Discovering Symbolic Differential Equations with Symmetry Invariants

Anonymous Author(s)
Affiliation
Address
email

Abstract

Discovering symbolic differential equations from data uncovers fundamental dynamical laws underlying complex systems. However, existing methods often struggle with the vast search space of equations and may produce equations that violate known physical laws. In this work, we address these problems by introducing the concept of *symmetry invariants* in equation discovery. We leverage the fact that differential equations admitting a symmetry group can be expressed in terms of differential invariants of symmetry transformations. Thus, we propose to use these invariants as atomic entities in equation discovery, ensuring the discovered equations satisfy the specified symmetry. Our approach integrates seamlessly with existing equation discovery methods such as sparse regression and genetic programming, improving their accuracy and efficiency. We validate the proposed method through applications to various physical systems, such as fluid and reaction-diffusion, demonstrating its ability to recover parsimonious and interpretable equations that respect the laws of physics.

1 Introduction

Differential equations describe relationships between functions representing physical quantities and their derivatives. They are crucial in modeling a wide range of phenomena, from fluid dynamics and electromagnetic fields to chemical reactions and biological processes, as they succinctly capture the underlying principles governing the behavior of complex systems. The discovery of governing equations in symbolic forms from observational data bridges the gap between raw data and fundamental understanding of physical systems. Unlike black-box machine learning models, symbolic equations provide interpretable insights into the structure and dynamics of the systems of interest. In this paper, we aim to discover symbolic partial differential equations (PDEs) in the form

$$F(\mathbf{x}, u^{(n)}) = 0, \tag{1}$$

where $\mathbf{x}$ denotes the independent variables, $u^{(n)}$ consists of the dependent variable u and all of its up-to-nth order partial derivatives.

While it has long been an exclusive task for human experts to identify governing equations, symbolic regression (SR) has emerged as an increasingly popular approach to automate the discovery.[1] SR constructs expressions from a predefined set of atomic entities, such as variables, constants, and mathematical operators, and fits the expressions to data by numerical optimization. Common methods include sparse regression (Brunton et al., 2016; Champion et al., 2019), genetic programming (Cranmer et al., 2019, 2020; Cranmer, 2023), neural networks (Kamienny et al., 2022), etc.

[1]While some literature uses *symbolic regression* specifically for GP-based methods, we use the term interchangeably with *equation discovery* to refer to all algorithms for learning symbolic equations.

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

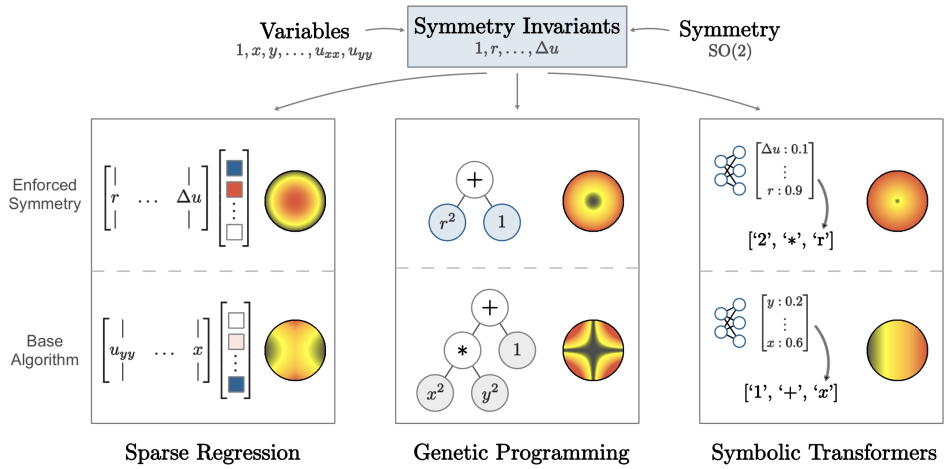

Figure 1: Our framework enforces symmetry in equation discovery by using symmetry invariants. We highlight three discovery algorithms in their original form (bottom row) and when constrained to only use symmetry invariants (top row). The colored circles visualize the predicted functions on a circular domain and demonstrate that using symmetry invariants guarantees a symmetric output.

However, symbolic regression algorithms may fail due to the vastness of the search space or produce more complex, less interpretable equations that overfit the data. A widely adopted remedy to these challenges is to incorporate inductive biases derived from physical laws, such as symmetry and conserved quantities, into equation discovery algorithms. Implementing these physical constraints narrows the space for equations and expedites the search process, and it also rules out physically invalid or unnecessarily complex equations.

Among the various physical constraints, symmetry plays a fundamental role in physical systems, governing their invariances under transformations such as rotations, translations, and scaling. Previous research has shown the benefit of incorporating symmetry in equation discovery, such as reducing the dimensionality of the search space and promoting parsimony in the discovered equation (Yang et al., 2024). However, the scopes of existing works exploiting symmetry are limited in terms of the types of equations they can handle, the compatible base algorithms, etc. For example, Udrescu & Tegmark (2020) deals with algebraic equations; Otto et al. (2023) deals with ODE systems; Yang et al. (2024) applies to sparse regression but not other SR algorithms.

In this paper, we propose a general procedure based on *symmetry invariants* to enforce the inductive bias of symmetry with minimal restrictions in the types of equations and SR algorithms. Specifically, we leverage the fact that a differential equation can be written in terms of the invariants of symmetry transformations if it admits a certain symmetry group. Thus, instead of operating on the original variables, our method uses the symmetry invariants as the atomic entities in symbolic regression, as depicted in Figure 1. These invariants encapsulate the essential information while automatically satisfying the symmetry constraints. Consequently, the discovered equations are guaranteed to preserve the specified symmetry. In summary, our main contributions are listed as follows:

- We propose a general framework to enforce symmetry in differential equation discovery based on the theory of differential invariants.
- Our approach can be easily integrated with existing symbolic regression methods, such as sparse regression and genetic programming, and improves their accuracy and efficiency for differential equation discovery.
- We show that our symmetry-based approach is robust in challenging setups in equation discovery, such as noisy data and imperfect symmetry.

Notations. Throughout the paper, subscripts are usually reserved for partial derivatives, e.g. $u_t := \partial u/\partial t$, and $u_{xx} := \partial^2 u/\partial x^2$. Superscripts are used for indexing vector components or list elements. We use Einstein notation, where repeated indices are summed over. Matrices, vectors and scalars are denoted by capital, bold and regular letters, respectively, e.g. $W, \mathbf{w}, w$. These conventions may admit exceptions for clarity or context. See Table 2 for a full description of notations.

2 Background

2.1 PDE Symmetry

This section introduces the basic concepts about partial differential equations and their symmetry. For a more thorough understanding of Lie point symmetry of PDEs, we refer the readers to Olver (1993).

Partial Differential Equations. We consider PDEs in the form $F(\mathbf{x}, u^{(n)}) = 0$, as given in (1). We restrict ourselves to a single equation and a single dependent variable, though generalization is possible. We use $\mathbf{x} \in X \subset \mathbb{R}^p$ to denote all independent variables. For example, $\mathbf{x} = (t, x)$ for a system evolving in 1D space. Note that the bold $\mathbf{x}$ refers to the collection of all independent variables while the regular x denotes the spatial variable. Then, $u = u(\mathbf{x}) \in U \subset \mathbb{R}$ is the dependent variable; $u^{(n)} = (u, u_x, ...)$ denotes all up to nth-order partial derivatives of u; $(\mathbf{x}, u^{(n)}) \in M^{(n)} \subset X \times U^{(n)}$, where $M^{(n)}$ is the nth order **jet space** of the total space $X \times U$. $M^{(n)}$ and $u^{(n)}$ are also known as the nth-order **prolongation** of $X \times U$ and u, respectively.

Symmetry of a PDE. A point symmetry g is a local diffeomorphism on the total space $E = X \times U$:

$$g \cdot (\mathbf{x}, u) = (\tilde{\mathbf{x}}(\mathbf{x}, u), \tilde{u}(\mathbf{x}, u)), \tag{2}$$

where $\tilde{\mathbf{x}}$ and $\tilde{u}$ are functions of E. The action of g on the function $u(\mathbf{x})$ is induced from (2) by applying it to the graph of $u : X \to U$. Specifically, denote the domain of u as $\Omega \subset X$ and its graph as $\Gamma_u = \{(\mathbf{x}, u(\mathbf{x})) : \mathbf{x} \in \Omega\}$. The group element g transforms the graph Γ_u as $\tilde{\Gamma}_u \coloneqq g \cdot \Gamma_u = \{(\tilde{\mathbf{x}}, \tilde{u}) = g \cdot (\mathbf{x}, u) : (\mathbf{x}, u) \in \Gamma_u\}$.

Since g transforms both independent and dependent variables, $\tilde{\Gamma}_u$ does not necessarily correspond to the graph of any single-valued function. Nevertheless, by suitably shrinking the domain Ω_X, we can ensure that the transformations close to the identity transform Γ_u to the graph of another function. This function with the transformed graph $\tilde{\Gamma}_u$ is then defined to be the transformed function of the original solution u, i.e. $g \cdot u = \tilde{u}$ s.t. $\Gamma_{\tilde{u}} = \tilde{\Gamma}_u$. The symmetry of the PDE (1) is then defined:

Definition 2.1. A symmetry group of $F(\mathbf{x}, u^{(n)}) = 0$ is a local group of transformations G acting on an open subset of the total space $X \times U$ such that, for any solution u to $F = 0$ and any $g \in G$, the function $\tilde{u} = (g \cdot u)(\mathbf{x})$ is also a solution of $F = 0$ wherever it is defined.

Infinitesimal Generators. Often, the symmetry group of a PDE is a continuous Lie group. In practice, one needs to compute with infinitesimal generators of continuous symmetries, i.e., vector fields. In more detail, we will write vector fields $\mathbf{v} : E \to TE$ on $E = X \times U$ as

$$\mathbf{v} = \xi^j(\mathbf{x}, u)\frac{\partial}{\partial x^j} + \phi(\mathbf{x}, u)\frac{\partial}{\partial u}. \tag{3}$$

Any such vector field generates a one-parameter group of symmetries of the total space $\{\exp(\epsilon\mathbf{v}) : \epsilon \in \mathbb{R}\}$. The symmetries arising from the exponentiation of a vector field moves a point in the total space along the directions given by the vector field. We will specify symmetries by vector fields in the following sections. For instance, $\mathbf{v} = x\partial_y - y\partial_x$ represents the rotation in (x, y)-plane; $\mathbf{v} = \partial_t$ corresponds to time translation.

Since we deal with PDEs, we need to consider the **prolonged** group actions and infinitesimal actions on the nth-order jet space, induced from (2) and (3). These are denoted $g^{(n)}$ and $\mathbf{v}^{(n)}$, respectively. A more detailed discussion on prolongation of group actions is deferred to Appendix B.2. To introduce our method, it suffices to note that the prolongation of the vector field (3) can be described explicitly by ξ^j and ϕ and their derivatives via the *prolongation formula* (10).

2.2 Symbolic Regression Algorithms

Given the data $\{(x^i, y^i)\} \subset X \times Y$, the objective of symbolic regression is to find a symbolic expression for the function $y = f(x)$. Although this original formulation is for algebraic equations, it can be generalized to differential equations like (1). To discover a PDE from the dataset of its observed solutions on a grid Ω, i.e., $\{(\mathbf{x}, u(\mathbf{x})) : \mathbf{x} \in \Omega\}$, we estimate the partial derivative terms and add them to the dataset: $\{(\mathbf{x}, u^{(n)}) : \mathbf{x} \in \Omega\}$. One of the variables in the variable set $(\mathbf{x}, u^{(n)})$

is used as the LHS of the equation, i.e., the role of the label y in symbolic regression, while other variables serve as features. The precise set of derivatives added to symbolic regression and the choice of the equation LHS requires prior knowledge or speculations about the underlying system.

We briefly review two classes of symbolic regression algorithms: sparse regression (SINDy) and genetic programming (GP). A more detailed discussion of related works is found in Appendix A.

Sparse regression (Brunton et al., 2016) is specifically designed for discovering differential equations. It assumes the LHS ℓ of the equation is a fixed term, e.g. $\ell = u_t$, and the RHS of the equation can be written as a *linear combination* of m predefined functions θ^j with trainable coefficients $\mathbf{w} \in \mathbb{R}^m$, i.e.,

$$\ell(\mathbf{x}, u^{(n)}) = w^j \theta^j(\mathbf{x}, u^{(n)}),\ \theta^j : M^{(n)} \to \mathbb{R}. \tag{4}$$

The equation is found by solving for $\mathbf{w}$ that minimizes the objective $\|\mathbf{L} - \mathbf{R}\|_2^2 + \lambda\|\mathbf{w}\|_0$, where $\mathbf{L}$ and $\mathbf{R}$ are obtained by evaluating ℓ and $w^j\theta^j$ on all data points and concatenating them into column vectors, and $\|\mathbf{w}\|_0$ regularizes the number of nonzero terms. This formulation can be easily extended to q equations and dependent variables ($q > 1$): $\ell^i(\mathbf{x}, \mathbf{u}^{(n)}) = W^{ij}\theta^j(\mathbf{x}, \mathbf{u}^{(n)})$, $W \in \mathbb{R}^{q\times m}$.

One problem with sparse regression is that its assumptions about the form of the equation are restrictive. Many equations cannot be expressed in the form of (4), e.g. $y = \frac{1}{x+a}$ where a could be any constant. Also, the success of sparse regression relies on the proper choice of the predefined function library $\{\theta^j\}$. If any term in the true equation were not included, sparse regression would fail to identify the correct equation.

Genetic programming offers an alternative solution for equation discovery (Cranmer, 2023), which is capable of learning equations in more general forms. It represents each expression as a tree and instantiates a population of individual expressions. At each iteration, it randomly samples a subset of expressions and selects one of the expressions that best fits the data; the selected expression is then mutated by a random mutation, a crossover with another expression, or a constant optimization; the mutated expression replaces a weaker expression in the population that does not fit the data well. The algorithm repeats this process to search for different combinations of variables, constants, and operators and finally returns the "fittest" expression. Genetic programming can be less efficient than sparse regression when the equation can be expressed in the form (4) due to its larger search space. However, we will show that it is a promising alternative to discover PDEs of generic forms, and our approach further boosts its efficiency.

3 Symbolic Regression with Symmetry Invariants

Symmetry offers a natural inductive bias for the search space of symbolic regression in differential equations. It reduces the dimensionality of the space and encourages parsimony of the resulting equations. To enforce symmetry in PDE discovery, we aim to find the maximal set of equations admitting a *given* symmetry and search in that set with symbolic regression methods.

3.1 Differential Invariants and Symmetry Conditions

To achieve this, our general strategy is to replace the original variable set with a complete set of *invariant functions* of the given symmetry group. Since we consider PDEs containing partial derivatives, the invariant functions refer to the differential invariants defined as follows.

Definition 3.1 (Def 2.51, Olver (1993))**.** Let G be a local group of transformations acting on $X \times U$. Any $g \in G$ gives a prolonged group action $\mathrm{pr}^{(n)}g$ on the jet space $M^{(n)} \subset X \times U^{(n)}$. An nth order differential invariant of G is a smooth function $\eta : M^{(n)} \to \mathbb{R}$, such that for all $g \in G$ and all $(\mathbf{x}, u^{(n)}) \in M^{(n)}$, $\eta(g^{(n)} \cdot (\mathbf{x}, u^{(n)})) = \eta(\mathbf{x}, u^{(n)})$ whenever $g^{(n)} \cdot (\mathbf{x}, u^{(n)})$ is defined.

In other words, differential invariants are functions of all variables and partial derivatives that remain invariant under prolonged group actions. Equivalently, if G is generated by a set of infinitesimal generators $\mathcal{B} = \{\mathbf{v}_a\}$, then a function η is a differential invariant of G iff $\mathbf{v}_a^{(n)}(\eta) = 0$ for all $\mathbf{v}_a \in \mathcal{B}$.

The following theorem guarantees that any differential equation admitting a symmetry group can be expressed solely in terms of the group invariants.

Theorem 3.2 (Prop 2.56, Olver (1993))**.** *Let G be a local group of transformations acting on $X \times U$. Let $\{\eta^1(\mathbf{x}, u^{(n)}), ..., \eta^k(\mathbf{x}, u^{(n)})\}$ be a complete set of functionally independent nth-order differential*

invariants of G. An nth-order differential equation (1) *admits G as a symmetry group if and only if it is equivalent to an equation of the form*

$$\tilde{F}(\eta^1, ..., \eta^k) = 0. \tag{5}$$

Consequently, symbolic regression with a complete set of invariants precisely searches within the space of all symmetric differential equations, while automatically excluding equations violating the specified symmetry.

Our strategy of using differential invariants applies broadly to various equation discovery algorithms. For instance, in sparse regression, we can construct the function library using invariants rather than raw variables and derivatives. Similarly, in genetic programming, the variable set can be redefined to include only invariant functions. In each case, the key benefit is the same: the search space is restricted to symmetry-respecting equations by construction. The reduced complexity of the equation search also leads to increased accuracy and efficiency.

Next, we describe how to construct a complete set of differential invariants (Section 3.2), and how to incorporate them into specific SR algorithms (Section 3.3).

3.2 Constructing a Complete Set of Invariants

Despite the simplicity of our strategy, we still need a concrete method for computing the invariants. In this subsection, we provide a general guideline to construct a complete set of differential invariants up to a required order given the group action.

By definition of differential invariants, we look for functions $\eta(\mathbf{x}, u^{(n)})$ satisfying $\mathbf{v}^{(n)}(\eta) = 0$ given a prolonged vector field $\mathbf{v}^{(n)}$. This is a first-order linear PDE that can be solved by the method of characteristics. However, in practice, if $E = X \times U \simeq \mathbb{R}^p \times \mathbb{R}$, there are $\binom{p+n-1}{n}$ partial derivatives of the independent variable u of order exactly n. Therefore, as n grows, it quickly becomes impractical to solve directly for nth-order differential invariants. The computation of higher-order differential invariants, if necessary, is made tractable by the following result, where higher-order invariants are computed recursively from lower-order ones.

Proposition 3.3. *Let G be a local group of transformations acting on $X \times U \simeq \mathbb{R}^p \times \mathbb{R}$. Let $\eta^1, \eta^2, \cdots, \eta^p$ be any p differential invariants of G whose horizontal Jacobian $J = [D_i \eta^j]$ is non-degenerate on an open subset $\Omega \subset M^{(n)}$. If there are a maximal number of independent, strictly nth-order differential invariants $\zeta^1, \cdots, \zeta^{q_n}$, $q_n = \binom{p+n-1}{n}$, then the following set contains a complete set of independent, strictly $(n+1)$th-order differential invariants defined on Ω:*

$$\frac{\det(D_i \tilde{\eta}^j_{(k,k')})}{\det(D_i \eta^j)}, \ \forall k \in [p], k' \in [q_n], \tag{6}$$

where $i, j \in [p]$ are matrix indices, D_i denotes the total derivative w.r.t i-th independent variable and $\tilde{\eta}^j_{(k,k')} = [\eta^1, ..., \eta^{k-1}, \zeta^{k'}, \eta^{k+1}, ..., \eta^p]$.

In practice, we first solve for $\mathrm{pr}\ \mathbf{v}(\eta) = 0$ to obtain a sufficient number of lower-order invariants as required in Proposition 3.3, starting from which we can construct complete sets of invariants of arbitrary orders. Direct results from applying (6) can be complicated, especially for higher-order invariants. Thus, we often combine them to get simpler invariants, which we later use as the variable set in equation discovery. In Appendix B.4, we provide two examples of different symmetry groups and their differential invariants. Those results will also be used in our experiments.

3.3 Implementation in SR Algorithms

Our symmetry principle characterizes a subspace of all equations with a given symmetry. Generally, this subspace partially overlaps with the hypothesis spaces of symbolic regression algorithms, conceptually visualized in Figure 2. As in Theorem 3.2, PDEs with symmetry can be expressed as *implicit* functions of all differential invariants. However, symbolic regression methods typically learn *explicit* functions mapping features to labels. Some algorithms, such as SINDy, impose even stronger constraints on equation forms. Therefore, adaptation is needed to implement our strategy of using differential invariants in specific symbolic regression algorithms, as detailed below.

General explicit SR We start with general SR methods that learn an *explicit* function $y = f(x)$ without additional assumptions about the form of f, e.g., genetic programming and symbolic transformer. When learning the equation in terms of differential invariants, we do not know which one of them should be used as the LHS of the equation, i.e., the label y in symbolic regression. Thus, we fit an equation for each invariant as LHS and choose the equation with the lowest data error, as described in Algorithm 1. We use the relative error to select the best equation because the scales of LHS terms differ.

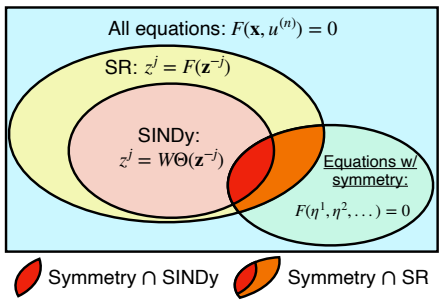

Figure 2: Venn diagram of hypothesis spaces from base SR methods and our symmetry principle.

Algorithm 1 General explicit SR for differential equations with symmetry invariants

Require: PDE order n, dataset $\{\mathbf{z}^i = (\mathbf{x}^i, (u^{(n)})^i) \in M^{(n)}\}_{i=1}^{N_D}$, base SR algorithm $\mathcal{S} : (\mathbf{X}, \mathbf{y}) \mapsto y = f(x)$, infinitesimal generators of the symmetry group $\mathcal{B} = \{\mathbf{v}_a\}$.
Ensure: A PDE admitting the given symmetry group.
Compute the symmetry invariants of $\mathcal{B}$ up to nth-order: $\eta^1, \cdots, \eta^K$. {Proposition 3.3}
Evaluate the invariant functions on the dataset: $\eta^{k,i} = \eta^k(\mathbf{z}^i)$, for $k \in [K], i \in [N_D]$.
Initialize a list of candidate equations and their risks: $\mathtt{E} = []$.
for k in $1 : K$ **do**
  Use the kth invariant as label and the rest as features: $\mathbf{y} = \eta^{k,:}$, $\mathbf{X} = \eta^{-k,:}$.
  Run $\mathcal{S}(\mathbf{X}, \mathbf{y})$ and get a candidate equation $\eta^k = f^k(\boldsymbol{\eta}^{-k})$.
  Evaluate $\mathcal{L}^k = \|\mathbf{y} - f^k(\mathbf{X})\|_1 / \|\mathbf{y}\|_1$ and set $\mathtt{E}[k] = (f^k, \mathcal{L}^k)$.
end for
Choose the equation in $\mathtt{E}$ with the lowest error: $k = \arg\min_j \mathtt{E}[j][2]$.
return $\eta^k = f^k(\boldsymbol{\eta}^{-k})$. {Optionally, expand all η^j in terms of original variables $\mathbf{z}$.}

Sparse regression SINDy assumes a linear equation form (4). Generally, its function library differs from the set of differential invariants. Also, SINDy fixes a LHS term, while we do not single out an invariant as the LHS of the equation when constructing the set of invariants.

Assume we are provided the SINDy configuration, i.e., the LHS term ℓ and the function library $\{\theta^j\}$. To implement sparse regression with symmetry invariants, we assign an invariant η^k that symbolically depends on ℓ, i.e., $\partial \eta^k / \partial \ell \neq 0$, as the LHS for the equation in terms of symmetry invariants. For example, if $\ell = u_{tt}$ and the set of invariants is given by (32), we use $\eta_{(0,2)} = u_{tt} u_x^{(b-2)/(a-b)}$ as the LHS since it is the only invariant that involves u_{tt}. The remaining invariants are included on the RHS, where they serve as inputs of the original SINDy library functions. In other words, the equation form is $\eta^k = \tilde{w}^j \theta^j(\boldsymbol{\eta}^{-k})$. Similar to Algorithm 1, we can expand all η variables to obtain the equation in original jet variables.

The above approach optimizes an unconstrained coefficient vector $\tilde{\mathbf{w}}$ for functions of symmetry invariants. Alternatively, we can use the original SINDy equation form (4) and implement the symmetry constraint as a constraint on the coefficient $\mathbf{w}$, as demonstrated in the following theorem. Here, we generalize the setup to multiple dependent variables and equations.

Proposition 3.4. *Let $\boldsymbol{\ell}(\mathbf{x}, \mathbf{u}^{(n)}) = W\boldsymbol{\theta}(\mathbf{x}, \mathbf{u}^{(n)})$ be a system of q differential equations admitting a symmetry group G, where $\mathbf{x} \in \mathbb{R}^p$, $\mathbf{u} \in \mathbb{R}^q$, $\boldsymbol{\theta} \in \mathbb{R}^m$. Assume that there exist some nth-order invariants of G, $\eta_0^{1:q}$ and $\eta^{1:K}$, s.t. (1) the system of differential equations can be expressed as $\boldsymbol{\eta}_0 = W'\boldsymbol{\theta}'(\boldsymbol{\eta})$, where $\boldsymbol{\eta}_0 = [\eta_0^{1:q}]$ and $\boldsymbol{\eta} = [\eta^{1:K}]$, and (2) $\eta_0^i = T^{ijk}\theta^k \ell^j$ and $(\theta')^i = S^{ij}\theta^j$, for some library functions $\boldsymbol{\theta}'(\boldsymbol{\eta})$ and some constant tensors W', T and S. Then, the space of all possible W is a linear subspace of $\mathbb{R}^{q \times m}$.*

Intuitively, the conditions in Proposition 3.4 state that the equations can be expressed as a linear combination of invariant terms, similar to the form in (4) w.r.t original jet variables. Also, every invariant term in $\boldsymbol{\eta}_0$ and $\boldsymbol{\theta}'(\boldsymbol{\eta})$ is already encoded in the original library $\boldsymbol{\theta}$. In practice, we need to choose a suitable set of invariants according to the SINDy configuration to meet these conditions. For example, when $\boldsymbol{\theta}$ contains all monomials on $M^{(n)}$ up to some degree, we can choose a set of

invariants where each invariant is a polynomial on $M^{(n)}$. The proof of Proposition 3.4 is deferred to Appendix B, where we explicitly identify the linear subspace for W entailed by the proposition.

Proposition 3.4 allows us to keep track of the original SINDy parameters W during optimization. This enables straightforward integration of symmetry constraints to variants of SINDy, e.g. Weak SINDy (Messenger & Bortz, 2021a,b) for noisy data. For example, if the constrained subspace has a basis $Q \in \mathbb{R}^{r\times q\times m}$, where r is the subspace dimension, we write $W^{jk} = Q^{ijk}\beta^i$. While we directly optimize β, we can still easily compute the objective of Weak SINDy which explicitly depends on W. In comparison, if we use the raw invariant terms for regression, e.g. the equations take the form $\boldsymbol{\eta}_0 = W'\boldsymbol{\theta}'(\boldsymbol{\eta})$, it is challenging to formulate the objective of Weak SINDy w.r.t W'.

More implementation details related to Section 3.3 can be found in Appendix C.

3.4 Constraint Relaxation for Systems with Imperfect Symmetry

Our approach discovers PDEs assuming perfect symmetry. However, it is common in reality that a system exhibits imperfect symmetry due to external forces, boundary conditions, etc. (Wang et al., 2022). In these cases, the previously mentioned method would fail to identify any symmetry-breaking factors. To address this, we propose to relax the symmetry constraints by allowing symmetry-breaking terms to appear in the equation, but at a higher cost.

We implement this idea in sparse regression, where the equation has a linear structure $\boldsymbol{\ell} = W\boldsymbol{\theta}$. We adopt the technique from Residual Pathway Prior (RPP) (Finzi et al., 2021), which is originally developed for equivariant linear layers in neural nets. Specifically, let Q be the basis of the parameter subspace that preserves symmetry and P be the orthogonal complement of Q. Instead of parameterizing W in this subspace, we define $W = A + B$ where $A^{jk} = Q^{ijk}\beta^i$ and $B^{jk} = P^{ijk}\gamma^i$ and place a stronger regularization on γ than on β. While the model still favors equations in the symmetry subspace spanned by Q, symmetry-breaking components in P can appear if it fits the data well.

4 Experiments

4.1 Datasets and Their Symmetries

We consider the following PDE systems, which cover different challenges in PDE discovery, such as high-order derivatives, generic equation form, multiple dependent variables and equations, noisy dataset, and imperfect symmetry. The datasets are generated by simulating the ground truth equation from specified initial conditions, with detailed procedures described in Appendix E.1.

Water Wave. Consider the Boussinesq equation describing the unidirectional propagation of a solitary wave in shallow water (Newell, 1985):

$$u_{tt} + uu_{xx} + u_x^2 + u_{xxxx} = 0 \tag{7}$$

This equation has a scaling symmetry $\mathbf{v}_1 = 2t\partial_t + x\partial_x - 2u\partial_u$ and the translation symmetries in space and time. As shown in Appendix B.4, the differential invariants are given by $\eta_{(\alpha,\beta)} = u_{x^{(\alpha)}t^{(\beta)}}u_x^{-(2+\alpha+2\beta)/3}$ where α and β are the orders of partial derivatives in x and t, respectively. To discover the 4th-order equation, we compute all $\eta_{(\alpha,\beta)}$ for $0 \le \alpha+\beta \le 4$, except for $\eta_{(1,0)} = 1$.

Darcy Flow. The following PDE describes the steady state of a 2D Darcy flow (Takamoto et al., 2022) with spatially varying viscosity $a(x, y) = e^{-4(x^2+y^2)}$ and a constant force term $f(x) = 1$:

$$-\nabla(e^{-4(x^2+y^2)}\nabla u) = 1 \tag{8}$$

This equation admits an $\mathrm{SO}(2)$ rotation symmetry $\mathbf{v} = y\partial_x - x\partial_y$. A detailed calculation of the differential invariants of this group can be found in Example B.5. In our experiment, we use the following complete set of 2nd-order invariants: $\{\frac{1}{2}(x^2+y^2), u, xu_y - yu_x, xu_x + yu_y, u_{xx} + u_{yy}, u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2, x^2u_{xx} + y^2u_{yy} + 2xyu_{xy}\}$.

Reaction-Diffusion. We consider the following system of PDEs from Champion et al. (2019):

$$\begin{aligned} u_t &= d_1\nabla^2 u + (1-u^2-v^2)u + (u^2+v^2)v \\ v_t &= d_2\nabla^2 v - (u^2+v^2)u + (1-u^2-v^2)v \end{aligned} \tag{9}$$

In the default setup, we use $d_1 = d_2 = 0.1$. The system then exhibits rotational symmetry in the phase space: $\mathbf{v} = u\partial_v - v\partial_u$. The ordinary invariants are $\{t, x, y, u^2+v^2\}$. The higher-order invariants are $\{\mathbf{u}\cdot\mathbf{u}_\mu, \mathbf{u}^\perp\cdot\mathbf{u}_\mu\}$, where $\mathbf{u} = (u, v)^T$ and μ is any multi-index of t, x and y.

We also consider the following cases where the rotation symmetry is broken due to different factors:

- **Unequal diffusivities** We use different diffusion coefficients for the two components: $d_1 = 0.1$, $d_2 = 0.1 + \epsilon$. This can happen, for example, when two chemical species described by the equation diffuse at different rates due to molecular size, charge, or solvent interactions.
- **External forcing** The ground truth equation (9)is modified by adding $-\epsilon v$ to the RHS of u_t and $-\epsilon u$ to the RHS of v_t. This can reflect a weak parametric forcing on the system.

4.2 Methods and Evaluation Criteria

We consider three classes of algorithms for equation discovery: sparse regression (PySINDy, de Silva et al. (2020); Kaptanoglu et al. (2022)), genetic programming (PySR, Cranmer (2023)), and a pretrained symbolic transformer (E2E, Kamienny et al. (2022)). For each class, we compare the original algorithm using the regular jet space variables (i.e., $(\mathbf{x}, u^{(n)})$) and our method using symmetry invariants. Our method will be referenced as SI (**S**ymmetry **I**nvariants) in the results.

To evaluate an equation discovery algorithm, we run it 100 times with randomly sampled data subsets and randomly initialized models if applicable. We record its *success probability* (SP) of discovering the correct equation. Specifically, we expand the ground truth equation into $\sum_i c^i f^i(\mathbf{z}) = 0$, where c^i are nonzero coefficients, $\mathbf{z}$ denotes the variables involved in the algorithm, i.e., original jet variables $(\mathbf{x}, u^{(n)})$ for baselines and symmetry invariants for our method, and f^i are functions of $\mathbf{z}$. Also, the discovered equation is expanded as $\sum_i \hat{c}^i \hat{f}^i(\mathbf{z}) = 0$, where $\hat{c}^i \neq 0$. The discovered equation is considered correct if all the terms with nonzero coefficients match the ground truth, i.e., $\{f^i\} = \{\hat{f}^i\}$. We also report the *prediction error* (PE), which measures how well the discovered equation fits the data. For evolution equations, we simulate each discovered equation from an initial condition and measure its difference from the ground truth solution at a specific timestep in terms of root mean square error (RMSE). Otherwise, we just report the RMSE of the discovered equation evaluated on all test data points.

4.3 Results on Clean Data with Perfect Symmetry

Table 1: Equation discovery results on clean data. $\mathcal{C}$, standing for *complexity*, refers to the effective parameter space dimension in sparse regression and the number of variables in GP/Transformer. SP and PE stands for *success probability* and *prediction error*, as explained in Section 4.2. The entries "-" suggest that the method does not apply to the specific PDE system, or the result is not meaningful.

Method		Boussinesq (7)			Darcy flow (8)			Reaction-diffusion (9)		
		$\mathcal{C}\downarrow$	SP $\uparrow$	PE $\downarrow$	$\mathcal{C}\downarrow$	SP $\uparrow$	PE $\downarrow$	$\mathcal{C}\downarrow$	SP $\uparrow$	PE $\downarrow$
Sparse Regression	PySINDy	15	0.00	0.373	-	-	-	38	0.53	0.021
	SI	**13**	**1.00**	**0.098**	-	-	-	28	**0.54**	**0.008**
Genetic Programming	PySR	17	0.90	0.098	8	0.00	0.187	17	0.00	-
	SI	**14**	**1.00**	0.098	**7**	**0.79**	**0.051**	**16**	**0.81**	**0.023**
Transformer	E2E	10	0.53	0.132	8	0.00	-	17	0.00	-
	SI	**7**	**0.85**	**0.104**	**7**	0.00	-	**16**	0.00	-

Table 1 summarizes the performance of all methods on the three PDE systems. For prediction errors (PE), we report the median, instead of the average, of 100 runs for each algorithm, because some incorrectly discovered equations yield tremendous prediction errors. Comparisons are made within each class of methods. Generally, using symmetry invariants reduces the complexity of equation discovery and improves the chance of finding the correct equations compared to the baselines.

Specifically, in sparse regression, our method using symmetry invariants is only slightly better than PySINDy in the reaction-diffusion system, but constantly succeeds in the Boussinesq equation where PySINDy fails. The failure of PySINDy is because the u_x^2 term in (7) is not supported by its function library, showing that SINDy's success relies heavily on the choice of function library. On the other hand, by enforcing the equation to be expressed in invariants, our method automatically identifies the proper function library. Appendix D.1 provides results for other variants of sparse regression.

For GP-based methods, Table 1 displays the results with a fixed number of GP iterations for each dataset. We also include results with different numbers of iterations in Appendix D.2. Generally, GP with invariants can identify the correct equation with fewer iterations and is considered more efficient.

4.4 Results on Noisy Data and Imperfect Symmetry

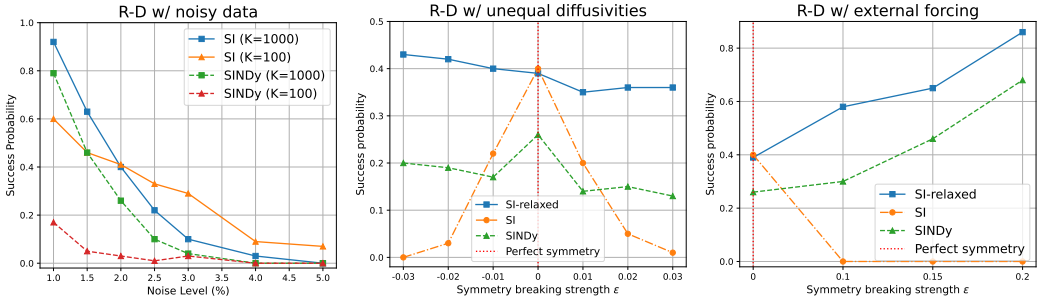

Figure 3: Success probabilities of sparse regression methods on the reaction-diffusion system with noisy data (left), unequal diffusivities (center) and external forcing (right). Under noisy data, our method (SI) consistently outperforms SINDy under the same number of test functions. For systems with imperfect symmetry, strictly enforcing symmetry (SI) can hurt performance, but a relaxed symmetry constraint (SI-relaxed) is still better than no inductive bias (SINDy).

We test the robustness of our method under two challenging scenarios: (1) noise in observed data, and (2) PDE with imperfect symmetry.

In the first experiment, we add different levels of white noise to the simulated solution of the reaction-diffusion system. Since the derivatives estimated by finite difference is inaccurate with the noisy solution, we use the weak formulation of SINDy (Messenger & Bortz, 2021a), which does not require derivative estimation. The success probabilities of our method (SI) and SINDy are shown in Figure 3 (left), where K is the number of test functions in weak SINDy. With the same K, our method consistently achieves higher success probability at different noise levels. Notably, when the noise level is high, our symmetry-constrained model performs better with fewer test functions ($K = 100$).

In the second experiment, we simulate the two variants of (9) (unequal diffusivities and external forcing) with different values for the symmetry-breaking parameter ϵ and add 2% noise to the numerical solutions. We compare three models: (1) our model with strictly enforced symmetry (SI), (2) our model with relaxed symmetry (SI-relaxed), and (3) weak SINDy as the baseline. The results for the two systems with symmetry breaking are shown in Figure 3 (center & right). As expected, SI has a much lower success probability when the symmetry-breaking factor becomes significant. Meanwhile, SI-relaxed remains highly competitive. It also has a clear advantage over baseline SINDy, showing that even if the inductive bias of symmetry is slightly inaccurate, our model with relaxed constraints is still better than a model without any knowledge of symmetry.

More comprehensive results, e.g., samples of discovered equations, are provided in Appendix D.

5 Discussion

In this paper, we propose to enforce symmetry in general methods for discovering symbolic differential equations by using the differential invariants of the symmetry group as the variable set in symbolic regression algorithms. We implement this general strategy in different classes of algorithms and observe improved accuracy, efficiency and robustness of equation discovery, especially in challenging scenarios such as noisy data and imperfect symmetry.

It should be noted that our method assumes the symmetry group is already given. This assumption aligns with common practice–physicists often begin by hypothesizing the symmetries of a system and seek governing equations allowed by those symmetries. However, our current framework cannot be applied if symmetry is unknown, and will produce incorrect results with misspecified symmetry. This can be potentially addressed by incorporating automated symmetry discovery methods for differential equations (Yang et al., 2024; Ko et al., 2024), which we leave for future work.

Another caveat of our method is the calculation of differential invariants. While solving for $\mathbf{v}^{(n)}(\eta) = 0$ and applying the formula (6) is easy with any symbolic computation package, the resulting differential invariants may be complicated and require ad-hoc adjustment for better interpretability and compatibility with specific algorithm implementations (e.g., conditions in Proposition 3.4). Fortunately, this only requires a one-time effort. Once we have derived the invariants for a symmetry group, the results can be reused for any equation admitting the same symmetry.

References

Akhound-Sadegh, T., Perreault-Levasseur, L., Brandstetter, J., Welling, M., and Ravanbakhsh, S. Lie point symmetry and physics informed networks. *arXiv preprint arXiv:2311.04293*, 2023.

Bakarji, J., Callaham, J., Brunton, S. L., and Kutz, J. N. Dimensionally consistent learning with buckingham pi. *Nature Computational Science*, 2:834–844, 12 2022. ISSN 2662-8457. doi: 10.1038/s43588-022-00355-5.

Biggio, L., Bendinelli, T., Neitz, A., Lucchi, A., and Parascandolo, G. Neural symbolic regression that scales. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 936–945. PMLR, 18–24 Jul 2021.

Brandstetter, J., Welling, M., and Worrall, D. E. Lie point symmetry data augmentation for neural pde solvers. In *International Conference on Machine Learning*, pp. 2241–2256. PMLR, 2022.

Brunton, S. L., Proctor, J. L., and Kutz, J. N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.

Champion, K., Lusch, B., Kutz, J. N., and Brunton, S. L. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019. doi: 10.1073/pnas.1906995116.

Cranmer, M. Interpretable machine learning for science with pysr and symbolicregression.jl, 2023.

Cranmer, M., Sanchez Gonzalez, A., Battaglia, P., Xu, R., Cranmer, K., Spergel, D., and Ho, S. Discovering symbolic models from deep learning with inductive biases. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 17429–17442. Curran Associates, Inc., 2020.

Cranmer, M. D., Xu, R., Battaglia, P., and Ho, S. Learning symbolic physics with graph networks, 2019.

Dalton, D., Husmeier, D., and Gao, H. Physics and lie symmetry informed gaussian processes. In *Forty-first International Conference on Machine Learning*, 2024.

de Silva, B., Champion, K., Quade, M., Loiseau, J.-C., Kutz, J., and Brunton, S. Pysindy: A python package for the sparse identification of nonlinear dynamical systems from data. *Journal of Open Source Software*, 5(49):2104, 2020. doi: 10.21105/joss.02104. URL `https://doi.org/10.21105/joss.02104`.

Dubčáková, R. Eureqa: software review, 2011.

Finzi, M., Benton, G., and Wilson, A. G. Residual pathway priors for soft equivariance constraints. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 30037–30049. Curran Associates, Inc., 2021. URL `https://proceedings.neurips.cc/paper_files/paper/2021/file/fc394e9935fbd62c8aedc372464e1965-Paper.pdf`.

Gaucel, S., Keijzer, M., Lutton, E., and Tonda, A. Learning dynamical systems using standard symbolic regression. In Nicolau, M., Krawiec, K., Heywood, M. I., Castelli, M., García-Sánchez, P., Merelo, J. J., Rivas Santos, V. M., and Sim, K. (eds.), *Genetic Programming*, pp. 25–36, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.

Grayeli, A., Sehgal, A., Costilla Reyes, O., Cranmer, M., and Chaudhuri, S. Symbolic regression with a learned concept library. *Advances in Neural Information Processing Systems*, 37:44678–44709, 2024.

Grundner, A., Beucler, T., Gentine, P., and Eyring, V. Data-driven equation discovery of a cloud cover parameterization. *arXiv preprint arXiv:2304.08063*, 2023.

Holt, S., Qian, Z., and van der Schaar, M. Deep generative symbolic regression. *arXiv preprint arXiv:2401.00282*, 2023.

Kaheman, K., Kutz, J. N., and Brunton, S. L. Sindy-pi: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proceedings of the Royal Society A*, 476(2242):20200279, 2020.

Kamienny, P.-A., d'Ascoli, S., Lample, G., and Charton, F. End-to-end symbolic regression with transformers. *Advances in Neural Information Processing Systems*, 35:10269–10281, 2022.

Kaptanoglu, A. A., de Silva, B. M., Fasel, U., Kaheman, K., Goldschmidt, A. J., Callaham, J., Delahunt, C. B., Nicolaou, Z. G., Champion, K., Loiseau, J.-C., Kutz, J. N., and Brunton, S. L. Pysindy: A comprehensive python package for robust sparse system identification. *Journal of Open Source Software*, 7(69):3994, 2022. doi: 10.21105/joss.03994. URL `https://doi.org/10.21105/joss.03994`.

Ko, G., Kim, H., and Lee, J. Learning infinitesimal generators of continuous symmetries from data. *arXiv preprint arXiv:2410.21853*, 2024.

Lee, K., Trask, N., and Stinis, P. Structure-preserving sparse identification of nonlinear dynamics for data-driven modeling. In Dong, B., Li, Q., Wang, L., and Xu, Z.-Q. J. (eds.), *Proceedings of Mathematical and Scientific Machine Learning*, volume 190 of *Proceedings of Machine Learning Research*, pp. 65–80. PMLR, 15–17 Aug 2022.

Martius, G. and Lampert, C. H. Extrapolation and learning equations. *arXiv preprint arXiv:1610.02995*, 2016.

Merler, M., Haitsiukevich, K., Dainese, N., and Marttinen, P. In-context symbolic regression: Leveraging large language models for function discovery. *arXiv preprint arXiv:2404.19094*, 2024.

Messenger, D. A. and Bortz, D. M. Weak sindy for partial differential equations. *Journal of Computational Physics*, 443:110525, 2021a.

Messenger, D. A. and Bortz, D. M. Weak sindy: Galerkin-based data-driven model selection. *Multiscale Modeling & Simulation*, 19(3):1474–1497, 2021b.

Mialon, G., Garrido, Q., Lawrence, H., Rehman, D., LeCun, Y., and Kiani, B. Self-supervised learning with lie symmetries for partial differential equations. *Advances in Neural Information Processing Systems*, 36:28973–29004, 2023.

Newell, A. C. *Solitons in mathematics and physics*. SIAM, 1985.

Olver, P. J. *Applications of Lie groups to differential equations*, volume 107. Springer Science & Business Media, 1993.

Olver, P. J. *Equivalence, invariants and symmetry*. Cambridge University Press, 1995.

Otto, S. E., Zolman, N., Kutz, J. N., and Brunton, S. L. A unified framework to enforce, discover, and promote symmetry in machine learning, 2023.

Petersen, B. K., Landajuela, M., Mundhenk, T. N., Santiago, C. P., Kim, S. K., and Kim, J. T. Deep symbolic regression: Recovering mathematical expressions from data via risk-seeking policy gradients. *arXiv preprint arXiv:1912.04871*, 2019.

Qian, Z., Kacprzyk, K., and van der Schaar, M. D-code: Discovering closed-form odes from observed trajectories. In *International Conference on Learning Representations*, 2022.

Rao, C., Ren, P., Liu, Y., and Sun, H. Discovering nonlinear pdes from scarce data with physics-encoded learning. *arXiv preprint arXiv:2201.12354*, 2022.

Rudy, S. H., Brunton, S. L., Proctor, J. L., and Kutz, J. N. Data-driven discovery of partial differential equations. *Science advances*, 3(4):e1602614, 2017.

Sahoo, S., Lampert, C., and Martius, G. Learning equations for extrapolation and control. In *International Conference on Machine Learning*, pp. 4442–4450. PMLR, 2018.

Schmidt, M. and Lipson, H. Distilling free-form natural laws from experimental data. *science*, 324 (5923):81–85, 2009.

Shojaee, P., Meidani, K., Gupta, S., Farimani, A. B., and Reddy, C. K. Llm-sr: Scientific equation discovery via programming with large language models. *arXiv preprint arXiv:2404.18400*, 2024.

Shojaee, P., Nguyen, N.-H., Meidani, K., Farimani, A. B., Doan, K. D., and Reddy, C. K. Llm-srbench: A new benchmark for scientific equation discovery with large language models. *arXiv preprint arXiv:2504.10415*, 2025.

Takamoto, M., Praditia, T., Leiteritz, R., MacKinlay, D., Alesiani, F., Pflüger, D., and Niepert, M. Pdebench: An extensive benchmark for scientific machine learning. *Advances in Neural Information Processing Systems*, 35:1596–1611, 2022.

Udrescu, S.-M. and Tegmark, M. Ai feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6(16):eaay2631, 2020.

Udrescu, S.-M., Tan, A., Feng, J., Neto, O., Wu, T., and Tegmark, M. Ai feynman 2.0: Pareto-optimal symbolic regression exploiting graph modularity. *Advances in Neural Information Processing Systems*, 33:4860–4871, 2020.

Wang, R., Walters, R., and Yu, R. Incorporating symmetry into deep dynamics models for improved generalization. In *International Conference on Learning Representations*, 2021.

Wang, R., Walters, R., and Yu, R. Approximately equivariant networks for imperfectly symmetric dynamics. In *International Conference on Machine Learning*. PMLR, 2022.

Wang, Y., Wagner, N., and Rondinelli, J. M. Symbolic regression in materials science. *MRS Communications*, 9(3):793–805, 2019.

Xie, X., Samaei, A., Guo, J., Liu, W. K., and Gan, Z. Data-driven discovery of dimensionless numbers and governing laws from scarce measurements. *Nature Communications*, 13(1):7562, 2022. doi: 10.1038/s41467-022-35084-w.

Yang, J., Rao, W., Dehmamy, N., Walters, R., and Yu, R. Symmetry-informed governing equation discovery. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Zhang, Z.-Y., Zhang, H., Zhang, L.-S., and Guo, L.-L. Enforcing continuous symmetries in physics-informed neural network for solving forward and inverse problems of partial differential equations. *Journal of Computational Physics*, 492:112415, 2023.

A Related Works

Symbolic Regression. Given the dataset $\{(x^i, y^i)\} \subset X \times Y$, symbolic regression (SR) aims to model the function $y = f(x)$ by a symbolic equation. A popular method for symbolic regression is genetic programming (GP) (Schmidt & Lipson, 2009; Gaucel et al., 2014), which leverages evolutionary algorithms to explore the space of possible equations and has demonstrated success in uncovering governing laws in various scientific domains such as material science (Wang et al., 2019), climate modeling (Grundner et al., 2023), cosmology (Cranmer et al., 2020), etc. Various software have been developed for GP-based symbolic regression, e.g. Eureqa (Dubčáková, 2011) and PySR (Cranmer, 2023).

Another class of methods is sparse regression (Brunton et al., 2016), which assumes the function to be discovered can be written as a linear combination of predefined candidate functions and solves for the coefficient matrix. It has also been extended to discover more general equations, such as equations in latent variables (Champion et al., 2019) and PDEs (Rudy et al., 2017).

Neural networks have also shown their potential in symbolic regression. Martius & Lampert (2016); Sahoo et al. (2018) represents a few earliest attempts, where they replace the activation functions in fully connected networks with math operators and functions, so the network itself translates to a symbolic formula. Other works represent mathematical expressions as sequences of tokens and train neural networks to predict the sequence given a dataset of input-output pairs. For example, Petersen et al. (2019) trains an RNN with policy gradients to minimize the regression error. Biggio et al. (2021), Kamienny et al. (2022) and Holt et al. (2023) pre-train an encoder-decoder network over a large amount of procedurally generated equations and query the pretrained model on a new dataset of input-output pairs at test time.

The aforementioned symbolic regression methods can be improved by incorporating specific domain knowledge. For example, AI Feynman (Udrescu & Tegmark, 2020; Udrescu et al., 2020) uses properties like separability and compositionality to simplify the data. Cranmer et al. (2020) specifies the overall skeleton of the equation and fits each part with genetic programming independently. The goal of this paper falls into this category – to use the knowledge of symmetry to reduce the search space of symbolic regression and improve its accuracy and efficiency.

Recently, Large Language Models (LLMs) have emerged as an alternative for SR, using pre-trained scientific priors to propose sequential hypothesis (Merler et al., 2024) or to guide genetic programming (Shojaee et al., 2024), balancing the efficiency of domain knowledge with the robustness of evolutionary search. However, current LLM-based methods often rely on memorizing known equations rather than facilitating genuine discovery, and their guidance lacks interpretability, specifically, the reasoning behind their suggestions, evidenced by a recent benchmark specially designed for LLM-SR (Shojaee et al., 2025). A recent effort sought to improve interpretability by binding symbolic evolution with natural language explanations (Grayeli et al., 2024). However, this method relies on frontier LLMs to conduct the evolution of the natural language components, rendering the process itself opaque. These limitations highlight the need for approaches that enhance the controllability and explainability of the prior knowledge injected, ensuring more transparent and trustworthy discovery.

Discovering Differential Equations. While it remains in the scope of symbolic regression, the discovery of differential equations poses additional challenges because the derivatives are not directly observed from data. Building upon the aforementioned SINDy sparse regression (Brunton et al., 2016), Messenger & Bortz (2021a,b) formulates an alternative optimization problem based on the variational form of differential equations and bypasses the need for derivative estimation. A similar variational approach is also applied to genetic programming (Qian et al., 2022). Various other improvements have been made, including refined training procedure (Rao et al., 2022), relaxed assumptions about the form of the equation (Kaheman et al., 2020), and the incorporation of physical priors (Xie et al., 2022; Bakarji et al., 2022; Lee et al., 2022).

PDE Symmetry in Machine Learning. Symmetry is an important inductive bias in machine learning. In the context of learning differential equation systems, many works encourage symmetry in their models through data augmentation (Brandstetter et al., 2022), regularization terms (Akhound-Sadegh et al., 2023; Zhang et al., 2023; Dalton et al., 2024), and self-supervised learning (Mialon et al., 2023). Strictly enforcing symmetry is also possible, but is often restricted to specific symmetries and systems (Wang et al., 2021). For more general symmetries and physical systems, enforcing symmetry

often requires additional assumptions on the form of equations, such as the linear combination form in sparse regression (Otto et al., 2023; Yang et al., 2024). To the best of our knowledge, our work is the first attempt to strictly enforce general symmetries of differential equations for general symbolic regression methods.

B Math

B.1 Notations

Table 2: Descriptions of symbols used throughout the paper. The three blocks include (1) basic notations for PDEs, (2) notations for Lie symmetry of PDEs, and (3) notations for symbolic regression algorithms and miscellaneous.

Symbols	Descriptions
p	Number of independent variables of a PDE.
q	Number of dependent variables of a PDE.
X	Space of independent variables of a PDE: $X \subset \mathbb{R}^p$. Also used to denote the feature space of SR algorithms.
U	Space of dependent variables of a PDE: $U \subset \mathbb{R}^q$. Assumed to be 1-dimensional unless otherwise stated.
E	Total space of all variables of a PDE: $E = X \times U$.
U_k	Space of strictly kth-order partial derivatives of variables in U w.r.t variables in X.
$U^{(n)}$	Space of all partial derivatives up to nth order (including the original variables in U): $U^{(n)} = U \times U_1 \times \cdots \times U_n$.
$M^{(n)}$	nth-order jet space: $M^{(n)} \subset X \times U^{(n)}$.
$T\mathcal{M}$	The tangent bundle of a manifold $\mathcal{M}$.
$\mathbf{x}$	Independent variables of a PDE: $\mathbf{x} \in \mathbb{R}^p$.
t	Time variable.
x, y	Spatial variables in PDE contexts. Also used to denote the features and labels of SR algorithms, where x can denote multi-dimensional features.
$u, \mathbf{u}$	Dependent variable(s) of a PDE: $u \in \mathbb{R}$ and $\mathbf{u} \in \mathbb{R}^q$.
$u^{(n)}, \mathbf{u}^{(n)}$	The collection of all up to n-th order partial derivatives of u or $\mathbf{u}$.
df	The (ordinary) differential of a function. For a differential function $f : M^{(n)} \to \mathbb{R}$, $df = \sum_j \frac{\partial f}{\partial x^j} dx^j + \sum_\alpha \frac{\partial f}{\partial u_\alpha} du_\alpha$.
$D_i f$	The total derivative of a differential function $f : M^{(n)} \to \mathbb{R}$ w.r.t the ith independent variable. For example, if $p = q = 1$, $D_1 f = \frac{\partial f}{\partial x} + \sum_{k=0}^{\infty} u_{k+1} \frac{\partial f}{\partial u_k}$, where $u_k \coloneqq \partial^k u / \partial x^k$.
Df	The total differential of a differential function $f : M^{(n)} \to \mathbb{R}$, i.e. $Df = D_i f \, dx^i$.
g	A group element with an action on E (2).
$\mathbf{v}$	A vector field on the total space E (3), representing an infinitesimal transformation. A list of multiple vector fields are indexed by subscripts.
$\mathrm{pr}^{(n)} g$	nth-order prolongation of g acting on $M^{(n)}$.
$\mathrm{pr}^{(n)} \mathbf{v}$	nth-order prolongation of $\mathbf{v}$ acting on $M^{(n)}$.
$g^{(n)}, \mathbf{v}^{(n)}$	Equivalent to $\mathrm{pr}^{(n)} g$ and $\mathrm{pr}^{(n)} \mathbf{v}$, respectively.
$\mathrm{pr}\ \mathbf{v}$	The (infinite) prolongation of $\mathbf{v}$. For an nth-order differential function $f(\mathbf{x}, \mathbf{u}^{(n)})$, $\mathrm{pr}\ \mathbf{v}(f) = \mathrm{pr}^{(n)} \mathbf{v}(f)$.
η, ζ, ϑ	Differential invariants of a symmetry group. η is used by default. The other letters are used to distinguish between invariants of different orders.
$\ell, \boldsymbol{\ell}$	The LHS of SINDy equation (4). Often assumed to be time derivatives.
$\boldsymbol{\theta}$	A column vector containing all SINDy library functions: $\boldsymbol{\theta} = [\theta^1, \cdots, \theta^m]$
$\mathbf{w}, W$	The SINDy parameters. For only one equation, $\mathbf{w} = [w^1, \cdots, w^m]$ is a row vector. For multiple equations, $W = [w^{ij}]$ is a $q \times m$ matrix.
$\mathbf{X}, \mathbf{y}$	Concatenated matrix/vector of features/labels of all datapoints for symbolic regression.
$[N]$	List of positive integers up to N, i.e. $[1, 2, \cdots, N]$ for any $N \in \mathbb{Z}^+$.
$1 : N$	Equivalent to $[N]$.
LHS, RHS	Left- and Right-hand side of an equation.

B.2 Extended Background on PDE Symmetry

References for the below material include Olver (1993), Olver (1995).

Prolonged group actions Let $E = X \times U \simeq \mathbb{R}^p \times \mathbb{R}^q$ be endowed with the action of a group G via point transformations. Then group elements $g \in G$ act locally on functions $\mathbf{u} = f(\mathbf{x})$, therefore also on derivatives of these functions. This in turn induces, at least pointwise, "prolonged" transformations on jet spaces: $(\tilde{\mathbf{x}}, \tilde{\mathbf{u}}^{(n)}) = \mathrm{pr}^{(n)} g \cdot (\mathbf{x}, \mathbf{u}^{(n)})$.

Let $J = (j_1, \ldots, j_n), 1 \leq j_\nu \leq p$ be an n-tuple of indices of independent variables and $1 \leq \alpha \leq q$. We will use the shorthand

$$u^\alpha_J := \frac{\partial^J u^\alpha}{\partial x^J} := \frac{\partial^{|J|} u^\alpha}{\partial x_{j_1} \cdots \partial x_{j_n}}$$

and

$$D_J := D_{j_1} \cdots D_{j_n}.$$

It is not practical to work explicitly with prolonged group transformations. Therefore one linearizes and considers the prolonged action of the infinitesimal generators of G. Explicitly, given a vector field

$$\mathbf{v} = \sum_{i=1}^{p} \xi^i(\mathbf{x}, \mathbf{u}) \frac{\partial}{\partial x^i} + \sum_{\alpha=1}^{q} \varphi^\alpha(\mathbf{x}, \mathbf{u}) \frac{\partial}{\partial u^\alpha},$$

its **characteristic** is a q-tuple $Q = (Q^1, \ldots, Q^q)$ of functions with

$$Q^\alpha(\mathbf{x}, \mathbf{u}^{(1)}) = \varphi^\alpha(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^{p} \xi^i(\mathbf{x}, \mathbf{u}) \frac{\partial u^\alpha}{\partial x^i}.$$

Now the prolongation of $\mathbf{v}$ to order n is defined by

$$\mathrm{pr}^{(n)} \mathbf{v} = \sum_{i=1}^{p} \xi^i(\mathbf{x}, \mathbf{u}) \frac{\partial}{\partial x^i} + \sum_{\alpha=1}^{q} \sum_{\#J=n} \varphi^\alpha_J(\mathbf{x}, \mathbf{u}^{(n)}) \frac{\partial}{\partial u^\alpha_J}. \tag{10}$$

Here J ranges over all n-tuples $J = (j_1, \ldots, j_n), 1 \leq j_\nu \leq p$ and the φ^α_J are given by

$$\varphi^\alpha_J = D_J Q^\alpha + \sum_{i=1}^{p} \xi^i \mathbf{u}^\alpha_{J,i}.$$

We remark that the prolongation of $\mathbf{v}$ has been described explicitly in terms of the coefficients of $\mathbf{v}$ and their derivatives.

B.3 Proof of Proposition 3.3

Olver (1995) provides the following general theorem to construct higher-order differential invariants from a contact-invariant coframe. We refer the readers to Chapter 5 of Olver (1995) for definitions of relevant concepts, e.g., contact forms and contact-invariant forms and coframes.

Theorem B.1 (Thm. 5.48, (Olver, 1995))**.** *Let G be a transformation group acting on a space with p independent variables and q dependent variables. Suppose $\omega^1, ..., \omega^p$ is a contact-invariant coframe for G, and let $\mathcal{D}_j$ be the associated invariant differential operators defined via $Df = D_j f \; dx^j = \mathcal{D}_j f \; \omega^j$. If there are a maximal number of independent, strictly nth-order differential invariants $\zeta^1, \cdots, \zeta^{q_n}$, $q_n = \binom{p+n-1}{n}$, then the set of differentiated invariants $\mathcal{D}_i \zeta^\nu$, $i \in [p]$, $\nu \in [q_n]$, contains a complete set of independent, strictly $(n+1)$th-order differential invariants.*

Specifically, the condition that there exist a maximal number of differential invariants of order exactly n is guaranteed if n is at least $\dim G$.

Our proposition is a derived result from the above theorem, which provides a concrete way of computation from lower-order invariants to higher-order ones:

Proposition B.2. *Let G be a local group acting on $X \times U \simeq \mathbb{R}^p \times \mathbb{R}$. Let $\eta^1, \eta^2, \cdots, \eta^p$ be any p differential invariants of G whose horizontal Jacobian $J = [D_i \eta^j]$ is non-degenerate on an open subset $\Omega \subset M^{(n)}$. If there are a maximal number of independent, strictly nth-order differential*

invariants $\zeta^1, \cdots, \zeta^{q_n}$, $q_n = \binom{p+n-1}{n}$, then the following set contains a complete set of independent, strictly $(n+1)$th-order differential invariants defined on Ω:

$$\frac{\det(D_i \tilde{\eta}^j_{(k,k')})}{\det(D_i \eta^j)}, \ \forall k \in [p], k' \in [q_n], \tag{11}$$

where $i, j \in [p]$ are matrix indices, D_i denotes the total derivative w.r.t i-th independent variable and $\tilde{\eta}^j_{(k,k')} = [\eta^1, ..., \eta^{k-1}, \zeta^{k'}, \eta^{k+1}, ..., \eta^p]$.

Proof. We show that the total differentials of the differential invariants $\eta^1, ..., \eta^p$ can be used to construct a *contact-invariant coframe* of G and then derive the associated invariant differential operators to complete the proof.

First, note that for any differential invariant η of G, its total differential $\omega = D\eta = D_j \eta \, dx^j$ can be written as

$$\omega = \omega_o + \theta, \tag{12}$$

where $\omega_o \coloneqq d\eta = \sum_{i \in [p]} \frac{\partial F}{\partial x^i} dx^i + \sum_{|\alpha| \leq n} \frac{\partial F}{\partial u_\alpha} du_\alpha$ is the ordinary differential of $\eta : M^{(n)} \to \mathbb{R}$ and θ is a contact form.

Since η is a differential invariant, its differential $\omega_o = d\eta$ is an invariant one-form on $M^{(n)}$, i.e. $(g^{(n)})^* \omega_o = \omega_o$.

Also, a prolonged group action maps contact forms to contact forms. To see this, note that a prolonged group action $g^{(n)}$ maps the prolonged graph of any function to the prolonged graph of a transformed function. Then, for any contact form θ, $(g^{(n)})^* \theta$ is annihilated by all prolonged functions $f^{(n)}$, thus a contact form by definition:

$$\begin{aligned}(f^{(n)})^*((g^{(n)})^* \theta) &= (g^{(n)} \circ f^{(n)})^* \theta \\ &= ((g \cdot f)^{(n)})^* \theta \\ &= 0.\end{aligned} \tag{13}$$

Then, from (12), we have

$$\begin{aligned}(g^{(n+1)})^* \omega &= (g^{(n)})^* \omega_o + (g^{(n+1)})^* \theta \\ &= \omega_o + \theta' \\ &= \omega + (\theta' - \theta)\end{aligned} \tag{14}$$

where θ' is some contact form and so is $\theta' - \theta$. Thus, ω is contact-invariant. For the p differential invariants $\eta^1, \cdots, \eta^p$, we have p contact-invariant one-forms $\omega^1, \cdots, \omega^p$, respectively.

Next, we prove that $\omega^1, \cdots, \omega^p$ are linearly independent and form a coframe. Assume there exists smooth coefficients c^j such that $\sum_j c^j \omega^j = 0$. Then, regrouping the coefficients of the horizontal forms dx^i, we have

$$0 = \sum_{i,j} c^j D_i \eta^j dx^i = \sum_i \left(\sum_j c^j D_i \eta^j \right) dx^i. \tag{15}$$

Because the dx^i are linearly independent, each coefficient of dx^i must vanish, i.e. $J_i{}^j c^j = 0$. Since the Jacobian $J = [D_i \eta^j]$ is non-degenerate, the only solution is $c^j = 0$ (on the open subset $\Omega \in M^{(n)}$). Thus, $\omega^1, \cdots, \omega^p$ form a contact-invariant coframe. According to Theorem B.1, the associated invariant differential operators of the coframe take a complete set of same-order invariants to a complete set of one-order-higher invariants.

The remaining step is to obtain the invariant differential operators explicitly in terms of η^j. Recall the formula in Theorem B.1 that defines the invariant differential operators:

$$D_i f \, dx^i = \mathcal{D}_j f \, \omega^j. \tag{16}$$

Expanding $\omega^j = D\eta^j = D_i\eta^j \, dx^i$, we have the following linear system of invariant differential operators $\mathcal{D}_j$:

$$\begin{bmatrix} D_1 \\ D_2 \\ \vdots \\ D_p \end{bmatrix} = \begin{bmatrix} D_1\eta^1 & D_1\eta^2 & \cdots & D_1\eta^p \\ D_2\eta^1 & D_2\eta^2 & \cdots & D_2\eta^p \\ \vdots & \vdots & & \vdots \\ D_p\eta^1 & D_p\eta^2 & \cdots & D_p\eta^p \end{bmatrix} \begin{bmatrix} \mathcal{D}_1 \\ \mathcal{D}_2 \\ \vdots \\ \mathcal{D}_p \end{bmatrix}. \tag{17}$$

Since $J = [D_i\eta^j]$ is non-degenerate, Cramer's rule yields

$$\mathcal{D}_k\zeta = \frac{\det(D_i\eta^1 \mid \cdots \mid D_i\eta^{k-1} \mid D_i\zeta \mid D_i\eta^{k+1} \mid \cdots \mid D_i\eta^p)}{\det(D_i\eta^j)}. \tag{18}$$

□

Remark B.3. We require that the differential invariants $\eta^1, \cdots, \eta^p$ has a nondegenerate horizontal Jacobian $[D_i\eta^j]$, which is a stronger condition than functional independence. Since the differential invariants are functions on the jet space, it is possible that a set of such functions is functionally independent, i.e., has a nondegenerate full Jacobian $[\partial_i\eta^j]$, where $i \in [q_n]$ indexes the jet space variables $(\mathbf{x}, u^{(n)})$, but has a lower-rank horizontal Jacobian. For example, consider $\eta^1 = u_x$ and $\eta^2 = u_y$. In the full Jacobian, $\partial\eta^j/\partial u_x$ and $\partial\eta^j/\partial u_y$ form the identity, so it has full rank. However, its horizontal Jacobian containing total derivatives is given by $\begin{bmatrix} u_{xx} & u_{xy} \\ u_{xy} & u_{yy} \end{bmatrix}$, which is not invertible on the subset of the jet space where $u_{xx}u_{yy} - u_{xy}^2 = 0$.

In practice, this non-degeneracy condition can be easily checked once we have the symbolic expressions of the p differential invariants.

Remark B.4. When $p = 1$, Proposition B.2 is equivalent to the following (Prop. 2.53, Olver (1993)): If $y = \eta(x, u^{(n)})$ and $w = \zeta(x, u^{(n)})$ are n-th order differential invariants of G, then $\frac{dw}{dy} \equiv \frac{D_x\zeta}{D_x\eta}$ is an $(n+1)$-th order differential invariant of G. Specifically, if $y = \eta(x, u)$ and $w = \zeta(x, u, u_x)$ form a complete set of functionally independent differential invariants of $\mathrm{pr}^{(1)}G$, the complete set of functionally independent differential invariants for $\mathrm{pr}^{(n)}G$ is then given by

$$y, w, dw/dy, ..., d^{n-1}w/dy^{n-1}. \tag{19}$$

B.4 Examples of Computing Differential Invariants

Example B.5. Consider the group $\mathrm{SO}(2)$ acting on $X \times U \simeq \mathbb{R}^2 \times \mathbb{R}$ by standard rotation in the 2D space of independent variable and trivial action on U, i.e. its infinitesimal generator given by $\mathbf{v} = y\partial_x - x\partial_y$.

First, we solve for a complete set of the ordinary and first-order invariants. The two ordinary invariants are given by $\eta_1(x, y, u) = \frac{1}{2}(x^2 + y^2)$ and $\eta_2(x, y, u) = u$. (6) dictates how we construct higher-order invariants using these two functionally independent invariants and another arbitrary invariant. For notational convenience, we convert (6) to operators defined according to η_2 and η_1, respectively:

$$\mathcal{O}_1 = \frac{xD_y - yD_x}{xu_y - yu_x} \tag{20}$$

$$\mathcal{O}_2 = \frac{u_yD_x - u_xD_y}{xu_y - yu_x} \tag{21}$$

Then, we need to find another new differential invariant, because applying these operators on η_1 and η_2 leads to trivial results. Since η_1 and η_2 generate all ordinary (zeroth-order) invariants, we must look for the first-order invariants. To do this, note the prolonged vector field is given by

$$\mathrm{pr}^{(1)}\mathbf{v} = \mathbf{v} + u_y\partial_{u_x} - u_x\partial_{u_y} \tag{22}$$

Solving for $\mathrm{pr}^{(1)}\mathbf{v}$ gives two first-order invariants, $\zeta_1 = xu_y - yu_x$ and $\zeta_2 = xu_x + yu_y$. Note that the differential invariant ζ_1 is exactly the common denominator in $\mathcal{O}_1$ and $\mathcal{O}_2$, so we can simplify $\mathcal{O}_1$

and $\mathcal{O}_2$ by using only their numerators, i.e.

$$\mathcal{O}_1 = xD_y - yD_x \tag{23}$$
$$\mathcal{O}_2 = u_yD_x - u_xD_y \tag{24}$$

Note that $\mathcal{O}_2$ has first-order coefficients, which may complicate things in the subsequent calculation. Denoting the space of all continuous functions of the existing four invariants as $\mathcal{I} = \mathcal{C}(\eta_1, \eta_2, \zeta_1, \zeta_2)$, we can choose any new operator within the $\mathcal{I}$-module spanned by $\mathcal{O}_1$ and $\mathcal{O}_2$ that makes things easier. Specifically, we use the following operator

$$\begin{aligned}\tilde{\mathcal{O}}_2 &= \frac{\zeta_2}{\zeta_1}\mathcal{O}_1 + \frac{2\eta_1}{\zeta_1}\mathcal{O}_2 \\ &= xD_x + yD_y\end{aligned} \tag{25}$$

Then, we apply these operators to the first-order invariants, which raise the order by one and give us the second-order invariants. For example, applying O_1 to ζ_1, we have

$$\begin{aligned}\mathcal{O}_1\zeta_1 &= xD_y\zeta_1 - yD_x\zeta_1 \\ &= x(xu_{yy} - u_x - yu_{xy}) - y(u_y + xu_{xy} - yu_{xx}) \\ &= x^2u_{yy} + y^2u_{xx} - xu_x - yu_y - 2xyu_{xy}\end{aligned} \tag{26}$$

Note that $\zeta_2 = xu_x + yu_y$ is a first-order invariant, so we can further remove it from the formula and get a simplified second-order invariant

$$\vartheta_1 = x^2u_{yy} + y^2u_{xx} - 2xyu_{xy} \tag{27}$$

Similarly, we compute $\mathcal{O}_1\zeta_2$, $\tilde{\mathcal{O}}_2\zeta_1$ and $\tilde{\mathcal{O}}_2\zeta_2$ and obtain the following, respectively:

$$\begin{aligned}\vartheta_2 = \vartheta_3 &= \zeta_1 + xy(u_{yy} - u_{xx}) + (x^2 - y^2)u_{xy} \\ &\equiv xy(u_{yy} - u_{xx}) + (x^2 - y^2)u_{xy}\end{aligned} \tag{28}$$

$$\begin{aligned}\vartheta_4 &= \zeta_2 + x^2u_{xx} + y^2u_{yy} + 2xyu_{xy} \\ &\equiv x^2u_{xx} + y^2u_{yy} + 2xyu_{xy}\end{aligned} \tag{29}$$

The above 8 invariants should form a complete set of second-order differential invariants of $\mathbf{v} = x\partial_y - y\partial_x$. To verify, note that the Laplacian $\Delta u = u_{xx} + u_{yy}$, which is a well-known rotational invariant, can be written in terms of these differential functions:

$$\begin{aligned}\Delta u = u_{xx} + u_{yy} &= \frac{(x^2 + y^2)(u_{xx} + u_{yy})}{x^2 + y^2} \\ &= \frac{\vartheta_1 + \vartheta_4}{2\eta_1}\end{aligned} \tag{30}$$

Another second-order rotational invariant, the trace of the squared Hessian matrix, $u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2$, is recovered by

$$u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2 = \frac{\vartheta_1^2 + 2\vartheta_2^2 + \vartheta_4^2}{4\eta_1^2} \tag{31}$$

On the other hand, these 8 invariants are apparently not functionally independent - note that $\vartheta_2 = \mathcal{O}_1\zeta_2$ and $\vartheta_3 = \tilde{\mathcal{O}}_2\zeta_1$ are the same. While this may be some coincidence, eventually it is not surprising because we would expect to see 3 functionally independent strictly second-order differential invariants instead of 4, since $(u_{xx}, u_{yy}, u_{xy}) \in U_2$ is only 3-dimensional.

Example B.6 (Scaling and translation). Consider the vector field $\mathbf{v}_1 = t\partial_t + ax\partial_x + bu\partial_u$. It generates the scaling symmetry $t \mapsto \lambda t, x \mapsto \lambda^a x, u \mapsto \lambda^b u$. The ordinary invariants of this symmetry are t^bu^{-1} and x^au^{-1}. The higher-order invariants are given by $\eta_{(\alpha,\beta)} = x^\alpha t^\beta u_{x^{(\alpha)}t^{(\beta)}}u^{-1}$, where α and β denote the orders of partial derivatives w.r.t t and x, e.g. $u_{x^{(2)}t^{(1)}} \coloneqq u_{xxt}$.

Besides the scaling symmetry, we can consider other common symmetries simultaneously, e.g. translation symmetries in both space and time, $\mathbf{v}_2 = \partial_x$ and $\mathbf{v}_3 = \partial_t$. These symmetries, along with the scaling symmetry $\mathbf{v}_1$, span a three-dimensional symmetry group. There are no ordinary invariants due to the translation symmetries. A convenient maximal set of functionally independent differential invariants is given by

$$\eta_{(\alpha,\beta)} = u_{x^{(\alpha)}t^{(\beta)}}u_x^{\frac{b-a\alpha-\beta}{a-b}},\ \alpha \geq 0, \beta \geq 0. \tag{32}$$

B.5 Proof of Proposition 3.4

Proposition 3.4, restated below, aligns our symmetry constraint into the SINDy framework and results in a set of constraints on the SINDy parameters.

Proposition B.7. *Let $\boldsymbol{\ell}(\mathbf{x}, \mathbf{u}^{(n)}) = W\boldsymbol{\theta}(\mathbf{x}, \mathbf{u}^{(n)})$ be a system of q differential equations admitting a symmetry group G, where $\mathbf{x} \in \mathbb{R}^p$, $\mathbf{u} \in \mathbb{R}^q$, $\boldsymbol{\theta} \in \mathbb{R}^m$. Assume that there exist some nth-order invariants of G, $\eta_0^{1:q}$ and $\eta^{1:K}$, s.t. (1) the system of differential equations can be expressed as $\boldsymbol{\eta}_0 = W'\boldsymbol{\theta}'(\boldsymbol{\eta})$, where $\boldsymbol{\eta}_0 = [\eta_0^{1:q}]$ and $\boldsymbol{\eta} = [\eta^{1:K}]$, and (2) $\eta_0^i = T^{ijk}\theta^k\ell^j$ and $(\theta')^i = S^{ij}\theta^j$, for some library functions $\boldsymbol{\theta}'(\boldsymbol{\eta})$ and some constant tensors W', T and S. Then, the space of all possible W is a linear subspace of $\mathbb{R}^{q\times m}$.*

Proof. (Note: In this proof, we do not distinguish between superscripts and subscripts. All are used for tensor indices, not partial derivatives.)

For simplicity, we omit the dependency of functions and write

$$\ell^i = W^{ij}\theta^j. \tag{33}$$

Combining the conditions about the differential invariants, we know that the equation can be equivalently expressed as

$$T^{ijk}\theta^k\ell^j = (W')^{ij}S^{jk}\theta^k \tag{34}$$

for some $W' \in \mathbb{R}^{q\times m'}$, where m' is the number of invariant functions in $\boldsymbol{\theta}'$.

Substituting (33) into (34) and rearranging the indices, the principle of symmetry invariants then translates to the following constraint on W: there exists some $W' \in \mathbb{R}^{q\times m'}$ s.t.

$$T_i{}^{rk}\theta_k W_r{}^l\theta_l = (W')_i{}^k S_k{}^j\theta_j, \forall \mathbf{x}, \mathbf{u}^{(n)}. \tag{35}$$

To solve for W, we first eliminate the dependency on the variables $\mathbf{x}$ and $\mathbf{u}^{(n)}$ from the equation. We adopt a procedure similar to Yang et al. (2024). Denote $\mathbf{z} = (\mathbf{x}, \mathbf{u}^{(n)})$. Define a functional $M_{\boldsymbol{\theta}}$ as mapping a function to its coordinate in the function space spanned by $\boldsymbol{\theta}$, i.e. $M_{\boldsymbol{\theta}} : (\mathbf{z} \mapsto c^j\theta_j(\mathbf{z})) \mapsto (c^1, c^2, \cdots, c^m)$. Before we proceed, note that the LHS of (35) contains the products of functions $\theta_k(\mathbf{z})\theta_l(\mathbf{z})$, which may or may not be included in the original function library $\boldsymbol{\theta}$. Therefore, we denote $\tilde{\boldsymbol{\theta}}(\mathbf{z}) = [\boldsymbol{\theta}(\mathbf{z}) \mid\mid \{\theta_k\theta_l \notin \boldsymbol{\theta}\}]$ as the collection of all library functions θ_k and all their products $\theta_k\theta_l$. The invariant functions $\boldsymbol{\theta}'(\boldsymbol{\eta})$ can also be rewritten in terms of the prolonged library: $\boldsymbol{\theta}'(\boldsymbol{\eta}) = \tilde{S}\tilde{\boldsymbol{\theta}}$, where $\tilde{S}_{1:m} = S$.

Then, applying $M_{\tilde{\boldsymbol{\theta}}}$ to (35), we have

$$M_{\tilde{\boldsymbol{\theta}}}(T_i{}^{rk}\theta_k W_r{}^l\theta_l) = (W')_i{}^k\tilde{S}_k{}^j. \tag{36}$$

Further expanding the LHS, we have

$$T_i{}^{rk}W_r{}^l\Gamma_{kl}{}^j = (W')_i{}^k\tilde{S}_k{}^j, \tag{37}$$

where Γ satisfies $\theta_k\theta_l = \Gamma_{kl}{}^j\tilde{\theta}_j$. In other words, the rows of the LHS fall in the row space of $\tilde{S}$. Let $\tilde{S}^\perp$ be the basis matrix for the null space of $\tilde{S}$, i.e. $\tilde{S}\tilde{S}^\perp = 0$, we have

$$T_i{}^{rk}W_r{}^l\Gamma_{kl}{}^j(\tilde{S}^\perp)_{js} = 0, \tag{38}$$

suggesting that W must lie in a linear subspace of $\mathbb{R}^{q\times m}$.

□

Remark B.8. In practice, to solve for (38), we first rearrange (38) into $M\mathrm{vec}(W) = 0$, where M has shape $(\tilde{S}.\text{shape}[2] \times q,\ q \times m)$. Then, we perform SVD on M and apply a threshold of 10^{-6} to the singular values. The right singular vectors corresponding to the singular values smaller than the threshold then form a basis of the linear subspace $\mathrm{vec}(W)$ lies in.

C Implementation Details

This section discusses some detailed considerations in implementing the sparse regression-based methods described in Section 3.3 and 3.4. Contents include:

- Appendix C.1: An algorithmic description of direct sparse regression with symmetry invariants.
- Appendix C.2: Converting the symmetry invariant condition as linear constraints on the sparse regression parameters.
- Appendix C.3: Using differential invariants in weak SINDy via the linear constraints, as well as other considerations.

C.1 Direct Sparse Regression With Symmetry Invariants

The first approach to enforcing symmetry in sparse regression, as discussed in Section 3.3, is to directly use the symmetry invariants as the variables and their functions specified by a function library as the RHS features. Similar to Algorithm 1 for general symbolic regression methods, we provide a detailed algorithm for sparse regression below. Following the setup from SINDy, we aim to discover a system of q differential equations for q dependent variables.

Algorithm 2 Sparse regression with symmetry invariants

Require: PDE order n, dataset $\{\mathbf{z}^i = (\mathbf{x}^i, (\mathbf{u}^{(n)})^i) \in M^{(n)}\}_{i=1}^{N_D}$, SINDy LHS $\boldsymbol{\ell}$, SINDy function library $\{\theta^j\}$, infinitesimal generators of the symmetry group $\mathcal{B} = \{\mathbf{v}_a\}$.
Ensure: A PDE system admitting the given symmetry group.
Compute the symmetry invariants of $\mathcal{B}$ up to nth-order: $\eta^1, \cdots, \eta^K$. {Proposition 3.3}
Choose an invariant function η^{k_i} s.t. $\partial \eta^{k_i} / \partial \ell^i \neq 0$ for SINDy LHS component ℓ^i.
Let $\boldsymbol{\eta}_0 = [\eta^{k_1}, ..., \eta^{k_q}]^T$ and $\boldsymbol{\eta}$ denote the column vector containing the remaining $K-q$ invariants.
Instantiate the sparse regression model as $\boldsymbol{\eta}_0 = W\boldsymbol{\theta}(\boldsymbol{\eta})$.
Optimize W with the SINDy objective: $\sum_i \|\boldsymbol{\eta}_0(\mathbf{z}^i) - W\boldsymbol{\theta}(\boldsymbol{\eta}(\mathbf{z}^i))\|^2 + \lambda \|W\|_0$.
return $\boldsymbol{\eta}_0 = W\boldsymbol{\theta}(\boldsymbol{\eta})$. {Optionally, expand all η^j in terms of original variables $\mathbf{z}$.}

The configuration from the original SINDy model, i.e., the LHS $\boldsymbol{\ell}$ and the function library $\{\theta^j\}$, are used to construct a new equation model in terms of the invariants. It should be noted that the functions in the SINDy function library does not specify their input variables. For example, in the PySINDy (Kaptanoglu et al., 2022) implementation, a function θ is provided in a lambda format `lambda x, y: x * y`. Thus, θ can be applied to both the original variables, e.g. $\theta(z_1, z_2) = z_1 z_2$, and the invariant functions, e.g. $\theta(\eta_1, \eta_2) = \eta_1 \eta_2$.

C.2 Symmetry Invariant Condition as Linear Constraints

Instead of directly using the invariant functions η as the features and labels for regression, we can derive a set of linear constraints from the fact that the equation can be rewritten in terms of invariant functions. As shown in Appendix B.5, a basis Q of the constrained parameter space can be obtained from the right singular vectors of a constraint matrix M. We rearrange Q to a tensor of shape (r, q, m), where r is the dimension of the constrained parameter space, and (q, m) is the original shape of the parameter matrix W. Then, we can parameterize W by $W^{jk} = Q^{ijk}\beta^i$, where β is the learnable parameter, and discover the equation using the original SINDy objective as described in Section 2.2.

In practice, we observe that the basis Q obtained from SVD is not sparse. Indeed, SVD does not inherently encourage sparsity in the singular vectors. The lack of sparsity can pose a problem when we perform sequential thresholding in sparse regression. Specifically, in SINDy, the entries in W that are close to zero are filtered out at the end of each iteration, which serves as a proxy to the L_0 regularization. Since we fix Q and only optimize β, a straightforward modification to the sequential thresholding procedure is to threshold the entries in β instead of those in W. However, if Q is dense, even a sparse vector β can lead to a dense W, which contradicts the purpose of sparse regression.

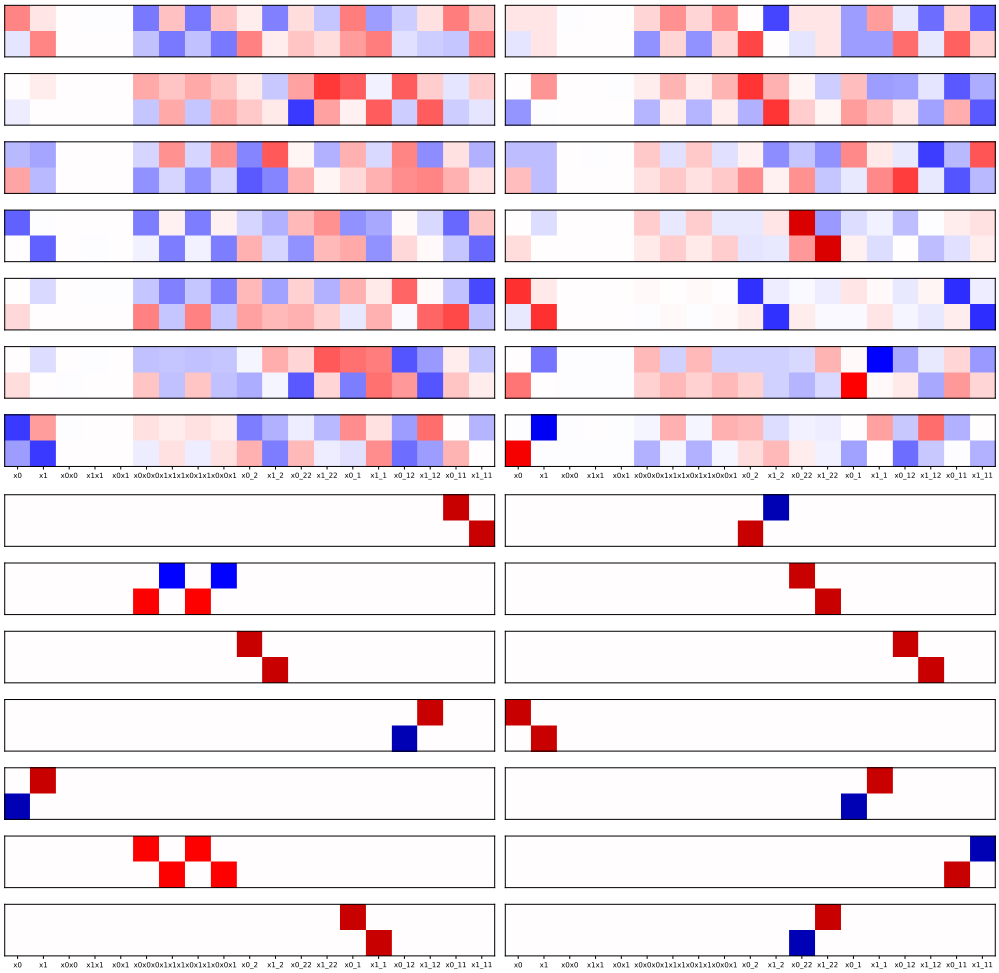

Figure 4: Basis for the SINDy parameter subspace that preserves $\mathrm{SO}(2)$ symmetry $\mathbf{v} = -v\partial_u + u\partial_v$. The SINDy parameter W has dimension 2×19. The two rows correspond to the two equations with u_t and v_t as the LHSs. The RHS contains 19 features, including all monomials of u, v up to degree 3 and their spatial derivatives up to order 2. The set of symmetry invariants used to compute the basis is given by $\{t, x, y, u^2 + v^2\} \bigcup \{\mathbf{u} \cdot \mathbf{u}_\mu\} \bigcup \{\mathbf{u}^\perp \cdot \mathbf{u}_\mu\}$, where $\mathbf{u} = (u, v)^T$ and μ is a multiindex of t, x, y with order no more than 2. The top 7×2 grid displays the original basis solved from SVD, and the bottom 7×2 grid displays the sparsified basis.

Therefore, after performing SVD, we apply a Sparse PCA to Q to obtain a sparsified basis, also of shape (r, q, m):

```
spca = SparsePCA(n_components=r)
spca.fit(Q.reshape(r, q*m))
Q_sparse = spca.components.reshape(r,q,m)
```

Figure 4 shows an example of the original basis solved from SVD (top 7×2 grid) and the sparsified basis using sparse PCA (bottom 7×2 grid). This is used in our experiment on the reaction-diffusion system (9).

C.3 Using Differential Invariants in Weak SINDy

In this subsection, we discuss the formulation of weak SINDy and how to implement our strategy of using differential invariants within the weak SINDy framework. To maintain a similar notation to the original works on weak SINDy (Messenger & Bortz, 2021a,b), we use D_{α_s} to denote partial

derivative operators, where $\alpha_s = (s_1, s_2, ..., s_p)$ is a multi-index, instead of using subscripts for partial derivatives. Thus, we no longer strictly differentiate subscripts and superscripts–both can be used for indexing lists, vectors, etc.

Given a differential equation in the form

$$D_{\alpha_0} u = \sum_{s,j} W_{sj} D_{\alpha_s} f_j(u), \tag{39}$$

we can perform integration by parts (i.e., divergence theorem) to move the derivatives from u to some analytic test function and thus bypass the need to estimate derivatives numerically. First, we multiply both sides of (39) by a test function ϕ with compact support $\mathcal{B} \subset X$ and integrate over the spacetime domain:

$$\int_X D_{\alpha_0} u(\mathbf{x}) \phi(\mathbf{x}) d\mathbf{x} = \sum_{s,j} W_{sj} \int_X D_{\alpha_s} f_j(u(\mathbf{x})) \phi(\mathbf{x}) d\mathbf{x} \tag{40}$$

WLOG, assume that $s_1 \neq 0$, and denote $\alpha_{s'} = (s_1 - 1, s_2, ..., s_p)$. Then, each term in the RHS can be integrated by parts as

$$\begin{aligned}
\int_X D_{\alpha_s} f_j(u(\mathbf{x})) \phi(\mathbf{x}) d\mathbf{x} &= \int_{\mathcal{B}} D_{\alpha_s} f_j(u(\mathbf{x})) \phi(\mathbf{x}) d\mathbf{x} \\
&= -\int_{\mathcal{B}} D_{\alpha_{s'}} f_j(u(\mathbf{x})) D_1 \phi(\mathbf{x}) d\mathbf{x} + \int_{\partial \mathcal{B}} \nu_1 D_{\alpha_{s'}} f_j(u(\mathbf{x})) \phi(\mathbf{x}) d\mathbf{x} \\
&= -\int_{\mathcal{B}} D_{\alpha_{s'}} f_j(u(\mathbf{x})) D_1 \phi(\mathbf{x}) d\mathbf{x},
\end{aligned} \tag{41}$$

where D_1 denotes the partial derivative operator w.r.t the first independent variable, and ν_1 is the first component of the unit outward normal vector.

Repeating this process until all the derivative operations move from $f_j(u)$ to the test function ϕ, we have

$$\int_X D_{\alpha_s} f_j(u(\mathbf{x})) \phi(\mathbf{x}) d\mathbf{x} = (-1)^{|\alpha_s|} \int_X f_j(u(\mathbf{x})) D_{\alpha_s} \phi(\mathbf{x}) d\mathbf{x} \tag{42}$$

Similarly for the LHS:

$$\int_X D_{\alpha_0} u(\mathbf{x}) \phi(\mathbf{x}) d\mathbf{x} = (-1)^{|\alpha_0|} \int_X u(\mathbf{x}) D_{\alpha_0} \phi(\mathbf{x}) d\mathbf{x} \tag{43}$$

The final optimization problem is to solve for $\mathbf{b} = \mathbf{G}\mathbf{w}$, where $\mathbf{w}$ is the vectorized coefficient matrix W, and each row in $\mathbf{b}$ and $\mathbf{G}$ is given by computing the integrals in (42) and (43) against a single test function. The number of rows equals the number of different test functions used.

Direct integration of symmetry via linear constraints As we have discussed in Appendix C.2, we can enforce symmetry by converting it to a set of linear constraints on the parameter W. With this approach, we can directly incorporate symmetry in weak SINDy. Specifically, we just parameterize W as in terms of a precomputed basis Q and a trainable vector β and directly substitute this parameterization of W into the optimization problem of weak SINDy. We adopt this strategy in our experiments concerning weak SINDy.

Expressing the equations with differential invariants The above approach is only possible when the conditions in Proposition 3.4 about the selected set of symmetry invariants hold. We should note that it is not always possible to find a set of invariants so that the symmetry condition can be converted to linear constraints on the parameter W via the procedure in the proof of Proposition 3.4. One may ask the following question: can we simply express the equations in terms of differential invariants and apply weak SINDy, similar to Algorithm 2 for the original SINDy formulation? Here, we do not provide a definite conclusion for this question, but only discuss several cases where directly using differential invariants in equations might succeed or fail in weak SINDy.

To adapt to the weak SINDy formulation (39), it is more helpful to consider the symmetry invariants as generated by some fundamental invariants and some invariant differential operators, instead of

specifying a complete set of differential invariants for every order. Concretely, there exists a set of invariant differential operators $\{\mathcal{O}_j\}$ and a set of fundamental differential invariants $\mathbf{I} = \{\eta_k\}$ s.t. every differential invariant can be written as $\mathcal{O}_{j_1}...\mathcal{O}_{j_n}\eta_k$. For the $\mathrm{SO}(2)$ symmetry group in Example B.5, one possible choice is

$$\eta_1 = \frac{1}{2}(x^2 + y^2),\ \eta_2 = u,\ \mathcal{O}_1 = xD_y - yD_x,\ \mathcal{O}_2 = xD_x + yD_y. \tag{44}$$

We can compose these generating invariant operators to obtain a full library of eligible differential operators up to some order, denoted $\mathbf{D} = \{\mathcal{D}_j\}$. The exact compositions can vary and we can choose the most convenient one for subsequent calculations. For the above $\mathrm{SO}(2)$ example, for up to second-order differential operators, we can choose $\{\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_1^2, \mathcal{O}_2^2, \frac{2}{\eta_1}(\mathcal{O}_1^2 + \mathcal{O}_2^2)\}$. Note the last operator is exactly the Laplacian.

Then, the complete set of eligible terms (respecting the symmetry) in the equation is $\{\mathcal{D}_j\eta_k : \mathcal{D}_j \in \mathbf{D}, \eta_k \in \mathbf{I}\}$. If we assume, as in SINDy, that the governing equation can be written in linear combination of these symmetry invariants, then we can assign a weight for each $\mathcal{D}_j\eta_k$ and form a coefficient matrix $W = [W_{jk}]$. That is,

$$\mathcal{D}_{j_0}\eta_{k_0} = \sum_{(j,k)\neq(j_0,k_0)} W_{jk}\mathcal{D}_j\eta_k. \tag{45}$$

Then, multiplying each side by a test function $\phi(\mathbf{x})$, we have

$$\int_X \mathcal{D}_{j_0}\eta_{k_0}\phi(\mathbf{x})d\mathbf{x} = \sum_{(j,k)\neq(j_0,k_0)} W_{jk}\int_X \mathcal{D}_j\eta_k\phi(\mathbf{x})d\mathbf{x}. \tag{46}$$

The question then boils down to whether we can apply the technique of integration by parts similarly to this set of differential operators and differential functions, since the original algorithm only deals with partial derivative operators D_{α_s} and ordinary functions $f_j(u)$.

To check this, let us explicitly write out the dependency of these operators and fundamental invariants.

Case 1 A relatively simple case is when all invariant operators take the form $\mathcal{D}_j = \sum_s a_s(\mathbf{x})D_{\alpha_s}$ and $\eta_k = \eta_k(\mathbf{x}, u(\mathbf{x}))$. Each term in the RHS of (46) can be expanded as

$$\begin{aligned}\int_X \mathcal{D}_j\eta_k\phi(\mathbf{x})d\mathbf{x} &= \sum_s \int_X a_s(\mathbf{x})D_{\alpha_s}\eta_k(\mathbf{x}, u(\mathbf{x}))\phi(\mathbf{x})d\mathbf{x} \\ &= \sum_s (-1)^{|\alpha_s|}\int_X \eta_k(\mathbf{x}, u(\mathbf{x}))D_{\alpha_s}[a_s(\mathbf{x})\phi(\mathbf{x})]d\mathbf{x}\end{aligned} \tag{47}$$

Evaluating (47) does not require estimating partial derivatives of u. Therefore, weak SINDy can be applied to this case quite straightforwardly.

Case 2 However, it is not always possible to have all $\mathcal{D}_j$ as classical linear differential operators and all η_k as ordinary functions. For instance, in Example B.6, there are no ordinary symmetry invariants due to the constraint of translation symmetry.

If we still have linear operators $\mathcal{D}_j = \sum_s a_s(\mathbf{x})D_{\alpha_s}$, but on the other hand we have differential functions $\eta_k = \eta_k(\mathbf{x}, u^{(n)})$, we can still perform integration by parts as in (47), but the final result becomes

$$\sum_s (-1)^{|\alpha_s|}\int_X \eta_k(\mathbf{x}, u^{(n)})D_{\alpha_s}[a_s(\mathbf{x})\phi(\mathbf{x})]d\mathbf{x}, \tag{48}$$

meaning we still have to evaluate whatever partial derivatives remain in η_k. It is possible that we can decrease the order of partial derivatives compared to vanilla sparse regression, but we cannot eliminate all partial derivatives compared to Weak SINDy without any symmetry information.

Case 3 The most challenging case is when the invariant differential operators explicitly involve the partial derivative, such as $\mathcal{D}_j = \sum_s a_s(\mathbf{x}, u^{(n)}) D_{\alpha_s}$. Then, similar to (48), integration by parts yields:

$$\sum_s (-1)^{|\alpha_s|} \int_X \eta_k(\mathbf{x}, u^{(n)}) D_{\alpha_s}[a_s(\mathbf{x}, u^{(n)})\phi(\mathbf{x})] d\mathbf{x}. \tag{49}$$

In this case, we still need to compute the partial derivatives, not only those in η_k, but also those arising from a_s and $D_{\alpha_s}(a_s)$. The latter might involve higher-order derivatives and the benefit of using the weak formulation may further diminish.

D Additional Experiment Results

Contents of this section include:

- Appendix D.1: Results for some variants of the sparse regression models considered in Table 1.
- Appendix D.2: Results for genetic programming-based algorithms under different computational budgets.
- Appendix D.3: Samples of equations discovered by different methods.
- Appendix D.4: Visualized prediction errors of equations discovered by different methods.

D.1 Variant Sparse Regression Models

Table 3: Results of sparse regression models on the Boussinesq equation and the reaction-diffusion system. $\mathcal{C}$ stands for complexity, i.e., the dimensionality of the parameter space. SP stands for success probability. The PySINDy and SI rows present the same results as the corresponding rows in Table 1.

Method	Boussinesq (7)		Reaction-diffusion (9)	
	$\mathcal{C}\downarrow$	SP $\uparrow$	$\mathcal{C}\downarrow$	SP $\uparrow$
PySINDy	15	0.00	38	0.53
PySINDy*	21	**1.00**	468	0.00
SI	**13**	**1.00**	28	0.54
SI-aligned	-	-	**14**	**0.56**

The original implementation of PySINDy (de Silva et al., 2020; Kaptanoglu et al., 2022) does not allow functions to be applied to partial derivative terms. As a result, terms such as u_x^2 cannot be modeled. This leads to its failure to discover the Boussinesq equation (7), as we have shown in Table 1.

We modify the implementation and include an additional set of results with different libraries, denoted as PySINDy* in Table 3. The PySINDy* model supports a wider range of library functions, including functions of partial derivatives, e.g., u_x^2. A complete description of the hypothesis spaces of different sparse regression-based methods is available in Appendix E.5.

As Table 3 shows, PySINDy* succeeds in the Boussinesq equation. However, it fails in the reaction-diffusion system because its parameter space becomes too large due to a higher-dimensional total space $X \times U \simeq \mathbb{R}^2 \times \mathbb{R}^2$. This augments the point that SINDy's success relies on an appropriate choice of function library. If the library is too small to contain all the terms appearing in the equation of interest, the discovery is sure to fail. If the library is too large, the optimization problem becomes more difficult in the high-dimensional parameter space. On the other hand, by introducing the inductive bias of symmetry, our method automatically identifies a proper function library that contains all the necessary terms for a PDE with a specific symmetry group, but not other redundant terms.

We include another model in Table 3, SI-aligned, where we derive a set of linear constraints on the sparse regression parameters from the fact that the equations can be expressed in terms of symmetry invariants. In this way, we still optimize the original parameters (though in a constrained subspace) as in the base SINDy model without symmetry, effectively "aligning" the hypotheses about equations

from symmetry and the base SINDy model. This method is discussed in detail in Section 3.3 and Appendix C.2. We should also note that this method is mainly developed for incorporating the symmetry constraints into the weak formulation of SINDy. However, it is perfectly acceptable to implement it in the original formulation of SINDy, so we provide its results in Table 3 for reference.

For the reaction-diffusion system, SI-aligned has a 14-dimensional parameter space. The basis for its parameter space is visualized in Figure 4. It achieves a slightly higher success probability than SI (regression with symmetry invariants) and PySINDy (without symmetry information). We do not apply SI-aligned to the Boussinesq equation, because it is not necessary to align the hypotheses from SINDy and symmetry in this case. We can readily convert any equation discovered from SI (regression with symmetry invariants) by multiplying both sides by u_x^2.

We note that the results on the reaction-diffusion system in Table 3 are for models with the original SINDy formulation, in contrast to the weak SINDy formulation used in Figure 3. Therefore, the results in Figure 3 should not be directly compared to those in Table 1 and Table 3.

D.2 Genetic Programming

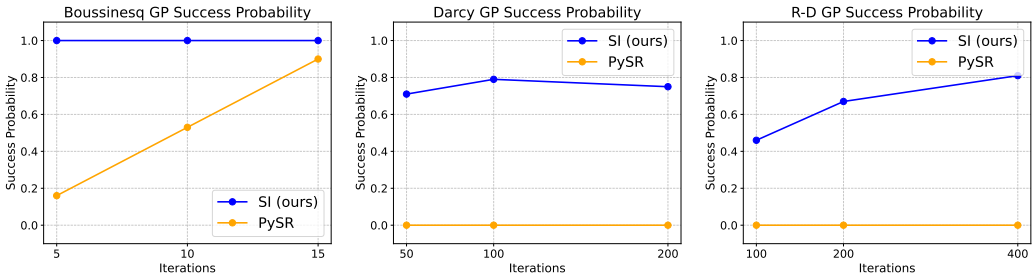

Figure 5: Success Probabilities of GP-based methods on different systems. Our method with symmetry invariants can discover the correct equations with fewer iterations.

For each system in Section 4.1, we run the genetic programming discovery algorithm with three different iteration counts, but otherwise keep all hyperparameters constant. In Figure 5, we plot the success probability as a function of the iteration count for both the base GP algorithm and our method that uses symmetry invariants.

In all cases, we find that using symmetry invariants results in a higher success probability in comparison to unmodified PySR. Specifically, for the Boussinesq equation, our method achieves a 100% chance of discovery with 5 iterations, whereas even with 3 times the number of iterations, PySR only yields a 90% success probability. This highlights that using invariants improves the efficiency of equation discovery. For Darcy flow and Reaction-Diffusion, we find that the base genetic programming algorithm fails to ever make a correct prediction. On the other hand, using symmetry invariants leads to a successful discovery the majority of the time.

We finally note that increasing the number of iterations to 200 for Darcy flow slightly lowers the success probability when using symmetry invariants. We hypothesize this is because at higher iterations, the search process begins to overfit and introduces extraneous low-order terms. While we already drop some terms with small enough coefficients, future works may consider a more refined filtration process.

D.3 Samples of Discovered Equations

In Table 4, we list some randomly selected equations discovered by different methods for the Boussinesq equation (7). Some methods almost consistently discover correct/incorrect equations (i.e., have success probabilities close to 1 or 0), so we only select one sample for each. For other methods with a large variance in the discovered equations, we display two samples: a correct equation and an incorrect one.

The ground truth equation in the original variables is given in (7). The ground truth equation in the symmetry invariants is given by

$$\eta_{(0,2)} + \eta_{(0,0)}\eta_{(2,0)} + \eta_{(4,0)} + 1 = 0 \tag{50}$$

Table 4: Samples of discovered equations from the observed solution of the Boussinesq equation (7). For GP-based methods, we include results from different numbers of iterations (indicated by "N its"). For transformer-based methods, we include two samples for each method because of the large variance of discovered equations from different runs.

Method		Equation sample(s)
Sparse regression	PySINDy	$u_{tt} = -1.01u_{xxxx} - 0.79uu_{xx}$
	PySINDy*	$u_{tt} = -1.01u_{xxxx} - 0.99u_x^2 - 0.98uu_{xx}$
	SI	$\eta_{(0,2)} = -1.00 - 1.00\eta_{(4,0)} - 1.00\eta_{(0,0)}\eta_{(2,0)}$
Genetic programming	PySR (5 its)	$uu_{xx} + 1.00u_{tt} + u_{xxxx} = 0$
	PySR (15 its)	$uu_{xx} + u_{tt} + u_x^2 + 1.00u_{xxxx} = 0$
	SI (5 its)	$1.00\eta_{(0,0)}\eta_{(2,0)} + 1.00\eta_{(0,2)} + 1.00\eta_{(4,0)} + 1 = 0$
Transformer	E2E	(1) $u_{tt} = -1.13uu_{xx} - 0.98u_{xxxx} - 0.30\lvert u_x\rvert$
		(2) $u_{tt} = -0.85uu_{xx} - 0.75u_x^2 - 0.99u_{xxxx}$
	SI	(1) $\eta_{(0,2)} = -1.05\eta_{(0,0)}\eta_{(2,0)} - 1.00\eta_{(4,0)} - 0.96$
		(2) $\eta_{(0,2)} = -0.81\eta_{(0,0)}\eta_{(2,0)} - 0.40\eta_{(0,0)} - 0.98\eta_{(4,0)} - 0.90$

Table 5 lists the equation samples discovered from the Darcy flow dataset. The ground truth equation in original variables is given in (8), and the ground truth equation in symmetry invariants is given by

$$8\zeta_2 - \Delta u - e^{4R^2} = 0, \tag{51}$$

where $\zeta_2 = xu_x + yu_y$, $\Delta u = u_{xx} + u_{yy}$, and $R^2 = x^2 + y^2$ are among the rotational invariants used in symbolic regression.

Table 5: Samples of discovered equations for the Darcy flow dataset.

Method		Equation sample
Genetic programming	PySR	$u - 0.47x^2y^2 - 0.38e^{0.09(u_{xx}+u_{yy})} + 0.20 = 0$
	SI	$\zeta_2 - 0.13\Delta u - 0.13e^{4.01R^2} = 0$
Transformer	E2E	$u_{xx} = -7.43\sqrt{u^2 + 0.65u_x^2}$
	SI	$\Delta u = -2.56u + 0.85\zeta_2 + 0.29$

Finally, Table 6 lists the equation samples discovered from the reaction-diffusion dataset. The ground truth equation in original variables is given in (9) with $d_1 = d_2 = 0.1$, and the ground truth equation in symmetry invariants is given by

$$\begin{aligned} I_t &= 0.1(I_{xx} + I_{yy}) + A(1 - A) \\ E_t &= 0.1(E_{xx} + E_{yy}) - A^2 \end{aligned} \tag{52}$$

where $I_\mu = uu_\mu + vv_\mu$ and $E_\mu = -vu_\mu + uv_\mu$ for any multiindex μ of t, x, y, and $A = u^2 + v^2$.

D.4 Prediction Errors of Discovered Equations

In Table 1, we report the prediction errors of the discovered equations on the three PDE systems. Specifically, for the Boussinesq equation and the reaction-diffusion system, we simulate the discovered PDE from an initial condition for a certain time period, e.g., $t \in [0, 20]$ for the Boussinesq equation and $t \in [0, 10]$ for the reaction-diffusion system. Then, we compare the numerical solution with the ground truth solution from the same initial condition at the end of the time period.

Table 6: Samples of discovered equations for the reaction-diffusion system dataset. Each discovered result contains two equations, since this is an evolution system with two dependent variables u, v.

Method		Equation sample
Sparse regression	PySINDy	$\begin{cases} u_t = 0.96u - 0.97u^3 + 1.00^3 - 0.97uv^2 + 1.00u^2v + 0.09u_{xx} + 0.09u_{yy} \\ v_t = 0.96v - 1.00u^3 - 0.97v^3 - 1.00uv^2 - 0.96u^2v + 0.09v_{xx} + 0.09v_{yy} \end{cases}$
	PySINDy*	$\begin{cases} u_t = 0.21u - 0.24u^3 + 1.00v^3 - 0.23uv^2 + 0.99u^2v \\ v_t = 0.21v - 1.01u^3 - 0.24v^3 - 0.99uv^2 - 0.23u^2v \end{cases}$
	SI	$\begin{cases} I_t = 0.10I_{xx} + 0.10I_{yy} + 0.96A - 0.96A^2 \\ E_t = 0.10E_{xx} + 0.10E_{yy} - 1.00A^2 \end{cases}$
	SI-aligned	$\begin{cases} u_t = 0.95u - 0.96u^3 + 1.00v^3 - 0.96uv^2 + 1.00u^2v + 0.09u_{xx} + 0.09u_{yy} \\ v_t = 0.95v - 1.00u^3 - 0.96v^3 - 1.00uv^2 - 0.96u^2v + 0.09v_{xx} + 0.09v_{yy} \end{cases}$
Genetic programming	PySR	$\begin{cases} u_t = 0.92v \\ v_t = -0.92u \end{cases}$
	SI	$\begin{cases} I_t = 0.10I_{xx} + 0.10I_{yy} + A - 1.00A^2 \\ E_t = 0.10E_{xx} + 0.10E_{yy} - 1.00A^2 \end{cases}$
Transformer	E2E	$\begin{cases} u_t = 0.89u_y \\ v_t = -0.91u \end{cases}$
	SI	$\begin{cases} I_t = 0 \\ E_t = 0.50 \arctan(0.45E_y - 0.31E_y/(-540.12AE_y + ...) + ...) + ... \end{cases}$

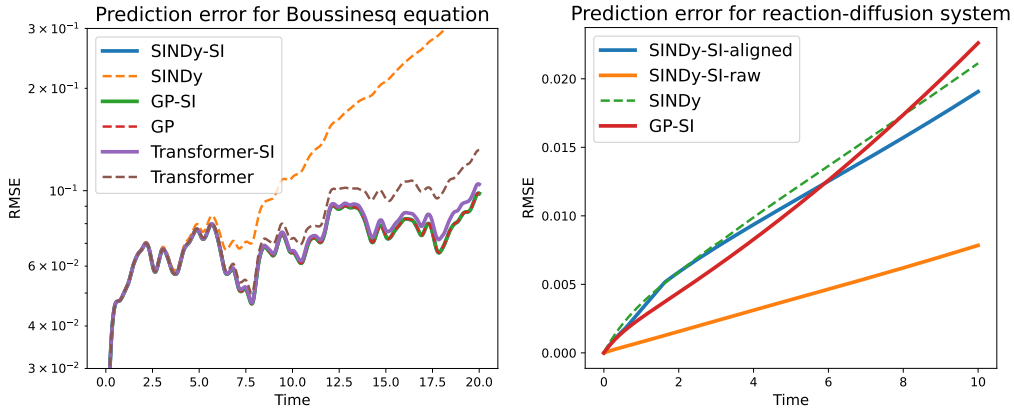

Figure 6: Prediction error over time using the discovered equations.

In addition to the prediction error at the end of the simulation time, Figure 6 shows the errors at each simulation timestep. We do not include methods whose error curves grow too fast due to the incorrectly identified equations. The results in Figure 6 are consistent with those in Table 1. Generally, the discovered equations with smaller prediction errors at the end of the simulation time also have lower prediction errors throughout the entire time interval.

For Darcy flow (8), since it describes the steady state of a system and does not involve time derivatives, we do not simulate the discovered PDEs. Instead, we evaluate each discovered PDE $F(\mathbf{x}, u^{(n)}) = 0$ on the test dataset $\{(\mathbf{x}, u^{(n)}) : \mathbf{x} \in \Omega\}$ and report the residual as the prediction error. In addition to the average error over all the spatial grid points reported in Table 1, we visualize the error heatmaps over the grid in Figure 7. It can be observed that the discovered equations with symmetry invariants have lower errors across the entire grid.

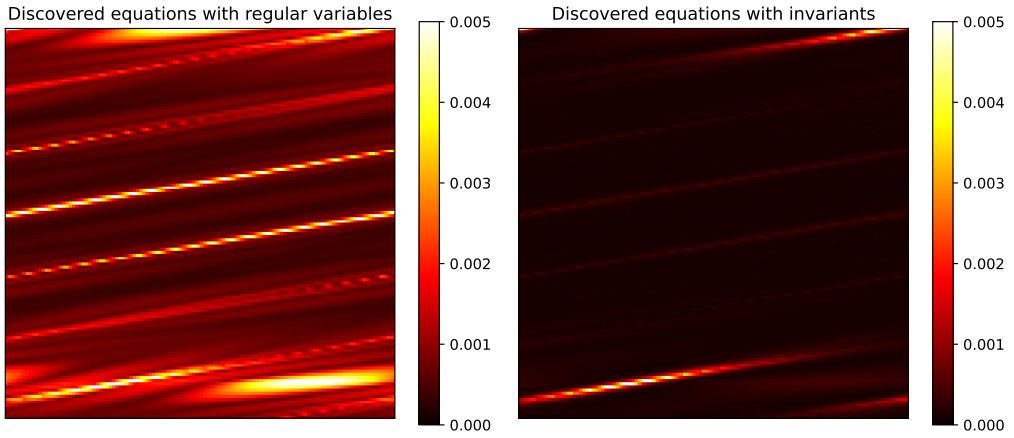

Figure 7: Prediction error of discovered equations from genetic programming methods for Darcy flow. Left: genetic programming with regular variables. Right: genetic programming with symmetry invariants.

E Experiment Details

In this section, we describe the experiment setups required to reproduce the experiments. In terms of computational resources, our experiments are conducted with 12 INTEL(R) XEON(R) PLATINUM 8558 CPUs and should be reproducible within minutes with any modern CPUs.

E.1 Data generation

Boussinesq equation The equation is solved using a Fourier pseudospectral method for spatial derivatives and a fourth-order Runge-Kutta (RK4) scheme for time integration. The solution is computed on a periodic spatial domain $[-L, L]$ with $N = 256$ grid points. The equation is reformulated as a first-order system in time by introducing $v = u_t$, and both u and v are evolved in time. Spatial derivatives are computed using the Fast Fourier Transform, and time derivatives of u up to the fourth order are derived analytically from the governing equation. At each time step, values of u are recorded in the dataset for equation discovery. The simulation starts from an initial condition of $u(x) = 0.5e^{-x^2}$ and $u_t = 0$ and proceeds up to a final time $T = 20$ with a time step of $\Delta t = 0.001$. Starting from the solution at $T = 20$, we simulate for another $T' = 20$ with the same configuration to obtain a test dataset for evaluating prediction errors of the discovered equations.

Darcy flow We use the data generation code[2] from PDEBench (Takamoto et al., 2022) to generate the steady-state solution of Darcy flow over a unit square. The solution is obtained by numerically solving a temporal evolution equation

$$u_t(\mathbf{x}, t) - \nabla(a(\mathbf{x})\nabla u(\mathbf{x}, t)) = f(x), \mathbf{x} \in (-0.5, 0.5)^2, \tag{53}$$

with $a(\mathbf{x}) = e^{-4\|\mathbf{x}\|_2^2}$ and $f(\mathbf{x}) = 1$.

Reaction-diffusion We use the data generation code[3] from PySINDy (de Silva et al., 2020; Kaptanoglu et al., 2022). The spatial domain is $[-10, 10] \times [-10, 10]$ with 128 grid points in each direction. The simulation proceeds up to a final time $T = 10$ with a time step $\Delta t = 0.05$. We perturb the numerical solution by a 0.05% noise and record the values of u, v to the dataset for equation discovery. Starting from the solution at $T = 10$, we simulate for another $T' = 10$ with the same configuration to obtain a test dataset for evaluating prediction errors of the discovered equations.

[2]https://github.com/pdebench/PDEBench/tree/main/pdebench/data_gen/data_gen_NLE/ReactionDiffusionEq
[3]https://github.com/dynamicslab/pysindy/blob/master/examples/10_PDEFIND_examples.ipynb

E.2 Sparse regression

Boussinesq equation For SINDy with original variables, we fix u_{tt} as the LHS of the equation and include functions of up to 4th-order derivatives on the RHS. For PySINDy in Table 1, the library contains monomials on $U^{(4)}$ with degree in u no larger than 2 and degree in any partial derivative terms u_α no larger than 1. For example, u^2u_x is included, but u^3, u_x^2 are not. For PySINDy*, the library contains all monomials on $U^{(4)}$ up to degree 2. For example, u_x^2 and uu_x are included. Note that the PySINDy* library does not contain all functions in the original PySINDy library, e.g., u^2u_x is not included because it has degree 3.

Our method, SI, uses the invariant set in Example B.6 for sparse regression. Specifically, $\eta_{(0,2)} = u_{tt}/u_x^2$ is used as the LHS of the equation, and the rest of the invariants are included in the RHS. The function library contains all monomials of these RHS invariants up to degree 2. Also, since the invariants contain rational functions with u_x on the denominator, we remove the data points with small $|u_x|$ to avoid numerical issues.

For all methods, we flatten the data on the spatiotemporal grid and randomly sample 2% of the data for each run. The data filtering process in SI-raw is performed after subsampling. The threshold value for sequential thresholding is set to 0.25, and the coefficient for L_2 regularization is set to 0.05.

Darcy flow Sparse regression-based methods are not directly applicable to Darcy flow (8) because there exist terms such as $e^{-4(x^2+y^2)}$. While it is still possible to include all necessary terms in the function library so that the equation can be written in the linear combination form (4), the knowledge of these complicated terms is nontrivial and should not be assumed available before running the equation discovery algorithm.

Reaction-Diffusion For SINDy with original variables, We fix u_t and v_t as the LHS of the equation and include functions of up to 2nd-order spatial derivatives on the RHS. In PySINDy, the library contains monomials of u, v up to degree 3 and all spatial derivatives up to order 2. In PySINDy*, the library contains all monomials of u, v and their up to second-order spatial derivatives up to degree 3.

Our method uses the invariant set $\{t, x, y, u^2 + v^2\} \bigcup \{\mathbf{u} \cdot \mathbf{u}_\mu\} \bigcup \{\mathbf{u}^\perp \cdot \mathbf{u}_\mu\}$, where $\mathbf{u} = (u, v)^T$ and μ is a multiindex of t, x, y. We will denote $I_\mu = \mathbf{u} \cdot \mathbf{u}_\mu$ and $E_\mu = \mathbf{u}^\perp \cdot \mathbf{u}_\mu$. We use I_t and E_t as the LHS of the equation, and the rest of the invariants are included in the RHS. The function library contains all monomials of these RHS invariants up to degree 2.

We randomly sample 10% of the data for each run. The threshold value for sequential thresholding is set to 0.05. The coefficient for L_2 regularization is set to 0 for SINDy with original variables and 0.1 for our method with symmetry invariants.

For the experiments with different levels of noise (Section 4.4), we use weak SINDy as the base algorithm. The function library is the same as PySINDy as described above. To enforce symmetry, instead of directly using the symmetry invariants, we derive a set of linear constraints on the sparse regression parameters to adapt to weak SINDy. This procedure is further described in Appendix C.3.

E.3 Genetic Programming

In all experiments, to determine if an equation matches the ground truth we first expand the prediction into a sum of monomial terms. We then eliminate all terms whose relative coefficient is below 0.01. For each term in the filtered expression, we see if it matches any term in the ground truth expression. This is done by randomly sampling 100 points from the standard normal distribution and evaluating both the prediction and candidate ground truth term on the generated points. Note that we drop the coefficients before evaluation. If all evaluations of the predicted term have a relative error of less than 5% from those of the ground truth, the terms are said to match. If there is a perfect matching between the terms in the ground truth and prediction, the prediction is listed as correct.

Rather than directly returning a single equation, PySR finally produces a hall-of-fame that consists of multiple candidate solutions with varying complexities. To finally pick a single prediction, we use a selection strategy equivalent to the "best" option from PySR.

Boussinesq equation For the Boussinesq equation (7), we first randomly subsample 10000 datapoints. We configure PySR to use the addition and multiplication operators, to have 127 populations of size 27, and to have the default fraction-replaced coefficient of 0.00036.

When running with ordinary variables, we sequentially try fixing the LHS to each variable in $(\mathbf{x}, u^{(4)})$ and allow the RHS to be a function of all remaining variables. Similarly, runs using invariants sequentially fix the LHS from the set given by Example B.6 and the RHS as a function of all other invariants.

For each iteration count of 5, 10, and 15, we run the algorithm using invariant or ordinary variables and report the number of correct predictions out of 100 trials.

Darcy flow In the Darcy experiment (8), we eliminate all points that are within 3 pixels from the border and then randomly subsample 10000 datapoints. We configure PySR to use the addition, multiplication, and exponential operators; to have 127 populations of size 64; and to have a fraction-replaced coefficient of 0.1. We further constrain it to disallow nested exponentials (e.g. $\exp(\exp(x) + 4)$.

We try all possible ordinary variables in $(\mathbf{x}, u^{(2)})$ for the LHS and the RHS is then a function of the unused variables. Likewise when using invariants, we fix the LHS to each possible invariant specified in Example B.5 and set the RHS as a function of the remaining invariants.

For each iteration count of 50, 100, and 200, we run the algorithm using invariant or ordinary variables and report the number of correct predictions out of 100 trials.

Reaction-Diffusion For the Reaction Diffusion equation (9), we remove all points that are within 3 pixels from the border or have timestamp greater than or equal to 40, and then randomly subsample 10000 datapoints. We configure PySR to use the addition and multiplication operators, to have 127 populations of size 64, and to have a fraction-replaced coefficient of 0.5.

In the ordinary variable case, we fix the LHS as either u_{tt} or v_{tt} and allow the RHS to be a function of all other variables in $(\mathbf{x}, u^{(2)})$. When using invariants, the LHS is fixed to be either I_t or E_t and the RHS is then a function of all remaining invariants.

For each iteration count of 100, 200, and 400, we run the algorithm using regular and ordinary variables and report the number of correct predictions out of 100 trials.

E.4 Symbolic Transformer

We use the pretrained symbolic transformer model provided in the official codebase[4] from Kamienny et al. (2022). The transformer-based symbolic regressor is initialized with 200 maximal input points and 100 expression trees to refine. The variable sets used in the symbolic transformer are the same as those described in the genetic programming experiments, except for the Boussinesq equation, where we remove all mixed derivative terms in both the original variable set and the symmetry invariant set. We find that the symbolic transformer can sometimes discover the correct equation under this further simplified setup, but fails when using the larger variable sets.

We also fix the LHS of the function and use the remaining variables as RHS features. For the Boussinesq equation, the LHS is fixed to u_{tt} for original variables and $\eta_{(0,2)}$ for symmetry invariants. For the Darcy flow, the LHS is fixed to u_{xx} for original variables and Δu for symmetry invariants. For the reaction-diffusion system, the LHS is fixed to u_t, v_t for original variables and I_t, E_t for symmetry invariants.

E.5 Hypothesis Spaces of Equation Discovery Algorithms

Table 7 and Table 8 describe the hypothesis spaces of different equation discovery algorithms when applied to the Boussinesq equation and the reaction-diffusion system.

[4] https://github.com/facebookresearch/symbolicregression/blob/main/Example.ipynb

Table 7: Hypothesis spaces of different equation discovery algorithms for the Boussinesq equation.

Method		Hypothesis space
Sparse Regression	PySINDy	$u_{tt} = W\boldsymbol{\theta}(u^{(4)}), \{\theta^j\} = \{ab : a \in \mathrm{Mono}_{\leq 2}(U), b \in \{1, u_x, ..., u_{xxxx}\}\}$
	PySINDy*	$u_{tt} = P(u^{(4)}) \in \mathrm{Poly}_{\leq 2}(U^{(4)})$
	SI	$\eta_{(0,2)} = P(\boldsymbol{\eta}) \in \mathrm{Poly}_{\leq 2}(\{\eta_{(\alpha,\beta)}\} \backslash \{\eta_{(0,2)}\})$
Genetic Programming	PySR	$z^j = f(\mathbf{z}^{-j})$ for $\mathbf{z} = (\mathbf{x}, u^{(4)})$ and some j
	SI	$\eta_{(\alpha_0,\beta_0)} = f(\boldsymbol{\eta}_{-(\alpha_0,\beta_0)})$ for $\boldsymbol{\eta} = \{\eta_{(\alpha,\beta)} : \alpha + \beta \leq 4\}$ and some (α_0, β_0)

Table 8: Hypothesis spaces of different equation discovery algorithms for 2D reaction-diffusion. $\mathbf{u}^{(n)} \in U^{(n)}$ denotes the collection of all up to nth order *spatial* derivatives. $\alpha = [\alpha_1, \alpha_2]$ is the multiindex for spatial variables. $\mathbf{x} = (x, y, t)$. $A = u^2 + v^2$.

Method		Hypothesis space
Sparse Regression	PySINDy	$\mathbf{u}_t = W\boldsymbol{\theta}(\mathbf{u}^{(2)}), \{\theta^j\} = \mathrm{Mono}_{\leq 3}(U) \bigcup \{\mathbf{u}_\alpha : \lvert\alpha\rvert \leq 2\}$
	PySINDy*	$\mathbf{u}_t = P(\mathbf{u}^{(2)}) \in \mathrm{Poly}_{\leq 3}(U^{(2)})$
	SI	$[I_t, E_t]^T = P \in \mathrm{Poly}_{\leq 2}(A, \mathbf{x}, I_\alpha, E_\alpha; \lvert\alpha\rvert \leq 2)$
	SI-aligned	$\mathbf{u}_t = W\boldsymbol{\theta}(\mathbf{u}^{(2)}), W^{jk} = Q^{\widetilde{ijk}}\beta^i$ for some precomputed Q
Genetic Programming	PySR	$\mathbf{u}_t = \mathbf{f}(\mathbf{x}, \mathbf{u}^{(2)})$
	SI	$[I_t, E_t]^T = \mathbf{f}(A, \mathbf{x}, I_\alpha, E_\alpha; \lvert\alpha\rvert \leq 2)$

F Broader Impacts

The method in this paper can potentially be used to expedite the process of discovering governing equations from data and aid researchers in other scientific domains. Equally important, equations inferred from imperfect or biased data may appear authoritative yet embed systematic errors. Thorough validation checks, uncertainty quantification, and domain-expert review protocols for the discovered equations are essential.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction clearly state the claims made, with a contribution list at the end of the introduction.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: See Section 5.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: Full assumptions and proofs are provided in Appendix B.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: We provide code and instructions to run the experiments in the supplementary material.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide code and instructions to run the experiments in the supplementary material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: See Appendix E and also supplementary materials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: In the main experiment in Table 1, we run each experiment for 100 times with different random seeds. The reported success probability is itself a random variable defined based on all runs, so error bar is not applicable. For the prediction error, we report the median across all runs instead of the mean and standard deviation, because there are outliers in the prediction errors arising from incorrectly identified equations, making the mean and standard deviation less meaningful.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: See Appendix E.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The authors have reviewed the NeurIPS Code of Ethics and confirmed that the research in this paper conforms with the Code.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Discussed at the end of the Appendix.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The original papers/licenses are properly cited/included.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: We will provide the codebase for our experiments in the supplementary material along with documentation.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.