (\mathbf{X}_v^{(s)}, \mathbf{y}^{(s)})$, where $\mathbf{X}^{(s)} \in \mathbb{R}^{n \times t}$, $\mathbf{y}^{(s)} \in \mathbb{R}^n$, and $\mathbf{X}_v^{(s)}$ denotes the v -th column of $\mathbf{X}^{(s)}$ (i.e., the data corresponding to the x_v variable). $\mathbf{X}^{(s)}$ is constructed in such a way that each of the variables in $\mathbf{x} \setminus \{x_v\}$ are assigned random values and held fixed for all samples.

$\mathbf{X}^{(s)}$ can be constructed by selecting n samples from \mathbf{X} that meet the criterion that the variable x_v is allowed to vary while the variables in $\mathbf{x} \setminus \{x_v\}$ are fixed.

In this case, the set $\mathbf{y}^{(s)}$ is known as it is the corresponding set of observed response values of $\mathbf{X}^{(s)}$. Conversely, if the set \mathbf{X} is not large enough, $\mathbf{X}^{(s)}$ can be generated artificially using the regression model, and the corresponding set of response values would be obtained as $\mathbf{y}^{(s)} = f(\mathbf{X}^{(s)})$. Since f is not known *a priori*, the black-box model \hat{f} that has been trained to approximate f (i.e., $\hat{f}(\mathbf{x}) \approx f(\mathbf{x})$) is used to estimate the response of $\mathbf{X}^{(s)}$ (cf. Sect. 3.1).

Since the variables in $\mathbf{x} \setminus \{x_v\}$ have been fixed to constant values to construct each set $\mathbf{D}^{(s)}$, the underlying function that explains the relationship between $\mathbf{X}_v^{(s)}$ and $\mathbf{y}^{(s)}$ could be expressed solely in terms of variable x_v . As such, the underlying function of the s -th set is denoted by $f^{(s)}(x_v)$. It is important to note that functions $f^{(1)}(x_v), \dots, f^{(N_S)}(x_v)$ have been derived from the same function $f(\mathbf{x})$ and only differ in their coefficient values due to the selection of different values for the variables in $\mathbf{x} \setminus \{x_v\}$ for each constructed set. Thus, if we apply the skeleton function $\kappa(\cdot)$ to functions $f^{(1)}(x_v), \dots, f^{(N_S)}(x_v)$, they should produce the same target symbolic skeleton $\mathbf{e}(x_v)$, in which the constant values have been replaced by placeholders; i.e., $\mathbf{e}(x_v) = \kappa(f^{(1)}(x_v)) = \dots = \kappa(f^{(N_S)}(x_v))$.

Then, the collection \mathbf{D} is fed as an input to the MSSP problem. The objective is to generate an skeleton $\hat{\mathbf{e}}(x_v)$ that characterizes the functional form of all input sets, and approximates the target skeleton $\mathbf{e}(x_v)$; i.e., $\hat{\mathbf{e}}(x_v) \approx \mathbf{e}(x_v)$. For the sake of generality, we define the MSSP problem as follows:

Definition 1. *The input of the Multi-Set Symbolic Skeleton Prediction (MSSP) problem consists of a collection of N_S sets, denoted as $\mathbf{D} = \{\mathbf{D}^{(1)}, \dots, \mathbf{D}^{(N_S)}\}$. Each set $\mathbf{D}^{(s)}$ comprises n input–response pairs such that $\mathbf{D}^{(s)} = (\mathbf{X}_v^{(s)}, \mathbf{y}^{(s)})$, where $\mathbf{X}_v^{(s)} \in \mathbb{R}^n$ and $\mathbf{y}^{(s)} \in \mathbb{R}^n$. The underlying function of the s -th set is denoted by $f^{(s)}$; i.e., $\mathbf{y}^{(s)} = f^{(s)}(\mathbf{X}_v^{(s)})$. The underlying functions of all input sets are assumed to share a common unknown symbolic skeleton, denoted as \mathbf{e} . Thus, the objective of the MSSP problem is to generate a symbolic skeleton $\hat{\mathbf{e}} \approx \mathbf{e}$ that characterizes the functional form of all input sets \mathbf{D} .*

Multi-set Transformer. Our method for solving the MSSP problem draws inspiration from the Set Transformer [21], which is an attention-based neural network that is derived from the transformer model [30]. The Set Transformer is designed for set-input problems and, as such, it is capable of processing input sets of varying sizes and exhibits permutation invariance. We propose a Multi-Set Transformer that presents modifications to address the limitations of the Set Transformer and adapt it to the specific requirements of the proposed research. Details about the differences between the Set Transformer and our Multi-Set Transformer are given in the **Supplementary Materials**¹.

Let the s -th input set be denoted as $\mathbf{D}^{(s)} = (\mathbf{X}_v^{(s)}, \mathbf{y}^{(s)}) = \{(x_{v,i}^{(s)}, y_i^{(s)})\}_{i=1}^n$. The first step involves arranging $\mathbf{D}^{(s)}$ in a manner analogous to the input structure of the Set Transformer, which consists of a matrix where each row represents a d_{in} -dimensional element of the input set. $\mathbf{S}^{(s)} \in \mathbb{R}^{n \times d_{in}}$ denotes the s -th input

¹ Please visit <https://github.com/NISL-MSU/MultiSetSR> for the supplementary file.

of our proposed Multi-Set Transformer such that its i -th row, $\mathbf{s}_i^{(s)}$, consists of the concatenation of the input value $x_{v,i}^{(s)}$ and its corresponding output $y_i^{(s)}$; i.e., $\mathbf{s}_i^{(s)} = [x_{v,i}^{(s)}, y_i^{(s)}]$. Hence, $\mathbf{S}^{(s)}$ is defined as a matrix with $d_{in} = 2$ columns.

Our Multi-Set Transformer comprises two primary components: an encoder and a decoder. The encoder maps the information of all input sets into a unique latent representation \mathbf{Z} . To do so, an encoder stack ϕ transforms each input set $\mathbf{S}^{(s)}$ into a latent representation $\mathbf{z}^{(s)} \in \mathbb{R}^d$ (where d is context vector length or the “embedding size”) individually. Our encoder, denoted as Φ , comprises the use of the encoder stack ϕ to generate N_S individual encodings $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(N_S)}$, which are then aggregated into a unique latent representation:

$$\mathbf{Z} = \Phi(\mathbf{S}^{(1)}, \dots, \mathbf{S}^{(N_S)}, \boldsymbol{\theta}_e) = \rho(\phi(\mathbf{S}^{(1)}, \boldsymbol{\theta}_e), \dots, \phi(\mathbf{S}^{(N_S)}, \boldsymbol{\theta}_e)),$$

where $\rho(\cdot)$ is a pooling function, $\boldsymbol{\theta}_e$ represents the trainable weights of the encoder stack, and $\mathbf{z}^{(s)} = \phi(\mathbf{S}^{(N_S)}, \boldsymbol{\theta}_e)$. We define ϕ as a stack of ℓ induced set attention blocks (ISABs) [21] to encode high-order interactions among the elements of an input set in a permutation-invariant way. ISAB is a multi-head attention layer with low complexity ($\mathcal{O}(mn)$) that acts as a universal approximator of permutation invariant functions. Furthermore, we include a pooling by multi-head attention (PMA) layer [21] in ϕ to aggregate the features extracted by the ISAB blocks, whose dimensionality is $n \times d$, into a single d -dimensional latent vector. Finally, the function $\rho(\cdot)$ that is used to aggregate the latent representations $\mathbf{z}^{(s)}$ is implemented using an additional PMA layer.

The decoder ψ generates sequences conditioned on the representation \mathbf{Z} generated by Φ . This objective is aligned with that of the transformer decoder [30] and, thus, the same architecture is used for our model. ψ consists of a stack of M identical blocks, each of which is composed of three main layers: a multi-head self-attention layer, an encoder–decoder attention layer, and a position-wise feedforward network. Let $\hat{\mathbf{e}} = \{\hat{e}_1, \dots, \hat{e}_{N_{out}}\}$ denote the output sequence produced by the model, which represents the skeleton as a sequence of indexed tokens in prefix notation. Each token in this sequence is transformed into a numerical index according to a pre-defined vocabulary that contains all unique symbols. During inference, each element \hat{e}_i is generated auto-regressively; that is, ψ produces a probability distribution over the elements of the vocabulary: $\sigma(\psi(\mathbf{Z}, \boldsymbol{\theta}_d | \hat{e}_1, \dots, \hat{e}_{i-1})) = P(\hat{e}_i | \hat{e}_1, \dots, \hat{e}_{i-1}, \mathbf{Z})$, where $\boldsymbol{\theta}_d$ represents the weights of ψ . This distribution is obtained by applying a softmax function $\sigma(\cdot)$ to the decoder’s output. \hat{e}_i is thus selected as the token with the highest probability.

Multi-set Transformer Training. We train our model on a large dataset of artificially generated MSSP problems. Let $\mathbf{D}^b = \{\mathbf{D}_1, \dots, \mathbf{D}_B\}$ denote a training batch with B samples, each of which represents a collection of N_S input sets; i.e., $\mathbf{D}_j = \{\mathbf{D}_j^{(1)}, \dots, \mathbf{D}_j^{(N_S)}\}$ ($j \in [1, \dots, B]$). In addition, $\mathbf{E}^b = \{\mathbf{e}_1, \dots, \mathbf{e}_B\}$ is the corresponding set of target skeletons, each of which represents a sequence of variable length; i.e., $\mathbf{e}_j = \{e_{j,1}, \dots, e_{j,N_j}\}$ and $N_j = |\mathbf{e}_j|$. The function computed by the model is denoted as $g(\cdot)$, and Θ denotes its weights. Note that $\Theta = [\boldsymbol{\theta}_e, \boldsymbol{\theta}_d]$

contains the weights of the encoder and the decoder stacks. Given an input set collection \mathbf{D}_j , $g(\mathbf{D}_j, \Theta)$ computes the estimated skeleton $\hat{\mathbf{e}}_j$ with length $N_{out,j}$. $g(\cdot)$ is trained to generate estimated skeletons so that $\hat{\mathbf{e}}_j \approx \mathbf{e}_j$.

During training, the model is provided with past elements of the target skeleton sequence as inputs for generating subsequent tokens. Hence, the normalized probability distribution produced by the decoder ψ for the i -th element of the predicted skeleton sequence of the j -th sample would be expressed as $P(\hat{e}_{j,i}|e_{j,1}, \dots, e_{j,i-1}, \mathbf{Z}_j)$. In cases where $\hat{\mathbf{e}}_j$ and \mathbf{e}_j differ in length, we use padding and masking to ensure that the loss is only calculated for valid tokens. Thus, our optimization objective is defined as the cross-entropy loss between the padded target and predicted skeleton sequences, which is calculated using their corresponding probability distributions over the set of possible tokens:

$$\mathcal{L} = -\frac{1}{B} \sum_{j=1}^B \sum_{i=1}^{T_j} w_{j,i} P(e_{j,i}) \log P(\hat{e}_{j,i}|e_{j,1}, \dots, e_{j,i-1}, \mathbf{Z}_j),$$

where $w_{j,i}$ is 0 for padding positions and 1 otherwise, and T_j is the length of the j -th sequence. Hence, the optimization problem is expressed as: $\Theta = \operatorname{argmin}_{\Theta} \mathcal{L}$.

We generate a training dataset of expressions stored in prefix notation. The generation process is explained in Section B of the **Supplementary Materials** (see footnote 1). Our method differs from the generation method used by Biggio *et al.* [4] and Lamplé *et al.* [19]. Their approach generates several expressions that contain binary operators exclusively and, thus, are limited to simple expressions such as $\hat{e}(x) = c_1x$. Conversely, we use a method that builds the expression tree recursively in preorder, which allows us to enforce certain conditions and constraints effectively. That is, we forbid certain combinations of operators and set a maximum limit on the nesting depth of unary operators within each other.

The training routine of the Multi-Set Transformer g takes as inputs the set of pre-generated expressions, denoted as \mathbf{Q} . The skeleton \mathbf{e}_j and its associated data collection \mathbf{D}_j , corresponding to the j -th expression in \mathbf{Q} , are generated using a function called `generateSets(Q[j], N_S , n)` (Algorithm 1). The estimated skeletons $\hat{\mathbf{e}}_j$ are obtained processing the input \mathbf{D}_j and target \mathbf{e}_j through the network g using a teacher forcing strategy. $\mathcal{L}(\mathbf{E}^B, \hat{\mathbf{E}}^B)$ represents the loss function while `update(g , L)` encompasses the conventional backpropagation and stochastic gradient descent processes used to update the weights of model g .

Algorithm 1 describes the function `generateSets`, which takes as input a pre-generated expression \mathbf{ex} . Function `getConstants(ex)` returns \mathbf{c} , the list of constant placeholders in \mathbf{ex} , and $n_c = |\mathbf{c}|$. Function `selectConstants(ex, \mathbf{c} , n_f)` retrieves an expression with n_f constant placeholders selected randomly ($2 \leq n_f \leq n_c$). Then, we generate the underlying functions $f^{(s)}$ corresponding to each of the N_S sets using the function `sampleConstants(ex)`, which samples the values of each constant in \mathbf{ex} from a uniform distribution $\mathcal{U}(-10, 10)$ independently. The n input points $\mathbf{X}_v^{(s)}$ are then sampled using the function `sampleSupport(n)` from a uniform distribution $\mathcal{U}(-x^{\text{limit}}, x^{\text{limit}})$, where $x^{\text{limit}} \sim \mathcal{U}(1, 10)$ (i.e., the domain range is randomly sampled each time). The

Algorithm 1. Multi-Set Transformer Data Generation

```

1: function GENERATESETS( $\mathbf{ex}, N_S, n$ )
2:    $\mathbf{c}, n_c \leftarrow \text{getConstants}(\mathbf{ex})$ 
3:    $\mathbf{ex} \leftarrow \text{selectConstants}(\mathbf{ex}, \mathbf{c}, n_f = \text{randInt}(2, n_c))$ 
4:    $\mathbf{D}, s, \mathbf{e} \leftarrow [], 1, \kappa(\text{sampleConstants}(\mathbf{ex}))$ 
5:   while  $s \leq N_S$  do
6:      $\mathbf{X}_v^{(s)} \leftarrow \text{sampleSupport}(n)$ 
7:      $\mathbf{X}_v^{(s)}, f^{(s)}, \_ \leftarrow \text{avoidNaNs}(\mathbf{X}_v^{(s)}, \text{sampleConstants}(\mathbf{ex}))$ 
8:     if  $s > 1$  and  $\kappa(f^{(s)}) \neq \mathbf{e}$  then
9:       continue ▷ Verify that all sets correspond to the same skeleton
10:     $\mathbf{y}^{(s)} \leftarrow f^{(s)}(\mathbf{X}_v^{(s)})$ 
11:     $\mathbf{D}.\text{append}((\mathbf{X}_v^{(s)}, \mathbf{y}^{(s)}))$ 
12:     $s \leftarrow s + 1$ 
13:  return  $\mathbf{D}, \mathbf{e}$ 

```

function $\text{avoidNaNs}(\mathbf{X}_v^{(s)}, f^{(s)})$ may modify the coefficients in $f^{(s)}$ or sample additional $\mathbf{X}_v^{(s)}$ values to avoid undefined results (see Section B of the **Supplementary Materials** (see footnote 1)).

Univariate Symbolic Skeleton Prediction. Suppose we are currently analyzing the v -th variable, x_v . As an MSSP, we generate N_S artificial sets of points $\{\tilde{\mathbf{X}}^{(1)}, \dots, \tilde{\mathbf{X}}^{(N_S)}\}$ where the variable x_v is allowed to vary while the variables $\mathbf{x} \setminus \{x_v\}$ are fixed to random values. Specifically, the s -th artificial set is denoted as $\tilde{\mathbf{X}}^{(s)} = \{\tilde{\mathbf{x}}_1^{(s)}, \dots, \tilde{\mathbf{x}}_n^{(s)}\}$. The value of the v -th dimension of the j -th sample is obtained by sampling from the distribution $\mathcal{U}(x_v^{\min}, x_v^{\max})$ whose lower and upper bounds, x_v^{\min} and x_v^{\max} , respectively, are calculated from the observed data. The values assigned to the remaining dimensions are sampled independently using similar uniform distributions; however, the same value is shared across all samples (i.e., $\tilde{\mathbf{x}}_{1,k}^{(s)} = \tilde{\mathbf{x}}_{2,k}^{(s)} = \dots = \tilde{\mathbf{x}}_{n,k}^{(s)}, \forall k \in [1, \dots, t]$ and $k \neq v$).

Our approach requires processing data sets where the variables not under analysis are held constant. However, it is not always possible to find subsets of data that meet this condition for all variables in real-world datasets. Even when such subsets exist, they may not be sufficiently large. Thus, we generate sets of data $\tilde{\mathbf{X}}^{(s)}$ with the desired behavior and estimate their response using a regression model that has been trained on observed data. Then, our method derives univariate skeletons based on multiple sets of input-estimated response pairs. It is important to use a prediction model that learns a function \hat{f} that is as close as possible to f so that it accurately estimates how the real system would respond to the artificial inputs in $\tilde{\mathbf{X}}^{(s)}$. As a consequence, our analysis can be regarded as an explainability method that generates univariate symbolic skeletons as explanations of the function approximated by the regression model.

The response of the samples in $\tilde{\mathbf{X}}^{(s)}$, $\tilde{\mathbf{y}}^{(s)}$, is estimated using the network \hat{f} as $\tilde{\mathbf{y}}^{(s)} = \hat{f}(\tilde{\mathbf{X}}^{(s)})$. In addition, $\tilde{\mathbf{D}}_v^{(s)} = (\tilde{\mathbf{X}}_v^{(s)}, \tilde{\mathbf{y}}^{(s)})$ denotes the set of n input–response pairs and is used to analyze the relationship between the system’s

response and the v -th variable. The collection of N_S sets $\tilde{\mathbf{D}}_v = \{\tilde{\mathbf{D}}_v^{(1)}, \dots, \tilde{\mathbf{D}}_v^{(N_S)}\}$ is then fed into the pre-trained Multi-Set Transformer g so that the estimated skeleton obtained for variable x_v is calculated as $\tilde{\mathbf{e}}(x_v) = g(\tilde{\mathbf{D}}_v, \Theta)$. This process is repeated for all the variables of the system to obtain their corresponding symbolic skeleton expressions with respect to the system’s response.

3.3 Performance Evaluation

Typically, SR methods are evaluated on benchmark datasets such as the Feynman SR Benchmark (FSRB) [28]. Their performance is compared based on the mean squared error (MSE) achieved by the learned expressions on a subset of the available data. Note that a function that was learned by minimizing the prediction error with respect to the system’s response does not necessarily correspond to the underlying functional form; e.g., $f(x) = \cos(-\frac{x}{10})^2$ and $\hat{f}(x) = -0.0093x^2 + 0.9983$ produce an MSE lower than 4×10^{-5} when $x \in [-5, 5]$. As such, conventional evaluation approaches may not align with our objectives.

Instead of minimizing prediction error, our method produces a set of univariate skeletons aiming to describe the functional form between each variable and the system’s response. We argue that an expression with minimum prediction error can be produced as a consequence of identifying the correct functional form. To the best of our knowledge, no previous work has addressed the problem of testing how well the learned expression’s functional form (i.e., its skeleton) matches the system’s underlying functional form. Thus, we present a method to test the similarity between the underlying skeleton corresponding to the variable x_v , represented as $\mathbf{e}(x_v) = \kappa(f(\mathbf{x}), x_v)$, and the estimated skeleton $\hat{\mathbf{e}}(x_v)$. The function $\kappa(\cdot, x_v)$ replaces the numerical constants of a given symbolic expression with the placeholder labeled with the suffix c . In addition, it considers the remaining variables $\mathbf{x} \setminus x_v$ as constants as they are irrelevant when describing the functional form between x_v and the system’s response. For example, if $f(\mathbf{x}) = 3x_1^2 + \sqrt{x_2 + 1}/e^{2x_3}$, then $\kappa(f(\mathbf{x}), x_1) = c_1 x_1^2 + c_2$.

We assign random numerical values to the coefficients of skeleton $\mathbf{e}(x_v)$ using the `sampleConstants` routine (Algorithm 1) in order to obtain a function $f_{\text{target}}(x_v)$. Let $f_{\text{est}}(x_v) = \text{setConstants}(\hat{\mathbf{e}}(x_v), \mathbf{c})$ denote the function obtained when replacing the n_c constant placeholders in $\hat{\mathbf{e}}(x_v)$ with the numerical values in a given set $\mathbf{c} = [c_1, \dots, c_{n_c}]$. If the functional form of $\hat{\mathbf{e}}(x_v)$ is mathematically equivalent to that of $\mathbf{e}(x_v)$, then there exists a set of values \mathbf{c} so that the difference between $f_{\text{target}}(x_v)$ and $f_{\text{est}}(x_v)$ is zero for all values of x_v . The optimal set \mathbf{c}^* is found as: $\mathbf{c}^* = \text{argmin}_{\mathbf{c}} \sum_{x_v \in \mathbf{X}_v^{\text{test}}} |f_{\text{target}}(x_v) - f_{\text{est}}(x_v)|$, where $\mathbf{X}_v^{\text{test}}$ is a test set of N_{test} elements whose elements are drawn from a distribution $\mathcal{U}(2x_v^{\text{min}}, 2x_v^{\text{max}})$. Note that the domain of $\mathbf{X}_v^{\text{test}}$ is larger than that used for training (i.e., $[x_v^{\text{min}}, x_v^{\text{max}}]$). If the estimated skeleton matches the system’s underlying functional form, it should do so regardless of the variable domain.

The coefficient fitting problem is solved using a simple genetic algorithm (GA) [14]. The individuals of our GA are arrays of n_c elements that represent potential \mathbf{c} sets. The population size is set to 500 and the optimization algorithm is stopped when the change of the objective function after 20 generations

Table 1. Equations used for experiments

Eq.	Underlying equation	Reference	Domain range
E1	$(3.0375x_1x_2 + 5.5 \sin(9/4(x_1 - 2/3)(x_2 - 2/3)))/5$	[15]	$[-5, 5]^2$
E2	$5.5 + (1 - x_1/4)^2 + \sqrt{x_2} + 10 \sin(x_3/5)$	–	$[-10, 10]^2$
E3	$(1.5e^{1.5x_1} + 5 \cos(3x_2))/10$	[15]	$[-5, 5]^2$
E4	$((1 - x_1)^2 + (1 - x_3)^2 + 100(x_2 - x_1^2)^2 + 100(x_4 - x_3^2)^2)/10000$	Rosenbrock-4D	$[-5, 5]^4$
E5	$\sin(x_1 + x_2x_3) + \exp(1.2x_4)$	–	$x_1 \in [-10, 10], x_2 \in [-5, 5], x_3 \in [-5, 5], x_4 \in [-3, 3]$
E6	$\tanh(x_1/2) + x_2 \cos(x_3^2/5)$	–	$[-10, 10]^3$
E7	$(1 - x_2^2)/(\sin(2\pi x_1) + 1.5)$	[31]	$[-5, 5]^2$
E8	$x_1^4/(x_1^4 + 1) + x_2^4/(x_2^4 + 1)$	[27]	$[-25, 25]^2$
E9	$\log(2x_2 + 1) - \log(4x_1^2 + 1)$	[27]	$[0, 5]^2$
E10	$\sin(x_1 e^{x_2})$	[3]	$x_1 \in [-2, 2], x_2 \in [-4, 4]$
E11	$x_1 \log(x_2^4)$	[3]	$[-5, 5]^2$
E12	$1 + x_1 \sin(1/x_2)$	[3]	$[-10, 10]^2$
E13	$\sqrt{x_1} \log(x_2^2)$	[3]	$x_1 \in [0, 20], x_2 \in [-5, 5]$

is less than 10^{-5} . We utilize tournament selection, binomial crossover, and generational replacement. We selected this configuration as it demonstrated effective optimization results across all experiments utilized in this work.

If $\hat{\mathbf{e}}(x_v)$ and $\mathbf{e}(x_v)$ are not equivalent, the error $r = \sum_{x_v \in \mathbf{X}_v^{\text{test}}} |f_{\text{target}}(x_v) - \text{setConstants}(\hat{\mathbf{e}}(x_v), \mathbf{c}^*)|$ is greater than 0. We use r as a performance metric that indicates the closeness between $\hat{\mathbf{e}}(x_v)$ and $\mathbf{e}(x_v)$ given the sampled values of the constants of $\mathbf{e}(x_v)$. Note that if $\hat{\mathbf{e}}(x_v)$ and $\mathbf{e}(x_v)$ are similar, the error r should be low regardless of the sampled values of the constants of $\mathbf{e}(x_v)$. Thus, for the sake of generality, we repeat this process 30 times; that is, we sample 30 different $f_{\text{target}}(x_v)$ functions and solve 30 optimization problems. Finally, we report the mean and the standard deviation of the 30 resulting error metrics.

4 Experimental Results

A training dataset² consisting of one million pre-generated expressions ($|\mathbf{Q}| = 10^6$) has been created to train the Multi-Set Transformer. These expressions allow up to one nested operation and contain a maximum of five unary operators. We also generated an independent validation set consisting of 10^5 expressions. For the model architecture, due to the high computational expense associated with training a single model, we used a one-factor-at-a-time approach to choose the following hyperparameters: the number of ISAB encoder blocks $\ell = 3$, the number of decoder blocks $M = 5$, an embedding size $d = 512$, and the number of heads $h = 8$. In addition, we set the number of input sets to $N_S = 10$ and the number of input-response pairs in each input set to $n = 3000$. Future work will examine the impact of altering the values of N_S and n on the final performance.

² The code and datasets are available at <https://github.com/NISL-MSU/MultiSetSR>.

4.1 Experiments with Synthetic Data

We assessed the performance of the skeletons generated by our Multi-Set Transformer using ten synthetic SR problems generated by equations inspired by previous works and equations proposed in this work, as reported in Table 1. Note that previous works used narrow domain ranges for all variables (e.g., $[-1, 1]$) while we used extended ranges (e.g., $[-5, 5]$ and $[-10, 10]$) to increase the difficulty of the problems. We adapted the benchmark equations proposed by Bertschinger *et al.* [3] (i.e., E10–E13) to a multivariate setting.

In all cases, the training datasets consisted of 10,000 points where each variable was sampled using a uniform distribution. For the estimated response functions, \hat{f} , we trained feed-forward NNs with varying depths: three hidden layers for problem E2; five hidden layers for problems E1, E4, E5, and E7; and four hidden layers for the other cases. Each layer consisted of 500 nodes with ReLU activation. We used 90% of the samples for training and 10% for validation.

We compared the skeletons produced by the Multi-Set Transformer (MST) to the ones extracted from the expressions generated by three other methods: two GP-based methods (PySR [7] and TaylorGP [12]) and two neural SR methods (NeSymReS [4] and E2E [17]). For NeSymReS and E2E, we used the pre-trained models available online. NeSymReS, E2E, and the Multi-Set Transformer were trained using vocabularies with the same unary and binary operators: $+$, \times , $/$, abs , acos , asin , atan , cos , cosh , exp , log , pow2 , pow3 , pow4 , pow5 , sin , sinh , sqrt , tan , tanh . Thus, PySR and TaylorGP were executed using the same set of operators. Our experimentation with the GP-based methods involved a maximum of 10,000 iterations, though convergence was consistently achieved in fewer iterations across all cases. Population sizes of 100, 200, 500, and 1000 were tested, with no discernible advantage observed beyond a size of 500.

The compared methods produce multivariate expressions, from which the skeleton variable corresponding to variable x_v is obtained using the skeleton function $\kappa(\cdot, x_v)$. NeSymReS could not be executed on E4 and E5 because its model was trained using expressions limited to three variables. Table 2 shows the target and estimated skeletons corresponding to each variable for problem E2. The skeletons obtained for the other problems are presented in Section C of the **Supplementary Materials** (see footnote 1). We also evaluated skeleton performance using the method described in Sect. 3.3. We set the size of the test sets to $N_{\text{test}} = 3000$. Using $N_{\text{test}} > 3000$ did not vary the obtained results. Table 3 reports the rounded mean and the standard deviation of the error metrics obtained after 30 repetitions of the proposed evaluation. The bold entries indicate the method that achieved the lowest mean error r and that its difference w.r.t. the values obtained by the other methods is statistically significant according to Tukey’s honestly significant difference test performed at the 0.05 significance level.

Table 2. Comparison of skeleton prediction results for problem E2

Method	x_1	x_2	x_3
PySR	$c_1 + c_2 + c_3 + x_1 $	c_1	$c_1 + c_2 x_3$
TaylorGP	$c_1 + c_2 x_1$	$c_1 + c_2 x_2$	$c_1 + c_2 x_3$
NeSymReS	$c_1 + c_2 x_1$	$c_1 + \exp(\exp(c_2 x_2))$	$c_1 + c_2 x_3$
E2E	$c_1 + c_2 x_1 + c_3(c_4 + c_5 x_1)^2$	$c_1 + c_2(c_3 + c_4 x_2)$	$c_1 + c_2 x_3 + c_3(c_4 + c_5 \cos(c_6 + c_7 x_3))$
MST	$c_1 + c_2(c_3 + c_4 x_1)^2$	$c_1 \sqrt{c_2 x_2 + c_3} + c_4$	$c_1 + c_2 \sin(c_3 x_3 + c_4)$
Target $\mathbf{e}(x)$	$c_1 + (c_2 + c_3 x_1)^2$	$c_1 \sqrt{x_2 + c_2} + c_3$	$c_1 + c_2 \sin(c_3 x_3)$

4.2 Analysis of Nitrogen Fertilizer Response Curves

In the previous section, we performed experiments based on synthetic datasets. Recall that the objective of this work is to demonstrate that our method can explain the relationship between each system’s variable and the system’s response. To do so, datasets with known underlying functions are needed. However, in this section, we show the applicability of our approach in a real-world application.

In precision agriculture, crop yield prediction is a critical task with a significant societal impact [13]. Let $y(i, j)$ represent the observed yield at a field site with coordinates (i, j) . Furthermore, let $\mathbf{x}(i, j)$ represent a set of multiple covariate factors, such as the nitrogen (N) fertilizer rate applied, and other topographic and meteorological features. It describes the state of the field at position (i, j) , and potentially its neighboring areas. The underlying yield function of the field is denoted as $f(\cdot)$ and $y(i, j) = f(\mathbf{x}(i, j))$. Thus, the early crop yield prediction problem consists of generating predicted yield values $\hat{y}(i, j)$ using a prediction model $\hat{f}(\cdot)$, such that $\hat{y}(i, j) = \hat{f}(\mathbf{x}(i, j))$ and $\hat{y}(i, j) \approx y(i, j)$.

In practice, f is a complex multivariate system with unknown functional form. Nevertheless, tasks like N-rate optimization, which allows for profit maximization and environmental impact maximization [13], do not require estimating the full functional form of $f(\mathbf{x}(i, j))$. Instead, N-rate optimization only analyzes the functional relationship between the N-rate variable and the predicted yield values. This relationship is typically represented using N fertilizer-yield response (N-response) curves, which depict the estimated crop yield values for a specific field site in response to various admissible fertilizer rates.

The experiments presented in this section aim to estimate the functional form of N-response curves for a winter wheat dryland field called “Field A”, discussed in a previous work [23]. Traditionally, N-response curves are fitted employing a single parametric function for the entire field [5, 16]. However, previous work [23] demonstrates that the N-response functional form varies within the field, which can be clustered into regions known as management zones (MZs) based on the fertilizer responsiveness of the field sites (i.e., the shape of the approximated N-response curves). For details about the dataset, the generation of approximated N-response curves, and the clustering method, please see [23].

Table 3. Skeleton evaluation performance comparison

Eq.	Var.	PySR	TaylorGP	NeSymReS	E2E	MST
E1	x_1	1.4 ± 0.8	1.4 ± 0.8	0.9 ± 0.7	0.2 ± 0.4	0.01 ± 0.02
	x_2	1.5 ± 0.9	1.5 ± 0.9	1.3 ± 0.8	1.5 ± 0.9	0 ± 0
E2	x_1	303.5 ± 167.3	310.0 ± 170.1	310.0 ± 170.1	0 ± 0	0 ± 0
	x_2	5.4 ± 5.0	4.2 ± 5.3	4.6 ± 5.0	4.2 ± 5.3	0.02 ± 0.03
	x_3	1.7 ± 1.0	1.7 ± 1.0	1.7 ± 1.0	0.01 ± 0.02	0 ± 0
E3	x_1	$2 \times 10^{12} \pm 5 \times 10^{12}$	939.4 ± 1419.9	1.9 ± 1.2	0.8 ± 1.8	0.8 ± 1.8
	x_2	1.3 ± 1.0	1.3 ± 1.0	0.8 ± 0.8	0 ± 0	0 ± 0
E4	x_1	4576.2 ± 2695.7	4581.5 ± 2697.4	–	2.3 ± 3.6	1.1 ± 0.7
	x_2	79.6 ± 41.3	80.2 ± 40.8	–	0 ± 0	0 ± 0
	x_3	3995.5 ± 2815.6	4304.6 ± 2843.7	–	2.0 ± 4.0	1.0 ± 0.9
	x_4	74.5 ± 48.0	75.5 ± 47.0	–	0 ± 0	0 ± 0
E5	x_1	0.6 ± 0.05	0.6 ± 0.05	–	0 ± 0	0 ± 0
	x_2	1.5 ± 1.0	1.5 ± 1.0	–	1.5 ± 1.0	0 ± 0
	x_3	0.6 ± 0.05	0.6 ± 0.05	–	0.6 ± 0.05	0 ± 0
	x_4	2.7 ± 1.2	487.4 ± 461.9	–	0.6 ± 0.8	0.6 ± 0.8
E6	x_1	0.8 ± 0.08	0.8 ± 0.08	0.3 ± 0.01	0.04 ± 0	0 ± 0
	x_2	16.8 ± 12.2	16.8 ± 12.2	13.8 ± 11.1	1.3 ± 0.9	0 ± 0
	x_3	2.9 ± 1.3	1.9 ± 0.7	1.6 ± 0.9	1.6 ± 0.9	0 ± 0
E7	x_1	29.5 ± 1.0	1.8 ± 2.0	1.8 ± 2.0	1.1 ± 1.3	0 ± 0
	x_2	63.6 ± 43.8	1.6 ± 1.0	42.4 ± 24.7	0 ± 0	0 ± 0
E8	x_1	0.02 ± 0.01	0.04 ± 0.01	0.8 ± 1.2	0.02 ± 0.02	0 ± 0
	x_2	0.02 ± 0.01	0.04 ± 0.01	0.9 ± 1.3	0.02 ± 0.01	0 ± 0
E9	x_1	271.3 ± 446.8	239.8 ± 428.2	375.1 ± 485.5	0 ± 0	0 ± 0
	x_2	0 ± 0	0.2 ± 0.09	2.7 ± 1.7	0.05 ± 0.01	0 ± 0
E10	x_1	0 ± 0	0.6 ± 0.2	0 ± 0	0 ± 0	0 ± 0
	x_2	0 ± 0	0.4 ± 0.06	0 ± 0	0 ± 0	0 ± 0
E11	x_1	0 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0
	x_2	0 ± 0	0 ± 0	0 ± 0	0 ± 0	0 ± 0
E12	x_1	21.8 ± 13.1	0 ± 0	0 ± 0	0 ± 0	0 ± 0
	x_2	0 ± 0	2.4 ± 1.5	2.5 ± 1.6	0 ± 0	0 ± 0
E13	x_1	0 ± 0	0.8 ± 0.5	0.7 ± 0.6	0.7 ± 0.8	0 ± 0
	x_2	0 ± 0	0 ± 0	3.8 ± 3.5	0 ± 0	0 ± 0

Figure 2 depicts Field A clustered into four MZs. The figure also shows fifty example response curves selected randomly from each MZ. Due to their shape similarity, we assume that all sites within an MZ share the same functional form. Consider that an N-response curve is a set of input–response pairs (N-rate Nr vs.




Fig. 2. MZs for Field A and corresponding N-response curves [23].

relative yield ry). Thus, the problem of estimating the functional form of the N-response curves of an MZ can be posed as an MSSP. To do so, we select $N_S = 10$ random field sites within an MZ and feed their corresponding N-response curves into the pre-trained MST model to obtain a univariate skeleton.

Table 4. Comparison of skeleton prediction results for Field A

Method	MZ	Functional Form	\bar{r}	Method	MZ	Functional Form	\bar{r}
Quadratic-plateau	1	$c_1 + c_2 (\min(Nr, c_3) + c_3)^2$	0.0680	Quadratic-plateau	3	$c_1 + c_2 (\min(Nr, c_3) + c_3)^2$	0.0879
Exponential		$c_1(1 - \exp(c_2 + c_3 Nr)) + c_4$	0.2303	Exponential		$c_1(1 - \exp(c_2 + c_3 Nr)) + c_4$	0.1965
MST		$c_1 + c_2(Nr - \log(c_3 + c_4 Nr))$	0.0489	MST		$c_1 + c_2 Nr + c_3 \cos(c_4 + c_5 Nr)$	0.0479
Quadratic-plateau	2	$c_1 + c_2 (\min(Nr, c_3) + c_3)^2$	0.0725	Quadratic-plateau	4	$c_1 + c_2 (\min(Nr, c_3) + c_3)^2$	0.0615
Exponential		$c_1(1 - \exp(c_2 + c_3 Nr)) + c_4$	0.1825	Exponential		$c_1(1 - \exp(c_2 + c_3 Nr)) + c_4$	0.2249
MST		$c_1 + c_2 \tanh(c_3 + c_4 Nr)$	0.0495	MST		$c_1 + c_2 Nr + c_3 \cos(c_4 + c_5 Nr)$	0.0237

Table 4 shows the skeletons derived by the MST for each MZ. The process was repeated multiple times for each MZ, with different random sites selected as inputs for the MSSP problem in each iteration, resulting in the generation of equivalent skeletons. Furthermore, we evaluated the suitability of the obtained skeletons for each MZ and compared them to two traditional N-response models: quadratic-plateau [5] and exponential [16]. For each method and field site, we fit the skeleton's coefficient values to minimize the distance to the corresponding N-response curve, following a similar procedure as the one described in Sect. 3.3. We report the average error \bar{r} obtained considering all sites within each MZ. Results indicate that different MZs may be modeled using different functional forms, unlike traditional approaches that assume one for the entire field.

5 Discussion

Our method involves training a feedforward NN, which is then used to solve an MSSP problem for each system's variable using a Multi-Set Transformer. This process produces univariate skeletons that describe the functional relationship between each variable and the system's response. After evaluation of our univariate skeleton prediction method across the tested problems, we observed that it generated skeletons that matched or were equivalent to the target skeleton

for all variables across all problems. For instance, for problem E11, the target skeleton for variable x_2 is given by $\mathbf{e}(x_2) = c_1 \log(x_2^4)$, and our method produces the skeleton $\hat{\mathbf{e}}_{MST}(x_2) = c'_1 + c'_2 \log(c'_3 x_2^2)$. Notice that these skeletons are equivalent if $c'_1 = 0$, $c'_2 = 2c_1$, and $c'_3 = 1$. In addition, from the skeleton performance evaluation shown in Table 3, we verified that our method consistently attained lower or comparable error metrics compared to other SR methods. These results strongly support our hypothesis that our method would generate univariate skeletons that are more similar to those corresponding to the underlying equations in comparison to other SR methods.

Note that E2E produced the correct skeleton for at least one of the variables in most of the cases. Recall that E2E generated expressions that minimize the prediction error; thus, it did not prioritize identifying the correct functional form of the variables that do not contribute substantially to the overall error. In addition, in some cases, E2E generated skeletons that were equivalent to the target skeletons but larger than those produced by MST. For example, in problem E4, E2E generated the skeleton $\hat{\mathbf{e}}_{E2E}(x_1) = c_1 + c_2|c_3 + c_4 x_1 + c_5 x_1^2 + c_6 x_1^3 + c_7 x_1^4|$ for variable x_1 , which is equivalent to the one produced by MST, $\hat{\mathbf{e}}_{MST}(x_1) = c_1 + c_2 x_1 + c_3 x_1^2 + c_4 x_1^3 + c_5 x_1^4$. Another advantage over the other neural SR methods is that MST requires 24.2 million parameters while E2E requires 93.5 million. NeSymReS requires 26.4 million parameters (i.e., 2.2 million more than MST), and Table 3 shows that it failed to identify the correct functional form in most cases and is limited to problems with up to three variables.

It is worth pointing out that in problems E5, E6, E8, E9, and E13, some compared methods achieved low error metrics but are not comparable to the ones achieved by our method. For example, in problem E6, E2E generated the skeleton $\hat{\mathbf{e}}_{E2E}(x_1) = c_1 + c_2 \operatorname{atan}(c_3 + c_4 x_1)$ for variable x_1 , which does not coincide with the functional form of the underlying skeleton $\mathbf{e}(x_1) = c_1 + \tanh(c_2 x_1)$. However, E2E achieved low error metrics because, during the coefficient fitting process, the GA found appropriate values for the constant that multiplies the argument of the `atan` function, stretching or compressing the curve, making it resemble the shape of `tanh` and minimizing the error. Hence, the skeleton generated by E2E produced low error metrics and is considered to be similar to the target skeleton. Conversely, the functional form of the skeleton generated by MST coincided with that of the target skeleton exactly (i.e., $\hat{\mathbf{e}}_{MST}(x_1) = c_1 + c_2 \tanh(c_3 x_1)$) and thus produced significantly lower error metrics.

We claim that our method can be regarded as a *post-hoc* explainability method that generates univariate skeletons as explanations of the function approximated by a black-box regression model. From an explainability perspective, the generation of more accurate univariate symbolic skeletons is crucial. By producing skeletons that align more closely with the underlying data, we offer transparent insights into how each variable influences the function's response.

Finally, we applied our method to a precision agriculture problem. Given a field that has been split into MZs, the problem consists of determining the functional form that describes the N-response curves corresponding to all sites within an MZ. This is modeled as an MSSP problem and the skeleton generated

for each MZ is reported in Table 4. Our results show that the generated skeletons yield lower fitting errors and, thus are more suitable when modeling the field’s N-response curves. In future work, the function fitted for each site will serve as a surrogate yield model that may aid in optimization tasks.

One potential limitation of our approach, as well as any neural SR method, lies in its ability to generate skeletons whose complexity is bounded inherently by the expressions produced during the pre-training phase of the Multi-Set Transformer. For instance, we would not be able to identify the skeleton $c_1 + c_2 x_2^2 / \sin(c_3 e^{c_4 x_2})$ as it requires eight operators, while our training set was limited to expressions with up to seven operators. However, it is feasible to overcome this limitation through transfer learning, so that the MST model can be trained on more complex tasks, potentially enabling the recognition of such complex skeletons. Another limitation is that we do not provide multivariate expressions as other methods do. This is because the objective of this work lies in discerning the functional relationship between each variable and the system’s response accurately. Future work will focus on using the generated univariate skeletons as building blocks to produce multivariate expressions that approximate the system’s entire underlying function. For instance, mathematical expressions explaining the full behavior of real-world systems will be generated, and their performance will be evaluated by comparing observed response values with those predicted by the generated expressions.

6 Conclusions

Symbolic regression aims to find symbolic expressions that represent the relationships within the observed data. As such, SR represents a promising avenue for building explainable models. By seeking to uncover mathematical equations that represent the relationships between input variables and their response, the resulting equations offer transparency and clear insights into model behavior.

Given a multivariate regression problem, the objective of the work reported in this paper is to produce symbolic skeletons that describe the relationship between each variable and the system’s response. To do so, we introduced a problem called Multi-Set Skeleton Prediction that aims to generate a symbolic skeleton expression that characterizes the functional form of multiple sets of input–response pairs. To solve this problem, we proposed a novel Multi-Set Transformer model, which was pre-trained on a large dataset of synthetic symbolic expressions. Then, the multivariate SR problem is tackled as a sequence of MSSP problems generated using a black-box regression model. The generated skeletons are regarded as explanations of the black-box model’s function. Experimental results showed that our method consistently produced more accurate univariate skeletons in comparison to two GP-based SR methods and two neural SR methods.

Future work will focus on merging the generated univariate skeletons into a multivariate symbolic expression that approximates the underlying function of the system. This task requires us to answer questions such as how to ensure that the skeletons generated for all variables are compatible and can be merged, or how to find the optimal order in which to merge the skeletons.

Acknowledgments. This research was supported by the Data Intensive Farm Management project (USDA-NIFA-AFRI 2016-68004-24769 and USDA-NRCS NR213A7500013G021). Computational efforts were performed on the Tempest HPC System, operated by University Information Technology Research Cyberinfrastructure at MSU.

References

1. Amir Haeri, M., Ebadzadeh, M.M., Folino, G.: Statistical genetic programming for symbolic regression. *Appl. Soft Comput.* **60**, 447–469 (2017)
2. Arnaldo, I., Krawiec, K., O’Reilly, U.M.: Multiple regression genetic programming. In: 2014 Annual Conference on Genetic and Evolutionary Computation, pp. 879–886 (2014)
3. Bertschinger, A., Davis, Q.T., Bagrow, J., Bongard, J.: The metric is the message: benchmarking challenges for neural symbolic regression. In: Koutra, D., Plant, C., Gomez Rodriguez, M., Baralis, E., Bonchi, F. (eds.) *ECML PKDD 2023*. LNCS, vol. 14172, pp. 161–177. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-43421-1_10
4. Biggio, L., Bendinelli, T., Neitz, A., Lucchi, A., Parascandolo, G.: Neural symbolic regression that scales. In: 38th International Conference on Machine Learning, vol. 139, pp. 936–945 (2021)
5. Bullock, D.G., Bullock, D.S.: Quadratic and quadratic-plus-plateau models for predicting optimal nitrogen rate of corn: a comparison. *Agron. J.* **86**(1), 191–195 (1994)
6. Chu, X., Zhao, H., Xu, E., Qi, H., Chen, M., Shao, H.: Neural symbolic regression using control variables. *ArXiv abs/2306.04718* (2023)
7. Cranmer, M.: Interpretable machine learning for science with PySR and SymbolicRegression.jl. *ArXiv abs/2305.01582* (2023)
8. Dignum, S., Poli, R.: Generalisation of the limiting distribution of program sizes in tree-based genetic programming and analysis of its effects on bloat. In: 9th Annual Conference on Genetic and Evolutionary Computation, pp. 1588–1595 (2007)
9. Ebner, M.: On the search space of genetic programming and its relation to nature’s search space. In: 1999 Congress on Evolutionary Computation, vol. 2, pp. 1357–1361 (1999)
10. Filho, R., Lacerda, A., Pappa, G.: Explaining symbolic regression predictions. In: 2020 IEEE Congress on Evolutionary Computation, pp. 1–8 (2020)
11. Fletcher, R.: *Practical Methods of Optimization*, 2nd edn. Wiley, New York (1987)
12. He, B., Lu, Q., Yang, Q., Luo, J., Wang, Z.: Taylor genetic programming for symbolic regression. In: *Genetic and Evolutionary Computation Conference*, pp. 946–954 (2022)
13. Hegedus, P., et al.: Towards a low-cost comprehensive process for on-farm precision experimentation and analysis. *Agriculture* **13**(3) (2023)
14. Holland, J.H.: *Genetic algorithms*. *Sci. Am.* (1992)
15. Jin, Y., Fu, W., Kang, J., Guo, J., Guo, J.: Bayesian symbolic regression. *ArXiv abs/1910.08892* (2020)
16. Kakimoto, S., Mieno, T., Tanaka, T., Bullock, D.: Causal forest approach for site-specific input management via on-farm precision experimentation. *Comput. Electron. Agric.* **199**, 107164 (2022)

17. Kamienny, P.A., d'Ascoli, S., Lample, G., Charton, F.: End-to-end symbolic regression with transformers. In: *Advances in Neural Information Processing Systems*, vol. 35, pp. 10269–10281 (2022)
18. La Cava, W., et al.: Contemporary symbolic regression methods and their relative performance. In: *Advances in Neural Information Processing Systems* (2021)
19. Lample, G., Charton, F.: Deep learning for symbolic mathematics. In: *International Conference on Learning Representations* (2020)
20. Langdon, W.B., Soule, T., Poli, R., Foster, J.A.: *The evolution of size and shape*. In: *Advances in Genetic Programming*. The MIT Press (1999)
21. Lee, J., Lee, Y., Kim, J., Kosiorek, A., Choi, S., Teh, Y.W.: Set Transformer: a framework for attention-based permutation-invariant neural networks. In: *36th International Conference on Machine Learning*, vol. 97, pp. 3744–3753 (2019)
22. Martius, G., Lampert, C.H.: Extrapolation and learning equations. *ArXiv abs/1610.02995* (2016)
23. Morales, G., Sheppard, J.W.: Counterfactual analysis of neural networks used to create fertilizer management zones. In: *2024 International Joint Conference on Neural Networks*. Yokohama, Japan (2024)
24. Orzechowski, P., La Cava, W., Moore, J.: Where are we now? A large benchmark study of recent symbolic regression methods. In: *Genetic and Evolutionary Computation Conference*, pp. 1183–1190 (2018)
25. Petersen, B.K., Landajuela, M., Mundhenk, T.N., Santiago, C.P., Kim, S.K., Kim, J.T.: Deep symbolic regression: recovering mathematical expressions from data via risk-seeking policy gradients. In: *International Conference on Learning Representations* (2021)
26. Sahoo, S., Lampert, C., Martius, G.: Learning equations for extrapolation and control. In: *35th International Conference on Machine Learning*, pp. 4442–4450 (2018)
27. Trujillo, L., Muñoz, L., Galván-López, E., Silva, S.: neat genetic programming: controlling bloat naturally. *Inf. Sci.* **333**, 21–43 (2016)
28. Udrescu, S.M., Tegmark, M.: AI Feynman: a physics-inspired method for symbolic regression. *Sci. Adv.* **6**(16), eaay2631 (2020)
29. Valipour, M., You, B., Panju, M., Ghodsi, A.: SymbolicGPT: a generative transformer model for symbolic regression. In: *Advances in Neural Information Processing Systems* (2022)
30. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Advances in Neural Information Processing Systems*. vol. 30 (2017)
31. Werner, M., Junginger, A., Hennig, P., Martius, G.: Informed equation learning. *ArXiv abs/2105.06331* (2021)