
Learning Infinitesimal Generators of Continuous Symmetries from Data

Gyeonghoon Ko, Hyunsu Kim, Juho Lee
Kim Jaechul Graduate School of AI
KAIST
Seoul, South Korea
{kog, kim.hyunsu, juholee}@kaist.ac.kr

Abstract

Exploiting symmetry inherent in data can significantly improve the sample efficiency of a learning procedure and the generalization of learned models. When data clearly reveals underlying symmetry, leveraging this symmetry can naturally inform the design of model architectures or learning strategies. Yet, in numerous real-world scenarios, identifying the specific symmetry within a given data distribution often proves ambiguous. To tackle this, some existing works learn symmetry in a data-driven manner, parameterizing and learning expected symmetry through data. However, these methods often rely on explicit knowledge, such as pre-defined Lie groups, which are typically restricted to linear or affine transformations. In this paper, we propose a novel symmetry learning algorithm based on transformations defined with one-parameter groups, continuously parameterized transformations flowing along the directions of vector fields called infinitesimal generators. Our method is built upon minimal inductive biases, encompassing not only commonly utilized symmetries rooted in Lie groups but also extending to symmetries derived from nonlinear generators. To learn these symmetries, we introduce a notion of a validity score that examine whether the transformed data is still valid for the given task. The validity score is designed to be fully differentiable and easily computable, enabling effective searches for transformations that achieve symmetries innate to the data. We apply our method mainly in two domains: image data and partial differential equations, and demonstrate its advantages. Our codes are available at <https://github.com/kogyehoon/learning-symmetry-from-scratch.git>.

1 Introduction

Symmetry is fundamental in many scientific disciplines, crucial for understanding the structure and dynamics of physical systems, datasets, and mathematical models. The ability to uncover and leverage symmetries has become increasingly important in machine learning and scientific research due to its potential to improve model efficiency, generalization, and interpretability. By capturing inherent symmetrical properties, models can learn more compact and informative representations, leading to improved performance in tasks like supervised learning [31, 29, 4, 7, 33], self-supervised learning [8, 15, 23], and generative models [19, 11, 18].

Previous methods for learning symmetry have often relied on the explicit parameterization of group representations based on predefined generators, which can be limited in capturing various symmetries, including transformations that do not align along the generators. For example, when searching for Lie group symmetries in images or physics data, existing methods [3, 34] parameterize a group action g as the matrix exponential of a linear combination of linear or affine Lie algebra generators L_i with their learnable coefficients w_i as $g = \exp(\sum_i w_i L_i)$. In the affine transformations of images in

(x_1, x_2) -coordinates, there are six generators, each corresponding to translation, scaling, and shearing operations with respect to the x_1 -axis and x_2 -axis. Although there exist some methods that directly learn the generators, they are either bound to the general linear group $GL(n)$, which cannot account for non-affine or non-linear transformations [24], or are not guaranteed to find the correct symmetry in real-world image datasets [9, 13].

When searching for symmetries in high-dimensional real-world datasets, we can take advantage of the fact that the data can be interpreted as a function $f : \mathcal{X} \rightarrow \mathcal{Y}$, such as images, which are functions from the 2D Euclidean space to the color space. Another notable example of such data is partial differential equations (PDEs), where the data take the form $u : \mathcal{X} \rightarrow \mathcal{U}$ and the Lie symmetries are defined as transformations on the space $\mathcal{X} \times \mathcal{U}$. There have been significant advances in Lie symmetry analysis in recent years, for both academic and industrial purposes, mostly involving extensive symbolic calculations and relying on computer algebra systems [26]. Discovering Lie symmetries of PDEs from data without prior knowledge is an unexplored topic, except for the work of Gabel et al. [14], which learns the symmetry generators of various PDEs in a supervised learning setup.

In this work, we propose a novel method for learning continuous symmetries, including non-affine transformations, from data without prior knowledge. By modeling one-parameter groups using Neural Ordinary Differential Equation (Neural ODE) [6], we establish a learnable infinitesimal generator capable of producing a sequence of transformed data through ODE integration. We design an appropriate *validity score* function that measures how much the transformation violates the invariance to certain criteria defined depending on the target task, and learn the generators by optimizing towards the validity score of the data transformed through ODE integration. For example, in an image classification dataset, we use a pre-trained feature extractor and define the validity score to be the cosine similarity between the features extracted from the original image and the transformed image. For PDEs, the validity score is defined by the numerical errors of the original equations after the transform. The validity scores are chosen based on the characteristics of the target tasks, and designed to be fully differentiable, so that the symmetry can be learned via gradient descent in an end-to-end fashion. We also incorporate two regularizations, orthonormality and Lipschitz loss, which prevent the learned generators from converging to a trivial solution.

Subsequently, we demonstrate that our method indeed discovers the correct symmetries in both image and PDE datasets. To the best of our knowledge, our research is the first to retrieve affine symmetry in the entire space of continuous transformations using the CIFAR-10 classification dataset, as shown in Figure 1. Moreover, our method excels in identifying non-affine symmetries and approximate symmetries in PDE tasks. We further demonstrate that the learned generators can be leveraged to develop automatic augmentation generators, which can be used to produce augmented training data for both image classification tasks and neural operator learning tasks of PDEs [21]. We provide empirical evidence that the models trained with data augmented by our learned generators perform competitively with those trained with traditional closed-form transforms such as Lie point symmetry (LPS) [4]. Moreover, we show that the *approximate symmetries* discovered by our method, which cannot be found by classical methods, can also boost the performance of the models, especially when the size of the training data is small.

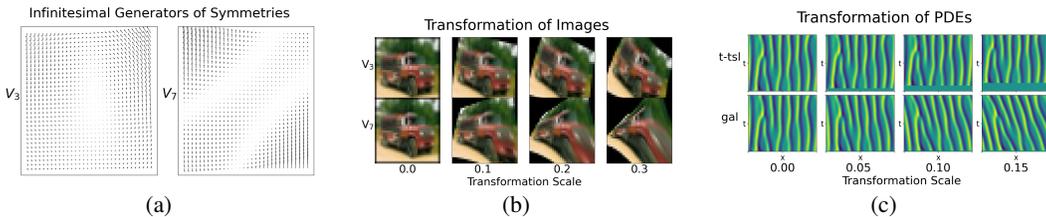


Figure 1: (a) Examples of the vector fields. V_3 is a learned symmetry which is approximately a rotation, while V_7 is not a symmetry, thus having a high validity score. (b) Transformed CIFAR-10 images using the learned generators. All the vector fields and transformations learned from CIFAR-10 are presented in Figure 8 of Appendix C. (c) Transformation of PDEs (KS equation) with learned symmetries: time translation (t-tsl) and Galilean boost (gal).

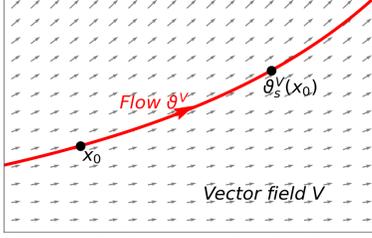


Figure 2: An example of a flow.

Table 1: Infinitesimal generators of the Affine group $\text{Aff}(2)$. The set of six generators $\{L_1, \dots, L_6\}$ forms a basis of the corresponding Lie algebra.

Generator	Expression	One-Parameter Group	Description
L_1	$(1, 0)$	$\mathbf{x} \mapsto \mathbf{x} + s(1, 0)$	translation in x_1 -axis
L_2	$(0, 1)$	$\mathbf{x} \mapsto \mathbf{x} + s(0, 1)$	translation in x_2 -axis
L_3	$(x_1, 0)$	$\mathbf{x} \mapsto \begin{pmatrix} e^s & 0 \\ 0 & 1 \end{pmatrix} \mathbf{x}$	scaling of x_1 -axis
L_4	$(0, x_2)$	$\mathbf{x} \mapsto \begin{pmatrix} 1 & 0 \\ 0 & e^s \end{pmatrix} \mathbf{x}$	scaling of x_2 -axis
L_5	$(x_2, 0)$	$\mathbf{x} \mapsto \begin{pmatrix} 1 & e^s \\ 0 & 1 \end{pmatrix} \mathbf{x}$	shear parallel to x_1 -axis
L_6	$(0, x_1)$	$\mathbf{x} \mapsto \begin{pmatrix} 1 & 0 \\ s & 1 \end{pmatrix} \mathbf{x}$	shear parallel to x_2 -axis
$L_3 + L_4$	(x_1, x_2)	$\mathbf{x} \mapsto e^s \mathbf{x}$	uniform scaling
$L_6 - L_5$	$(-x_2, x_1)$	$\mathbf{x} \mapsto \begin{pmatrix} \cos s & -\sin s \\ \sin s & \cos s \end{pmatrix} \mathbf{x}$	rotation

2 Preliminaries: One-parameter Group

In this section, we present the basic definitions of a one-parameter group, which we use to parameterize the symmetric transformations learned from the data.

Consider an unknown Euclidean domain $\mathcal{Z} \subseteq \mathbb{R}^n$ and a smooth vector field $V : \mathcal{Z} \rightarrow \mathbb{R}^n$. A path $\gamma : I = (a, b) \subseteq \mathbb{R} \rightarrow \mathcal{Z}$ satisfying $\frac{d}{ds}\gamma(s) = V(\gamma(s))$ for all $s \in I$ is a curve that travels around the domain \mathcal{Z} with a velocity given by the vector field V . Along the curve γ , a point $\mathbf{x}_0 = \gamma(a_0)$ can be transported to $\gamma(a_0 + s)$ by flowing along the vector field V by time s . We define the flow ϑ_s^V by $\vartheta_s^V(\mathbf{x}_0) = \gamma(a_0 + s)$ of V as in Figure 2. This flow is computed by solving an ODE

$$\frac{d}{ds}\vartheta_s^V(\mathbf{x}) = V(\mathbf{x}) \quad (1)$$

with initial condition $\vartheta_0^V(\mathbf{x}) = \mathbf{x}$ for all $\mathbf{x} \in \mathcal{Z}$ [20].

The flow ϑ^V is governed by an autonomous ODE, i.e., an ODE independent of the temporal variable s . Due to properties of autonomous ODEs, the flow ϑ^V exists uniquely and it is smooth in both variables s and \mathbf{x} . Assuming a mild condition on V , such as V extends to a compactly supported vector field \tilde{V} on \mathbb{R}^n , the ODE does not terminate in \mathbb{R}^n in finite time and hence ϑ_s^V is defined for all $s \in \mathbb{R}$. In that case, the flow satisfies a group equation

$$\vartheta_{s_1+s_2}^V(\mathbf{x}) = \vartheta_{s_1}^V \circ \vartheta_{s_2}^V(\mathbf{x}) \quad (2)$$

for all $s_1, s_2 \in \mathbb{R}$. It means that the flow can be regarded as a group action of \mathbb{R} on \mathbb{R}^n , transforming elements of $\mathcal{Z} \subseteq \mathbb{R}^n$. For this reason, ϑ_s^V is also called a *one-parameter group*, and the vector field V is called an *infinitesimal generator* of the one-parameter group.

On $\mathcal{Z} = \mathbb{R}^n$, a constant vector field $V(\mathbf{x}) = \mathbf{v} \in \mathbb{R}^n$ gives rise to a translation $\mathbf{x} \mapsto \mathbf{x} + s\mathbf{v}$. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, a vector field $V(\mathbf{x}) = \mathbf{A}\mathbf{x}$ gives rise to an affine transformation $\mathbf{x} \mapsto \exp(s\mathbf{A})\mathbf{x}$, where \exp is the matrix exponentiation. Multiple infinitesimal generators may span a vector space \mathfrak{g} , and if \mathfrak{g} satisfies some algebraic condition (closure under the Lie bracket), then \mathfrak{g} forms a *Lie algebra*. Composing the elements of one-parameter groups of elements in \mathfrak{g} gives rise to a *Lie group* G . The correspondence between G and \mathfrak{g} is called Lie group-Lie algebra correspondence.

Continuous symmetries are commonly defined by a Lie group G , acting on some domain and keeping the transformed objects *invariant* with respect to some criterion. We model symmetries by specifying their infinitesimal generators whose composition of one-parameter groups comprises the symmetries of that domain.

Below, we describe two representative examples that will be discussed extensively in the remainder of the paper: images (interpreted as functions on 2D planes) and PDEs.

2.1 Images and Their Symmetries

Consider a rescaled image of the form $f : \mathcal{X} = [-1, 1]^2 \subseteq \mathbb{R}^2 \rightarrow \mathcal{Y} = [0, 1]^3 \in \mathbb{R}^3$. The affine transformations on \mathbb{R}^2 have the form $\mathbf{x} = (x_1, x_2) \mapsto \mathbf{A}\mathbf{x} + \mathbf{b}$ for a matrix $\mathbf{A} \in \mathbb{R}^{2 \times 2}$ and a vector $\mathbf{b} \in \mathbb{R}^2$. The Affine transformations form the 6-dimensional Affine group $\text{Aff}(2)$, and it has a corresponding 6-dimensional Lie algebra having a basis $\{L_1, \dots, L_6\}$ given as in Table 1. The symmetries of images are often exploited as a data augmentation strategy for learning image classifiers, under an assumption that the transforms do not alter the identity or semantics of the images to be classified.

2.2 PDEs and Their Symmetries

Given an n -dimensional independent variable $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X} \subseteq \mathbb{R}^n$ and an m -dimensional dependent variable $\mathbf{u} = \mathbf{u}(\mathbf{x}) = (u_1(\mathbf{x}), \dots, u_m(\mathbf{x})) \in \mathcal{U} \subseteq \mathbb{R}^m$, we denote by $\mathbf{u}^{(i)}$ the collection of all i -th partial derivatives of \mathbf{u} with respect to \mathbf{x} . A *partial differential equation* Δ on $\mathbf{u}(\mathbf{x})$ of order k is defined by a set of algebraic equations $\Delta(\mathbf{x}, \mathbf{u}, \mathbf{u}^{(1)}, \dots, \mathbf{u}^{(k)}) = 0$ involving all the variables and their partial derivatives. For example, two scalar independent variables $x, t \in \mathbb{R}$ and one scalar dependent variable $u(x, t) \in \mathbb{R}$ governed by equation $\Delta = u_t + uu_x + u_{xxx} = 0$ gives the 1-dimensional Korteweg-de Vries (KdV) equation, where we denote partials using subscripts, e.g. $u_t = \frac{\partial u}{\partial t}$ and $u_{xxx} = \frac{\partial^3 u}{\partial x^3}$. The KdV equation is commonly used to model the dynamics of solitons, e.g. shallow water waves [35]. The KdV equation described above is an example of 1-dimensional scalar-valued evolution equation. Such an equation takes the form $u = u(x, t) \in \mathbb{R}$ with its governing equation of the form

$$u_t = F(x, t, u, u_x, u_{xx}, u_{xxx}, \dots) \quad (3)$$

for some function F . In this paper, we only deal with 1D scalar-valued evolution equation on a fixed periodic domain $x \in [0, L]$.

Continuous symmetries of PDEs are commonly parametrized by a one-parameter group on $\mathcal{X} \times \mathcal{U}$. Denote $(\xi, \mu) = (\xi[\mathbf{x}, \mathbf{u}], \mu[\mathbf{x}, \mathbf{u}])$ an infinitesimal generator defined on $\mathcal{X} \times \mathcal{U}$. Then the PDE Δ possesses the infinitesimal generator of symmetry (ξ, μ) if the equation is still satisfied after transforming both the independent variable \mathbf{x} and the dependent variable \mathbf{u} [27, 5, 26]. Symmetries of PDEs are categorized by how the generators (ξ, μ) depend on (\mathbf{x}, \mathbf{u}) . The symmetry is a Lie point symmetry (LPS) if the value of (ξ, μ) at each point $(\mathbf{x}, \mathbf{u}(\mathbf{x}))$ depends only on the point value $(\mathbf{x}, \mathbf{u}(\mathbf{x}))$ itself. If (ξ, μ) also depends on the derivatives $\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(k)}$ at that point, it is called a Lie-Bäcklund symmetry or generalized symmetry. If (ξ, μ) depends on integrals of \mathbf{u} , then it is called a nonlocal symmetry. Finding an LPS of a PDE Δ can be done algorithmically under some mild assumptions on Δ . However, there is no general recipe of finding Lie-Bäcklund symmetries or nonlocal symmetries, and discovering such symmetries remains an active area of research.

3 Related Work

Symmetry discovery. Approaches to learning symmetries can be categorized by addressing two questions: (a) *where do they search for symmetries*, and (b) *what are they aiming to learn*. One line of research aims to learn ranges, focusing on determining the ranges of transformation scales that enhance learning when employed as augmentation techniques. For example, Benton et al. [3] learns transformation ranges of predefined transformations by treating them as learnable parameters and backpropagating through differentiable transformations.

Another line of research aim to learn subgroups of bigger candidate groups, typically a linear group $GL(n)$ or an affine group $\text{Aff}(n)$. For example, Desai et al. [10] use the Generative Adversarial Network (GAN) to search for symmetries, with the generator transforming data by group elements sampled from the candidate group and the discriminator verifying whether the transformed data still lies in the data distribution. Similarly, Yang et al. [34] employ the GAN approach, but generator of GAN models infinitesimal generators instead of the subgroup itself. As an alternative, Moskalev et al. [24] proposed an idea of extracting symmetries from learned neural network by differentiating through it, and retrieved 2D rotation in the linear group using the rotation MNIST dataset.

Finally, *learning symmetries with minimal assumption*, i.e. without assuming the infinitesimal generators are linear or affine, is an area of large interest. An early attempt of Rao & Ruderman [28] models infinitesimal generator by a learnable matrix from the pixel space to the pixel space, and learn 2D rotation by solving a task that compares original images and rotated ones, where the images are 5×5 random pixels. Sohl-Dickstein et al. [32] takes the similar approach with eigen-decomposing the learnable matrix. Dehmamy et al. [9] builds a convolution operation whose kernel encodes learnable infinitesimal generators, and retrieved 2D rotation from random 7×7 images by comparing original and transformed ones. Our work closely aligns with Liu & Tegmark [22] and Forestano et al. [13], which model one-parameter groups by an MLP and learn the symmetries from an invariant scalar quantity. To the best of our knowledge, learning correct symmetries with minimal assumption was only achieved with toy datasets, far from real-world datasets such as CIFAR-10.

Utilizing symmetries in deep learning. An effective method for leveraging symmetries in deep learning is data augmentation [31]. In the image domain, there are numerous augmentation techniques available [29], most of which are based on geometric properties of images. Although data augmentation techniques have been primarily explored in the context of images, recent studies by [4, 23] have demonstrated that symmetries can also be used for augmenting data in the training of neural PDE solvers. In addition to data augmentation, some approaches involve designing new neural network architectures that inherently reflect the group symmetries of the input data [7]. Wang et al. [33] applied a similar strategy within the PDE domain.

4 Learning Continuous Symmetries with One-Parameter Groups

4.1 Training Process

Given a learning task with a dataset $\mathcal{D} \subset \mathcal{A}$ in an underlying space $\mathcal{A} = \{f|f : \mathcal{X} \rightarrow \mathcal{Y}\}$, we aim to model symmetry by a one-parameter group ϑ_s acting on \mathcal{A} , as explained in § 4.3. We define a continuous symmetry by stating that ϑ is a symmetry of this task if there exists some $\sigma > 0$ such that for any data point $f \in \mathcal{D}$ and transformation scale $s \in [-\sigma, \sigma]$, the transformed data point $\vartheta_s(f)$ remains valid for this task. We assume the existence of a differentiable *validity score* $S(\vartheta_s, f) \in \mathbb{R}$, such that $\vartheta_s(f) \in \mathcal{A}$ is valid if $S(\vartheta_s, f) < C$ for a certain threshold $C \in \mathbb{R}$. Then, a one-parameter group ϑ is a symmetry of the task if $S(\vartheta_s, f) < C$ for all $f \in \mathcal{D}$.

The validity score depends on the nature of the target task, though no strict criterion exists. As long as it is differentiable and the valid data aids learning, it is considered acceptable. For instance, we can define the validity based on a negative log-likelihood of a probabilistic model. In § 4.2, we discuss the validity scores to be used for image and PDE data.

Once a validity score is defined, we learn a symmetry ϑ^* by minimizing the validity scores of transformed data,

$$\vartheta^* = \arg \min_{\vartheta} \mathbb{E}_{f \sim \mathcal{D}, s \sim \text{Unif}([- \sigma, \sigma])} [S(\vartheta_s, f)], \quad (4)$$

where the $\arg \min$ is taken over the entire class of smooth one-parameter groups. Since the learning is performed in function space, we appropriately constrain the function space using a regularizer, as described in § 4.4. Once symmetries are learned, they reveal the symmetrical properties of the target task, which can then be exploited to augment the training data.

4.2 Task-specific Definition of Validity Score S

Images. In image-related tasks, we define a validity score using a pre-trained neural network. Let \mathcal{D} be an image classification dataset consisting of data of the form $(f, y) \in \mathcal{A} \times \mathbb{R}$, where f is an image and y is a label. Also let $H_{\text{cls}} \circ H_{\text{fext}} : \mathcal{A} \rightarrow \mathbb{R}$ be a learned neural network, where we denote by $H_{\text{fext}} : \mathcal{A} \rightarrow \mathbb{R}^k$ the feature extractor and $H_{\text{cls}} : \mathbb{R}^k \rightarrow \mathbb{R}$ the classifier. We define the validity score $S(\vartheta_s, f)$ as the cosine similarity between the features before and after the transformation:

$$S(\vartheta_s, f) = \text{sim}(H_{\text{fext}}(\vartheta_s(f)), H_{\text{fext}}(f)), \quad (5)$$

where sim is the cosine similarity defined as $\text{sim}(\mathbf{v}_1, \mathbf{v}_2) = \frac{|\mathbf{v}_1 \cdot \mathbf{v}_2|}{\|\mathbf{v}_1\| \|\mathbf{v}_2\|}$ for all $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^k \setminus \{\mathbf{0}\}$.

PDEs. Let $\mathbf{u}(\mathbf{x})$ be a solution of a given PDE Δ , discretized on a rectangular grid $\mathcal{X}_{\text{grid}} = \{\mathbf{x}_i\}_{i=1}^{N_{\text{grid}}}$. For a transformed data $\vartheta_s(\mathbf{u})$, we measure the violation of the equality of the $\Delta = 0$ to assess whether the transformed data is still a valid solution. Using an appropriate numerical differentiation method, we directly compute the value of the PDE, denoted as $\Delta(\vartheta_s(\mathbf{u}))$, which represents the error of $\vartheta_s(\mathbf{u})$ as a solution of Δ , taking a value $\Delta(\vartheta_s(\mathbf{u}))_i$ at grid point \mathbf{x}_i . The validity score is defined by the summation of all PDE errors across the grid points:

$$S(\vartheta_s, \mathbf{u}) = \sum_i |\Delta(\vartheta_s(\mathbf{u}))_i|. \quad (6)$$

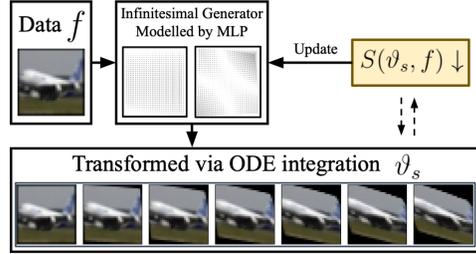


Figure 3: Process of learning symmetry.

For example, for a solution $u(x)$ of the 1D KdV equation, we examine whether the transformed solution $\tilde{u} = \vartheta_s(u)$ satisfies $\tilde{u}_t + \tilde{u}\tilde{u}_x + \tilde{u}_{xxx} = 0$ where the partials are computed using a numerical differentiation method.

4.3 Parametrization of One-Parameter Groups using Neural ODE

On a Euclidean domain $\mathcal{Z} \subseteq \mathbb{R}^n$, we model an infinitesimal generator with an MLP $\mathbf{h}_\theta : \mathcal{Z} \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$. The infinitesimal generator \mathbf{h}_θ gives rise to a one-parameter group $\vartheta_s^{\mathbf{h}_\theta}$. We sample a transformation scale $\alpha \sim \text{Unif}([- \sigma, \sigma])$ for a predefined hyperparameter $\sigma \in \mathbb{R}_{>0}$. To transform a point $x \in \mathcal{Z}$ along this one-parameter group by an amount $\alpha \geq 0$, we use a numerical ODE solver to solve the ODE for $\gamma : [0, \alpha] \rightarrow \mathcal{Z}$ satisfying

$$\gamma'(s) = \mathbf{h}_\theta(\gamma(s)), \quad \forall s \in [0, \alpha], \quad \gamma(0) = x \quad (7)$$

and obtain a transformed data point $\tilde{x} = \vartheta_\alpha^{\mathbf{h}_\theta}(x) = \gamma(\alpha)$. If $\alpha < 0$, we compute $\vartheta_\alpha^{\mathbf{h}_\theta}(x) = \vartheta_{-\alpha}^{-\mathbf{h}_\theta}(x)$ by integrating $-\mathbf{h}_\theta$ instead of \mathbf{h}_θ using the ODE solver. We can backpropagate through the numerical ODE solver using the adjoint method [6] to learn θ .

Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a data point on a domain \mathcal{A} . As \mathcal{A} is a space of functions, naively modeling symmetry on \mathcal{A} may ignore the geometry implied in the input space \mathcal{X} . Instead, we define two transformations: $\vartheta_{\mathcal{X}}$ on \mathcal{X} and $\vartheta_{\mathcal{Y}}$ on \mathcal{Y} , and induce a transformation of f by

$$(\vartheta_{\mathcal{X}}(f))(x) = f(\vartheta_{\mathcal{X}}^{-1}(x)), \quad (\vartheta_{\mathcal{Y}}(f))(x) = \vartheta_{\mathcal{Y}}(f(x)), \quad (8)$$

where we abuse notation and write the transformed function as $\vartheta_{\mathcal{X}}(f)$ and $\vartheta_{\mathcal{Y}}(f)$. For an image represented as a discretized function $f : \mathcal{X} \rightarrow \mathcal{Y}$ from $\mathcal{X} = [-1, 1]^2$ and $\mathcal{Y} = [0, 1]^3$, $\vartheta_{\mathcal{X}}$ corresponds to spatial transformations such as translation or rotation, and $\vartheta_{\mathcal{Y}}$ corresponds to color space transformations. For a PDE, a 1D scalar-valued evolution equation on a fixed periodic domain takes the form $u(x, t) \in \mathcal{U} = \mathbb{R}$ with $(x, t) \in [0, L] \times [0, T] = \mathcal{X} \subseteq \mathbb{R}^2$, and we parameterize an infinitesimal generator on a product space $\mathcal{X} \times \mathcal{U} \subseteq \mathbb{R}^3$ by an MLP. Then, a transformation on $(x, t, u) \in \mathcal{X} \times \mathcal{U}$ induces a transformation on the solution of the PDE $u(x, t)$.

4.4 Objective Functions

Symmetry loss. Let N_{sym} be the number of symmetries to be learned. Let $(\mathbf{h}_\theta^{(a)})_{a=1}^{N_{\text{sym}}}$ be the infinitesimal generators computed from a single MLP. For each $a \in \{1, \dots, N_{\text{sym}}\}$, we sample a transformation scale $s_a \sim \text{Unif}([- \sigma, \sigma])$ to transform f via numerical integration. The parameter θ is optimized by minimizing the average validity score over the training data,

$$\mathcal{L}_{\text{sym}}(\theta) = \sum_{a=1}^{N_{\text{sym}}} \mathbb{E}_{f \sim \mathcal{D}, s_a \sim \text{Unif}([- \sigma, \sigma])} \left[S\left(\vartheta_{s_a}^{\mathbf{h}_\theta^{(a)}}, f\right) \right]. \quad (9)$$

Orthonormality loss. Learning only with the symmetry loss may result in trivial solutions such as the zero vector field or the same vector field repeated in multiple slots. To prevent this, we introduce the orthonormality loss to regularize the model towards learning orthonormal vector fields. Specifically, given two vector fields $V_1, V_2 : \mathcal{Z} \rightarrow \mathbb{R}^n$, we define an inner product as,

$$\langle V_1, V_2 \rangle = \frac{1}{\text{vol}(\mathcal{Z})} \int_{\mathcal{Z}} \omega(\mathbf{x})(V_1(\mathbf{x}) \cdot V_2(\mathbf{x})) d\mathbf{x} \approx \frac{1}{|\mathcal{Z}_{\text{grid}}|} \sum_{\mathbf{x}_i \in \mathcal{Z}_{\text{grid}}} \omega(\mathbf{x}_i)(V_1(\mathbf{x}_i) \cdot V_2(\mathbf{x}_i)), \quad (10)$$

with a suitable weight function $\omega(x) : \mathcal{Z} \rightarrow \mathbb{R}$ and a discretized grid $\mathcal{Z}_{\text{grid}}$ of \mathcal{Z} of size $|\mathcal{Z}_{\text{grid}}|$. Given this definition, we first normalize each generator by its norm to ensure $\|\mathbf{h}_\theta^{(a)}\|^2 = 1$. Then we compute the orthonormality loss as,

$$\mathcal{L}_{\text{ortho}}(\theta) = \sum_{1 \leq a < b \leq N_{\text{sym}}} \left\langle \text{sg}(\mathbf{h}_\theta^{(a)}), \mathbf{h}_\theta^{(b)} \right\rangle, \quad (11)$$

where $\text{sg}(\cdot)$ denotes the stop-gradient operation to ensure that the constraint $\langle \mathbf{h}_\theta^{(a)}, \mathbf{h}_\theta^{(b)} \rangle = 0$ only affects the latter slot (b). By doing this, if the true number of symmetries N_{sym}^* is less than or equal to the assumed number of symmetries N_{sym} , the learned symmetries will be aligned in the first N_{sym}^* slots.

Lipschitz loss. We further introduce inductive biases to the infinitesimal generators we aim to learn. For instance, an infinitesimal generator moving only a single pixel near the boundary by a large scale would be undesirable. This idea can be implemented using Lipschitz continuity. For a grid point $\mathbf{x}_i \in \mathcal{Z}_{\text{grid}}$ and its neighboring point $\mathbf{x}_j \in \text{nbhd}(\mathbf{x}_i) \subset \mathcal{Z}_{\text{grid}}$, we expect the vector field V to satisfy the Lipschitz condition,

$$\text{Lips}(V; \mathbf{x}_i, \mathbf{x}_j) < \tau \text{ where } \text{Lips}(V; \mathbf{x}_i, \mathbf{x}_j) = \frac{\|V(\mathbf{x}_i) - V(\mathbf{x}_j)\|}{\|\mathbf{x}_i - \mathbf{x}_j\|}. \quad (12)$$

To regularize the model toward the Lipschitz condition, we introduce the Lipschitz loss,

$$\mathcal{L}_{\text{Lips}}(\boldsymbol{\theta}) = \sum_{a=1}^{N_{\text{sym}}} \sum_{\mathbf{x}_i \in \mathcal{Z}_{\text{grid}}, \mathbf{x}_j \in \text{nbhd}(\mathbf{x}_i)} \max(\text{Lips}(\mathbf{h}_{\boldsymbol{\theta}}^{(a)}; \mathbf{x}_i, \mathbf{x}_j) - \tau, 0). \quad (13)$$

Total loss and loss-scale-independent learning. We jointly minimize the three loss functions with suitable weights $w_{\text{sym}}, w_{\text{ortho}}, w_{\text{Lips}} > 0$ and learn the weights $\boldsymbol{\theta}$ of MLP using a stochastic gradient descent:

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} w_{\text{sym}} \mathcal{L}_{\text{sym}}(\boldsymbol{\theta}) + w_{\text{ortho}} \mathcal{L}_{\text{ortho}}(\boldsymbol{\theta}) + w_{\text{Lips}} \mathcal{L}_{\text{Lips}}(\boldsymbol{\theta}). \quad (14)$$

To minimize the computational burden of hyperparameter tuning, we ensure that all the loss terms have a *natural scale*, i.e. a dimensionless scale independent of the context. For example, when penalizing the inner product in Equation 11, we apply \arccos to the normalized inner product to ensure the loss term lies in $[0, \pi/2)$. Similarly, the scale of the PDE validity score $S(\vartheta_s, \mathbf{u})$ in Equation 9 depends on the scale of the data \mathbf{u} . When penalizing it, we apply the log function so that the gradients are scaled automatically as $\nabla_{\boldsymbol{\theta}} \log(S(\vartheta_s, \mathbf{u})) = \nabla_{\boldsymbol{\theta}} S(\vartheta_s, \mathbf{u}) / S(\vartheta_s, \mathbf{u})$.

Here we describe the generic training process, but the actual implementation requires non-trivial task-specific designs, such as the choice of the weighting function $w(\mathbf{x})$ or the method for locating the transformed data on the target grid. We defer these details for image and PDE tasks to Appendix A.

4.5 Comparison With Other Methods

Here, we compare our method with other recent symmetry discovery methods. The differences mainly arise from (a) what they aim to learn (e.g., transformation scales or subgroups from a larger group) and (b) their assumptions about prior knowledge (e.g., complete, partial, or no knowledge of symmetry generators). Another important distinction is the viewpoint on symmetry: some methods learn symmetries that raw datasets inherently possess (implicit), while others learn symmetries from datasets explicitly designed to carry such symmetries (explicit).

Some recent symmetry discovery works are listed in Table 2. We emphasize that our method excels in two key aspects: (a) our learning method reduces infinitely many degrees of freedom, (b) our method works with high-dimensional real-world datasets. For example, while LieGAN [33] and LieGG [24] reduce a 6-dim space (affine) to a 3-dim space (translation and rotation), ours reduces an ∞ -dim space to a finite one. L-conv [9] also does not assume any prior knowledge, but it is limited in finding rotation in a toy task, where it learns rotation angles of rotated images by comparing them with the original ones, which are 7x7 random pixel images.

5 Experiments

5.1 Images

We use images of size 32×32 from the CIFAR-10 classification task. Since our method does not model discrete symmetry, we use horizontal flip with 50% probability by default. We train a ResNet-18 model, which will be used as the feature extractor H_{fext} in Equation 5. The weight function on the pixels is computed as explained in Appendix A.1. We expect to find 6 affine generators and we use an MLP modeling 10 vector fields in the pixel space $[-1, 1]^2 \subseteq \mathbb{R}^2$, expecting the first six learned vector fields to be the affine generators. We learn the Equation 14 using stochastic gradient descent with $w_{\text{sym}} = 1$ and $w_{\text{ortho}}, w_{\text{Lips}} = 10$. The parameter σ , which controls the scale of transformation, is set to $\sigma = 0.4$, and the Lipschitz threshold τ is set to $\tau = 0.5$. Other details are described in Appendix B.1. We conducted three trials with random initializations and report the full results in

Table 2: Comparison with other symmetry discovery methods.

	Augerino	LieGAN	LieGG	L-conv	Forestano et al.	Ours
Symmetry generators	completely known	partially known (affine)	partially known (affine)	completely unknown	completely unknown	completely unknown
Learn what?	transformation scales	symmetry generator (rotation)	symmetry generator (rotation)	symmetry generator (rotation)	symmetry generator (in low-dim task)	symmetry generator (affine)
Verified with what?	raw CIFAR-10	rotation MNIST	rotation MNIST	random 7×7 pixel image	toy data (dim ≤ 10)	raw CIFAR-10 & PDEs
Implicit or explicit?	implicit	explicit	explicit	explicit	explicit	implicit
How?	optimize while training downstream task	compare fake/true data in GAN framework	extracts from learned NN using Lie derivative	compare rotated and original images	extracts from invariant oracle using Lie derivative	extracts from validity score using ODE integration

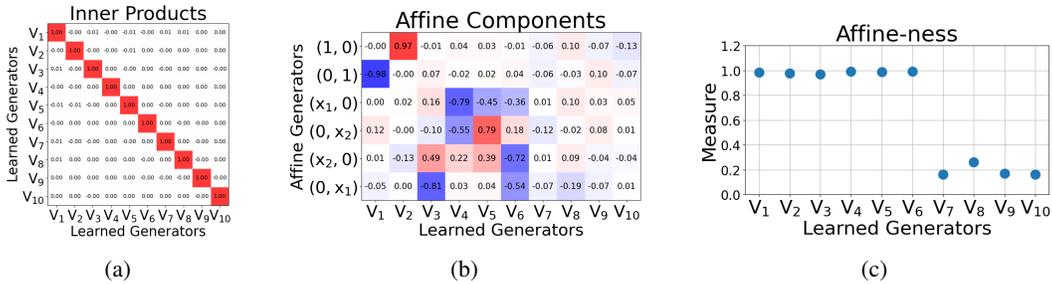


Figure 4: (a) Self inner-products of the learned generators. (b) Inner product comparison of the learned generators with the affine generators. (c) Affine-ness of learned generators.

[Appendix C.1](#). Furthermore, we also learn symmetries in the *color space*, and their results are shown in [Appendix G](#).

Learned symmetries. Since we expect to learn affine symmetries, we compare the results with the affine basis $\{L_1, \dots, L_6\}$ defined in [Table 1](#). We compute the inner products $\langle V, L_i \rangle$ of the learned vector field V with L_i for $i = 1, \dots, 6$ to measure how much the learned vector fields contain the affine basis and measure the affine-ness of vector field by $\text{Affine-ness}(V)^2 = \sum_{i=1}^6 \langle V, L_i \rangle^2$.

In all experiments, we successfully retrieve six linearly independent affine generators in the first six slots. [Figure 4a](#) shows that the learned generators are orthogonal to each other, as desired. [Figure 4b](#) shows the inner product between the learned generators and the affine generators. Since the affine-ness measure of the first 6 learned generators in [Figure 4c](#) is almost close to 1, we can read out the affine components in [Figure 4b](#) and say that e.g., $V_1 \approx (0, -0.98 + 0.12x_2)$. Notably, two translation generators are found in the first two slots, indicating that the two translations are *the most invariant* one-parameter group among the entire class of one-parameter groups on the pixel space. After the two translation generators, four affine generators are learned, indicating that affine transformations are *the next most invariant* transformations. In particular, the third and fourth generators are close to the rotation generator and the scaling generator, respectively. The remaining four generators fix pixels close to the center and transform boundary pixels by a large magnitude.

Analysis. Unlike other symmetry discovery researches [[24](#), [9](#)] that use datasets which are explicitly designed to be symmetric such as rotation MNIST, we discovered affine transformations from CIFAR-10 with no augmentation except horizontal flip. Moreover, we extract the symmetries from the ResNet trained with CIFAR-10 without augmentation. This implies that although CIFAR-10 is not explicitly composed to be symmetric under affine transformations, the dataset possesses intrinsic affine symmetry. This also implies that even the ResNet trained without augmentation possesses invariance under affine transformation. It is widely believed that the strong generalization power of neural networks is linked to the augmentation insensitivity of the neural networks [[25](#)]. Our results show that

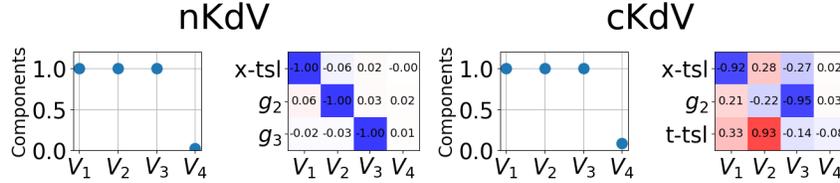


Figure 5: Inner products between the learned non-affine symmetry generators and the ground truth. The results including the affine symmetry generators are shown in Figure 10a.

Table 4: Definition of PDEs.

Name	Equation
KdV	$u_t + uu_x + u_{xxx} = 0$
KS	$u_t + u_{xx} + u_{xxxx} + uu_x = 0$
Burgers	$u_t + uu_x - \nu u_{xx} = 0$
nKdV	$e^{-\frac{t}{t_0}} u_t + uu_x + u_{xxx} = 0$
cKdV	$u_t + uu_x + u_{xxx} + \frac{u}{2(t+1)} = 0$

Table 5: LPS of PDEs.

Name	Lie Point Symmetries		
KdV [4], KS [4], Burgers [17]	$(1, 0, 0)$	$(0, 1, 0)$	$(t, 0, 1)$
nKdV (Appendix D.3)	$(1, 0, 0)$	$(0, e^{-\frac{t}{t_0}}, 0)$	$(t_0(e^{\frac{t}{t_0}} - 1), 0, 1)$
cKdV [30]	$(1, 0, 0)$	$(\sqrt{t+1}, \frac{1}{2\sqrt{t+1}}, 0)$	

ResNet is insensitive to affine transformation even when not explicitly designed to be so, supporting this hypothesis. This result implies that augmentation explicitly amplifies the insensitivity by using transformed data for training. Our analysis is tangential to that of Gruver et al. [16], which measures the extent of invariance using derivatives along the infinitesimal generators instead of ODE integration.

Augmentation results. We train a ResNet-18 model using CIFAR-10 classification data, applying the learned symmetries as data augmentation. We compare the results of no-augmentation, default augmentation (horizontal flip and random crop), and affine transformation, with transformation scale searched in $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. We conduct five experiments with random initialization for all the settings and report the results in Table 3.

Table 3: Test accuracy in CIFAR-10.

Method	Acc. (%)
No-aug	92.4 ± 0.3
Default	95.1 ± 0.1
Affine	95.1 ± 0.2
Learned	94.9 ± 0.1

5.2 PDEs

We follow the experimental setting of Brandstetter et al. [4], which use the Korteweg-de Vries (KdV) equation, the Kuramoto-Shivashinsky (KS) equation, and the Burgers' equation on a 1D periodic domain as experiments. They all have time translation, space translation, and the Galilean boost as LPSs. To consider PDEs with non-trivial and non-affine symmetries, we add two variants of the KdV, namely the nKdV equation and the cKdV equation. The nKdV is yielded by a nonlinear time translation of the original KdV, and the cKdV is the cylindrical KdV equation, having an extra time-dependent term. The equations are listed in Table 4, and their symmetries are listed in Table 5. Note that some symmetries are ruled out since we work within a fixed periodic domain, e.g., the scaling symmetry of the KdV. Since there are at most three symmetries, we open four slots in an MLP and learn symmetries using weights $w_{\text{sym}}, w_{\text{Lips}} = 1$ and $w_{\text{ortho}} = 3$.

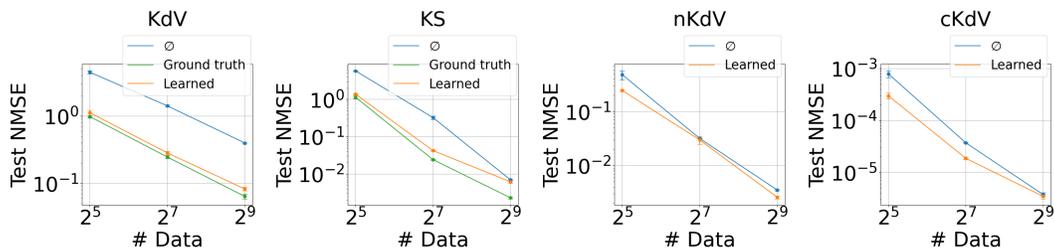


Figure 6: Comparison of augmentation performances using the ground truth symmetries and the learned symmetries with various numbers of data. The symbol \emptyset stands for no-augmentation.

Table 6: Test NMSE comparison of augmentation using LPS and AS in FNO learning for cKdV.

# Data	None	LPS	AS	LPS+AS
2^7	$(3.70 \pm 0.09) \times 10^{-6}$	$(3.49 \pm 1.06) \times 10^{-6}$	$(3.20 \pm 0.94) \times 10^{-6}$	$(2.66 \pm 0.36) \times 10^{-6}$
2^5	$(7.90 \pm 1.21) \times 10^{-4}$	$(5.90 \pm 0.17) \times 10^{-4}$	$(4.45 \pm 0.15) \times 10^{-4}$	$(3.70 \pm 0.15) \times 10^{-4}$

Learned symmetries. We compare the learned symmetries with the ground-truth symmetries using inner products. We found the ground truth symmetries in all the experiments as in Figure 5 and Figure 10.

Interestingly, for the Burgers’ equation and the cKdV equation, we additionally found the u -axis rescaling operation $(0, 0, u)$ and the time translation $(0, 1, 0)$ respectively. Applying the u -axis rescaling $u \mapsto cu$ for $c \approx 1$ to the Burgers’ equation $u_t + uu_x - \nu u_{xx} = 0$ gives $cu_t + c^2uu_x - \nu cu_{xx} = c(c - 1)uu_x$, leaving only the uu_x term. The uu_x term in Burgers’ is called the *convection term*, and it is approximately zero in most region and spikes in some small region. Similarly, the time translation $t \mapsto t + c$ for $c \approx 0$ fixes the first three terms in the cKdV equation $u_t + uu_x + u_{xxx} + u/(2(t + 1)) = 0$ and only changes the last term $u/(2(t + 1))$ by a negligible amount. These are not LPSs of the given equations, but the error of the PDE after transformation is smaller than the error of the numerical differentiation method. These are *approximate symmetries* (AS), and the theory of AS is also of great interest in the symmetry analysis of PDEs [2, 1].

Augmentation results. We use the learned symmetries as data augmentation and train Fourier Neural Operators (FNOs). The detailed experiment setting is described in Appendix B.2. Since FNOs are extremely sensitive to numerical error, we employ Whittaker-Shannon interpolation, explained in Appendix D.2, to resample the transformed results. The results are depicted in Figure 6. In all cases, data augmentation using the learned symmetries improve the performance, almost close to the results using the ground truth symmetries. The detailed results are in Appendix C.2. Additionally, we verify that the approximate symmetry of cKdV is also beneficial to training, as shown in Table 6, especially when the numbers of data points is low, proving the effectiveness of symmetries extracted from data.

Ablations. Additional ablation studies on numerical methods, such as numerical differentiation and interpolation, and hyperparameter sensitivity are conducted, and their results are presented in Appendices E and F.

6 Conclusion

We have introduced a novel method for learning continuous symmetries, including non-affine transformations, from data without prior knowledge. By leveraging Neural ODE, our approach models one-parameter groups to generate a sequence of transformations, guided by a task-specific validity score function. This approach captures both affine and non-affine symmetries in image and PDE datasets, enhancing automatic data augmentation in image classification and neural operator learning for PDEs. The learned generators produce augmented training data that improve model performance, particularly with limited training data.

Limitation. However, despite its flexibility, our method requires careful selection of numerical methods, such as numerical differentiation and interpolation, to ensure stable training and the ODE integration can be computationally large for augmentation generation compared to other augmentation methods. While we focus on image symmetries and LPSs of PDEs, the method could potentially model other symmetries and domains with proper validity scores, suggesting future applications in learning complex symmetries, including conditional and non-local symmetries, in various data types.

Acknowledgments and Disclosure of Funding

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.RS-2019-II190075, Artificial Intelligence Graduate School Program(KAIST), No.2022-0-00713, Meta-learning Applicable to Real-world Problems, No.2022-0-00184, Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics, No.RS-2024-00509279, Global AI Frontier Lab)

References

- [1] Baikov, V., Gazizov, R., and Ibragimov, N. K. Approximate symmetry and formal linearization. *Journal of Applied Mechanics and Technical Physics*, 30(2):204–212, 1989. [10](#)
- [2] Baumann, G. Approximate symmetries of partial differential equations. *Symmetry Analysis of Differential Equations with Mathematica®*, pp. 404–423, 2000. [10](#)
- [3] Benton, G. W., Finzi, M., Izmailov, P., and Wilson, A. G. Learning invariances in neural networks from training data. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2020. [1](#), [4](#)
- [4] 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. [1](#), [2](#), [5](#), [9](#), [14](#), [15](#)
- [5] Cantwell, B. J. *Introduction to symmetry analysis*. Cambridge University Press, 2002. [4](#)
- [6] Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018. [2](#), [6](#)
- [7] Cohen, T. and Welling, M. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016. [1](#), [5](#)
- [8] Dangovski, R., Jing, L., Loh, C., Han, S., Srivastava, A., Cheung, B., Agrawal, P., and Soljagic, M. Equivariant self-supervised learning: Encouraging equivariance in representations. In *International Conference on Learning Representations (ICLR)*, 2022. [1](#)
- [9] Dehmamy, N., Walters, R., Liu, Y., Wang, D., and Yu, R. Automatic symmetry discovery with lie algebra convolutional network. *Advances in Neural Information Processing Systems*, 34: 2503–2515, 2021. [2](#), [4](#), [7](#), [8](#)
- [10] Desai, K., Nachman, B., and Thaler, J. Symmetry discovery with deep learning. *Physical Review D*, 105(9):096031, 2022. [4](#)
- [11] Dey, N., Chen, A., and Ghafurian, S. Group equivariant generative adversarial networks. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021. [1](#)
- [12] Dumbser, M. and Käser, M. Arbitrary high order non-oscillatory finite volume schemes on unstructured meshes for linear hyperbolic systems. *Journal of Computational Physics*, 221(2): 693–723, 2007. [14](#), [16](#), [18](#)
- [13] Forestano, R. T., Matchev, K. T., Matcheva, K., Roman, A., Unlu, E. B., and Verner, S. Deep learning symmetries and their lie groups, algebras, and subalgebras from first principles. *Machine Learning: Science and Technology*, 4(2):025027, 2023. [2](#), [4](#)
- [14] Gabel, A., Quax, R., and Gavves, E. Data-driven lie point symmetry detection for continuous dynamical systems. *Machine Learning: Science and Technology*, 2024. [2](#)
- [15] Garrido, Q., Najman, L., and LeCun, Y. Self-supervised learning of split invariant equivariant representations. In *Proceedings of The 40th International Conference on Machine Learning (ICML 2023)*, 2023. [1](#)
- [16] Gruver, N., Finzi, M. A., Goldblum, M., and Wilson, A. G. The lie derivative for measuring learned equivariance. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023. [9](#)
- [17] Hoarau, E. and David, C. Lie group computation of finite difference schemes. *arXiv preprint math/0611895*, 2006. [9](#)
- [18] Hoogeboom, E., Satorras, V. G., Vignac, C., and Welling, M. Equivariant diffusion for molecule generation in 3d. In *Proceedings of The 39th International Conference on Machine Learning (ICML 2022)*, 2022. [1](#)

- [19] Huang, Y., Peng, X., Ma, J., and Zhang, M. 3dlinker: An $E(3)$ equivariant variational autoencoder for molecular linker design. In *Proceedings of The 39th International Conference on Machine Learning (ICML 2022)*, 2022. 1
- [20] Lee, J. M. *Smooth manifolds*. Springer, 2012. 3
- [21] Li, Z., Kovachki, N. B., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A. M., and Anandkumar, A. Fourier neural operator for parametric partial differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021. 2
- [22] Liu, Z. and Tegmark, M. Machine learning hidden symmetries. *Physical Review Letters*, 128(18):180201, 2022. 4
- [23] 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. 1, 5
- [24] Moskalev, A., Sepiarskaia, A., Sosnovik, I., and Smeulders, A. W. M. Liegg: Studying learned lie group generators. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022. 2, 4, 7, 8
- [25] Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Sensitivity and generalization in neural networks: an empirical study. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018. 8
- [26] Oliveri, F. Lie symmetries of differential equations: Classical results and recent contributions. *Symmetry*, 2(2):658–706, 2010. ISSN 2073-8994. doi: 10.3390/sym2020658. URL <https://www.mdpi.com/2073-8994/2/2/658>. 2, 4
- [27] Olver, P. J. *Applications of Lie groups to differential equations*, volume 107. Springer Science & Business Media, 1993. 4
- [28] Rao, R. and Ruderman, D. Learning lie groups for invariant visual perception. *Advances in neural information processing systems*, 11, 1998. 4
- [29] Shorten, C. and Khoshgoftaar, T. M. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019. 1, 5
- [30] Sil, S. Nonclassical symmetries, nonlinear self-adjointness, conservation laws and some new exact solutions of cylindrical kdv equation. *International Journal of Applied and Computational Mathematics*, 9(5):69, 2023. 9
- [31] Simard, P. Y., Steinkraus, D., Platt, J. C., et al. Best practices for convolutional neural networks applied to visual document analysis. In *Icdar*, volume 3. Edinburgh, 2003. 1, 5
- [32] Sohl-Dickstein, J., Wang, C. M., and Olshausen, B. A. An unsupervised algorithm for learning lie group transformations. *arXiv preprint arXiv:1001.1027*, 2010. 4
- [33] Wang, R., Walters, R., and Yu, R. Incorporating symmetry into deep dynamics models for improved generalization. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021. 1, 5, 7
- [34] Yang, J., Walters, R., Dehmamy, N., and Yu, R. Generative adversarial symmetry discovery. In *Proceedings of The 40th International Conference on Machine Learning (ICML 2023)*, 2023. 1, 4
- [35] Zabusky, N. J. and Kruskal, M. D. Interaction of "solitons" in a collisionless plasma and the recurrence of initial states. *Physical review letters*, 15(6):240, 1965. 4
- [36] Zhang, Y.-T. and Shu, C.-W. Eno and weno schemes. In *Handbook of numerical analysis*, volume 17, pp. 103–122. Elsevier, 2016. 14, 16

A Implementation Details

A.1 Image Dataset

Notation. Images are functions of the form $f : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{X} = [-1, 1]^2$ is the spatial domain and $\mathcal{Y} = [0, 1]^3$ is the normalized RGB domain. Pixels $\mathcal{X}_{\text{grid}} = \{\mathbf{x}_i\}_{i=1}^{N_{\text{grid}}}$ are discretized through a rectangular grid of \mathcal{X} , and images f are discretized on the pixel space by $f_i = f(\mathbf{x}_i)$ for all i .

$\omega(\mathbf{x})$ in orthonormality loss. We first learn a symmetry on the spatial domain \mathcal{X} using a feature extractor H_{hidden} taken from a pre-trained neural network. One obstacle is that in most image datasets, the main subjects of images are mostly located around the centers, and the regions close to the boundary are filled with backgrounds. In other words, each pixel has a different level of importance, and we may end up learning infinitesimal generators that only move boundary pixels and fix the center.

We take into account the importance of pixels using the weight function $\omega(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$. For a discretized image $f = \{f_i\}$ and for each pixel \mathbf{x}_i , we measure a *pixel sensitivity* of the image f at the i -th pixel up to the feature extractor H_{fext} as

$$\text{Sensitivity}(f, \mathbf{x}_i) = \left\| \frac{\partial H_{\text{fext}}(f)}{\partial \mathbf{x}_i} \right\| = \left\| \frac{\partial f_i}{\partial \mathbf{x}_i} \frac{\partial H_{\text{fext}}(f)}{\partial f_i} \right\|, \quad (15)$$

which can be computed by querying a Jacobian-vector product of H_{fext} with respect to the image gradient at each pixel $\frac{\partial f_i}{\partial \mathbf{x}_i}$. We define the weight $\omega(\mathbf{x})$ as the average of pixel sensitivity across the dataset:

$$\omega(\mathbf{x}_i) = \mathbb{E}_{f \sim \mathcal{D}}[\text{Sensitivity}(f, \mathbf{x}_i)] = \mathbb{E}_{f \sim \mathcal{D}} \left[\left\| \frac{\partial H_{\text{fext}}(f)}{\partial \mathbf{x}_i} \right\| \right]. \quad (16)$$

However, computing this weight function ω needs $N_{\text{grid}} \cdot N_{\text{data}}$ times of computation of Jacobian-vector products. Instead, we use a Gaussian kernel $\kappa(\mathbf{x}; \hat{\mathbf{x}}, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp(-\frac{\|\mathbf{x} - \hat{\mathbf{x}}\|^2}{2\sigma^2})$, with a fixed $\sigma = 0.1$ and the center $\hat{\mathbf{x}}$ sampled uniformly on \mathcal{X} , and approximate the weight function by

$$\omega(\mathbf{x}_i) \approx \mathbb{E}_{f \sim \mathcal{D}, \hat{\mathbf{x}} \sim \text{Unif}(\mathcal{X})} \left[\kappa(\mathbf{x}_i; \hat{\mathbf{x}}, \sigma) \left\| \sum_j \frac{\partial H_{\text{fext}}(f)}{\partial \mathbf{x}_j} \kappa(\mathbf{x}_j; \hat{\mathbf{x}}, \sigma) \right\| \right] \quad (17)$$

so that we compute the weight in a stochastic manner. Intuitively, it computes Jacobian-vector product for each larger pseudo-pixel represented by the Gaussian kernels instead of each individual pixel. We iterate over the dataset 20 times, and the computed weight function is shown in [Figure 7](#).

Training details. The infinitesimal generator $\mathbf{h}_{\theta}^{(a)}$ transforms each pixel \mathbf{x}_i into $\tilde{\mathbf{x}}_i$ via ODE integration. In this process, the transformed pixels $\{\tilde{\mathbf{x}}_i\}$ may no longer be located on the rectangular grid $\mathcal{X}_{\text{grid}}$, so we use the bilinear interpolation method to resample the transformed image $\vartheta_s(f)$ on $\mathcal{X}_{\text{grid}}$. Note that the bilinear interpolation operation is differentiable, thereby we can train the MLP by minimizing the validity score $S(\vartheta_s, f)$, which is the cosine similarity between the features of f and $\vartheta_s(f)$:

$$S(\vartheta_s, f) = \text{sim}(H_{\text{fext}}(\vartheta_s(f)), H_{\text{fext}}(f)). \quad (18)$$

A.2 PDE Dataset

Notation. Solutions of a 1D scalar-valued evolution equation on a periodic domain take the form $u(x, t) : \mathcal{X} \rightarrow \mathcal{U}$ for $\mathcal{X} = [0, L] \times [0, T]$ and $\mathcal{U} = \mathbb{R}$, where $[0, L]$ is the spatial domain and $[0, T]$ is the temporal domain. Similar to the image task, the domain \mathcal{X} is discretized by the rectangular grid $\mathcal{X}_{\text{grid}} = \{\mathbf{x}_i\}_{i=1}^{N_{\text{grid}}} = \{(x_i, t_i)\}_{i=1}^{N_{\text{grid}}}$. A discretized solution is a set of tuples (x_i, t_i, u_i) where $u_i = u(x_i, t_i)$.

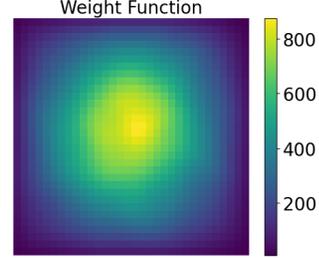


Figure 7: The weight function $\omega(\mathbf{x})$.

$\omega(x)$ in orthonormality loss. Unlike in images, we assume that all discretized grid points hold equal importance, so we set $\omega(x) = 1$ for all $x \in \mathcal{X}$.

Training details. The infinitesimal generators on the product space $\mathcal{X} \times \mathcal{U}$ transform a point $\mathbf{u} = \{(x_i, t_i, u_i)\}$ into $\tilde{\mathbf{u}} = \{(\tilde{x}_i, \tilde{t}_i, \tilde{u}_i)\}$. Therefore, the transformed solutions are no longer on the rectangular grid. To compute the PDE value $\Delta(\tilde{\mathbf{u}})$ for $\tilde{\mathbf{u}} = \vartheta_s(\mathbf{u})$ in Equation 6, we need to compute partial derivatives numerically such as $u_x(\tilde{x}_i, \tilde{t}_i)$, $u_t(\tilde{x}_i, \tilde{t}_i)$ or $u_{xx}(\tilde{x}_i, \tilde{t}_i)$. We use the *finite difference method*, in which the numerical derivative is approximated by finite differences, e.g.,

$$u_x(x_i, t_i) = \frac{u(x_i + \Delta x, t_i) - u(x_i - \Delta x, t_i)}{2\Delta x}. \quad (19)$$

for some small Δx .

In particular, we use the weighted essentially non-oscillating (WENO) scheme as a numerical differentiation method [36], with a careful choice of parameters as in Dumbser & Käser [12]. In the WENO method, multiple estimates for derivatives are made using multiple sets of neighboring gridpoints (called *stencils*). The multiple estimates are then averaged with weights (called *smoothness indicator*) that approximate how stable the derivative estimates are. We implement the WENO method working on a nonuniform grid, and the model learns symmetries by backpropagating through it. A detailed description on the WENO scheme is in Appendix D.1.

B Experiment Details

In this section, we present the detailed experimental settings of the experiments in § 5.

B.1 Images

Training ResNet-18. When training the ResNet-18 with CIFAR-10, both the feature extractor H_{fext} and models after augmentation, we train the model in 200 epochs with a batch size 128. The learning rate is set to 10^{-1} and decreases by a factor of 0.2 at the 60th, 120th, and 160th epoch. The model is trained by SGD optimizer with Nesterov momentum 0.9 and weight decay 0.4.

Learning symmetries. To learn the symmetry generators, we train the MLP using two shared hidden layers, each with a width of 256, followed by a hidden layer of width 32 for each output vector field. We use the swish activation function to ensure the learned vector fields are smooth. The MLP is trained for 50 epochs with a batch size 128 and fixed learning rate of 10^{-4} using the Adam optimizer. The learning process takes less than 10 hours on a GeForce RTX 2080 Ti GPU.

B.2 PDEs

Data generation. We follow the data generation method of Brandstetter et al. [4]. Given an 1D evolution equation $u_t = F(x, t, u, u_x, u_{xx}, \dots)$ for $u = u(x, t)$ on periodic domain $[0, L]$, we start with an initial condition $u(x, 0)$ by random Fourier series as

$$u(x, 0) = \sum_{p=1}^P A_p \sin(2\pi l_p x / L + \phi_p) \quad (20)$$

where P is the number of Fourier modes and (A_p, l_p, ϕ_p) are random coefficients. For time-stepping, we compute x -derivatives u_x, u_{xx}, \dots using pseudospectral method, which computes derivatives in Fourier domain and converts them back to the original domain. We use an ODE solver to compute the time evolution of $u(x, t)$ from $t = 0$ to $t = T$. The solution is discretized on regular grid of size $N_x \times N_t$ on $[0, L] \times [0, T]$, where $N_x = 256$ and $N_t = 140$ by default.

When simulating Burgers' equation, we instead solve the PDE for the heat equation $\phi_t = \phi_{xx}$ for $\phi = \nu\phi(x, t)$. The Burgers' equation $u(x, t)$ and the heat equation $\phi(x, t)$ is related via the Cole-Hopf transformation:

$$u = 2\nu \frac{\partial}{\partial x} \log(\phi). \quad (21)$$

After data generation, we transform the heat equation back to the Burgers' equation.

Learning symmetries. We use an MLP with the same architecture as described in [Appendix B.1](#). Since all the tuples (x, t, u) must pass through the MLP, totaling $N_x \times N_t$ for each data instance, we set a small batch size 4. Symmetries are learned using 1024 data instances over 50 epochs. We use the Adam optimizer and train the network with a learning rate of 10^{-4} in the first 25 epochs and 10^{-5} for the remaining 25 epochs. The learning process also takes less than 10 hours on a GeForce RTX 2080 Ti GPU.

Sobolev regularization. To ensure smoothness in the vector fields, we apply additional regularization using the Sobolev norm of order 2 in the x -domain. For a vector field $V(x, t)$, the Sobolev norm can be efficiently computed in Fourier domain:

$$\|V(\cdot, t)\|_{2,2}^2 = \sum_{i=0}^2 \left\| \frac{\partial^i V(\cdot, t)}{\partial x^i} \right\|_2^2 = \sum_{n=0}^{N_x-1} \left(1 + \left|\frac{n_{\text{freq}}}{L}\right|\right)^2 \hat{V}(n, t) \quad (22)$$

where $n_{\text{freq}} = \min(n, N_x - n)$ and $\hat{V}(\cdot, t)$ is the discrete Fourier transformation of $V(\cdot, t)$. Since we already enforce $\|V\| = 1$, we penalize towards the Sobolev norm excluding the zeroth order term:

$$\|V(\cdot, t)\|_{2,2}^2 - \|V(\cdot, t)\|_2^2 = \sum_{i=1}^2 \left\| \frac{\partial^i V(\cdot, t)}{\partial x^i} \right\|_2^2 = \sum_{n=0}^{N_x-1} \left(\left(1 + \left|\frac{n_{\text{freq}}}{L}\right|\right)^2 - 1 \right) \hat{V}(n, t). \quad (23)$$

We apply the Sobolev regularization during the final 10 epochs.

Training FNOs. For the KdV and KS equations, we train an autoregressive FNO solver, which takes 20 timesteps as input and predicts the subsequent 20 timesteps, following the experimental setup of Brandstetter et al. [4]. For the nKdV and cKdV equations, which are time-dependent and contain explicit t terms, we train FNOs as single-time neural operators. These models utilize the initial conditions of the equations to predict the states after 70 timesteps. We train the FNO over 40 epochs, with each epoch comprising 280 iterations across the dataset for the KdV and KS equation and 100 iterations for the nKdV and cKdV equations. The learning rate begins at 10^{-4} and decreases by a factor of 0.4 every 10 epochs.

C Experiment Results

In this section, we provide detailed results of the experiments outlined in [§ 5](#). For symmetry learning tasks, we conducted each experiment three times and randomly selected one for reporting in [§ 5](#). The complete set of results is provided here. Also, we report the detailed evaluation metrics in augmentation tasks.

C.1 Images

[Figure 8](#) is a visualization of the learned symmetries discussed in [§ 5.1](#). [Figure 9](#) displays results from experiments conducted under the same settings but with different model initializations compared to [Figure 4](#) in [§ 5.1](#). It is notable that affine symmetries consistently occupy the first six slots across all experiments.

C.2 PDEs

We report experiments results for learning symmetries, conducting three trials for each equation in [figure Figure 10](#). The ground truth symmetries and the approximate symmetries are consistently found in the former slots in all the experiments. The remaining fourth slots occasionally converge to learned symmetries despite the orthonormality loss or converge to some unknown vector fields with high validity scores. The results from the first trials are used in the augmentation experiments.

In [Table 7](#) and [Table 8](#), we provide the augmentation results of FNOs using the learned symmetries, which are illustrated in [Figure 6](#).

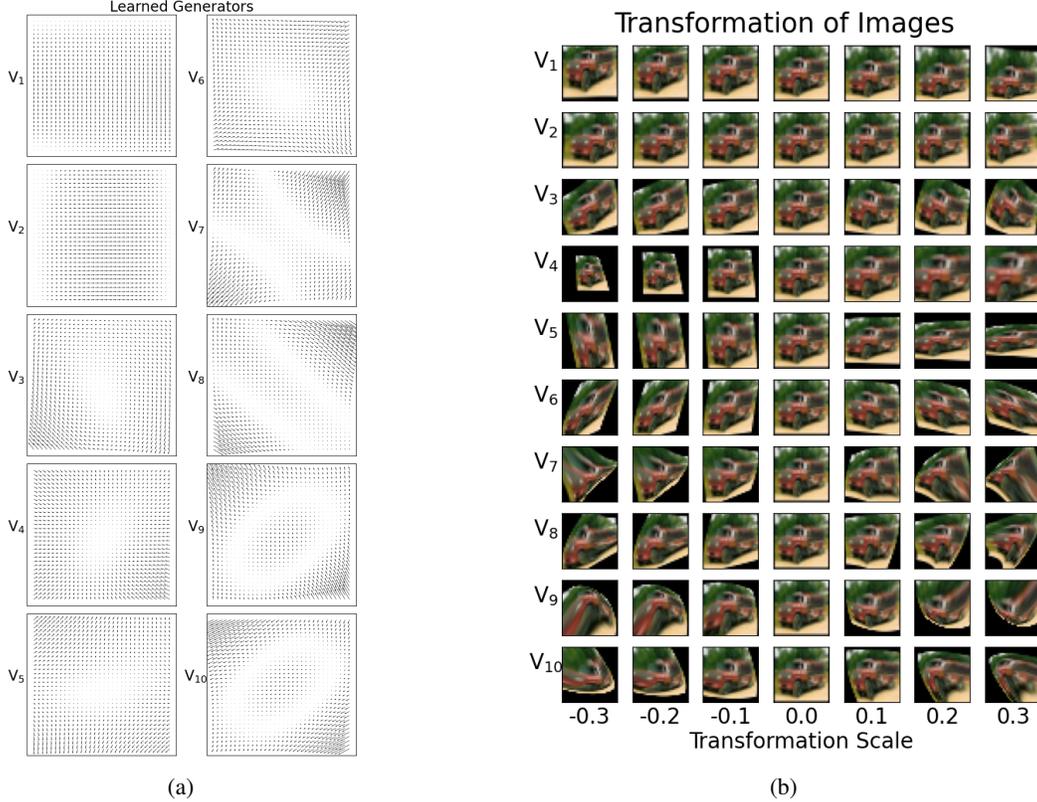


Figure 8: (a) The learned vector fields. (b) Transformed images using the learned generators. The images with transformation scale 0 are the original images.

Table 7: Comparison of augmentation methods with different numbers of data for KdV and KS

# Data	KdV			KS		
	None	Ground-truth	Ours	None	Ground-truth	Ours
2^9	0.398 ± 0.013	0.0640 ± 0.0050	0.0824 ± 0.0052	0.00693 ± 0.00039	0.00229 ± 0.00014	0.00614 ± 0.00051
2^7	1.42 ± 0.05	0.246 ± 0.012	0.283 ± 0.018	0.324 ± 0.031	0.0241 ± 0.0001	0.0422 ± 0.0002
2^5	4.47 ± 0.26	0.980 ± 0.039	1.12 ± 0.10	5.78 ± 0.13	1.14 ± 0.10	1.37 ± 0.04

D Technical Details

D.1 Weighted Essentially Non-Oscillating (WENO) Scheme

This section is largely based on Zhang & Shu [36], Dumbser & Käser [12].

WENO scheme in 1D. Consider a smooth function $f : \mathcal{X} \subseteq \mathbb{R} \rightarrow \mathbb{R}$, where we only have access to the function on the grid $\mathcal{X}_{\text{grid}} = \{x_1, \dots, x_{N_{\text{grid}}}\}, x_1 < \dots < x_{N_{\text{grid}}}$. To estimate the derivatives $f^{(k)}(x) = \frac{d^k f}{dx^k}(x)$ for some $x \in \mathcal{X}$, we use $k + 1$ neighboring points in $I = \{x_i, \dots, x_{i+k}\}$ and compute the unique k -th order polynomial $p_I(x)$ that interpolates the function f on the points in I . In other words, we ensure $p_I(x_i) = f(x_i), \dots, p_I(x_{i+k}) = f(x_{i+k})$. Then we can pick up the estimates of the derivatives using the polynomial as $f^{(k)}(x) \approx p_I^{(k)}(x)$. We call the set of neighboring points I a *stencil*, and the polynomial p_I the *reconstruction polynomial*. Usage of reconstruction polynomials is a fundamental concept in the numerical differentiation method known as the *finite difference method*.

In the weighted essentially non-oscillating (WENO) scheme, we compute multiple estimates $p_{I_1}^{(k)}(x), \dots, p_{I_{N_s}}^{(k)}(x)$ for the derivatives $f^{(k)}(x)$ using multiple stencils I_1, \dots, I_{N_s} . The estimates

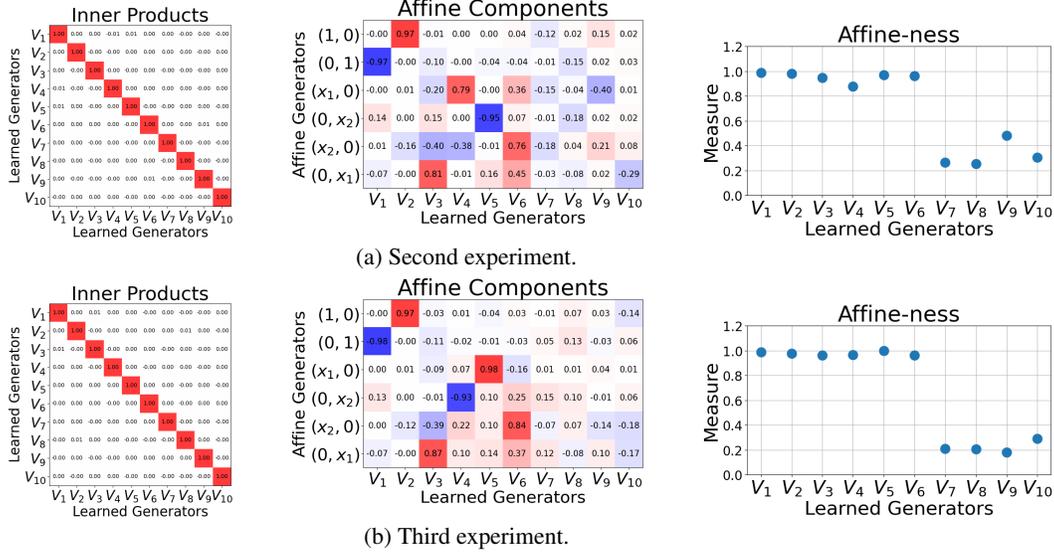


Figure 9: The same results (learned symmetries of images) under different initializations compared to Figure 4.

Table 8: Comparison of augmentation methods with different numbers of data for nKdV and cKdV

# Data	nKdV		cKdV	
	None	Ours	None	Ours
2^9	$(3.52 \pm 0.11) \times 10^{-3}$	$(2.58 \pm 0.17) \times 10^{-3}$	$(3.72 \pm 0.33) \times 10^{-6}$	$(3.44 \pm 0.37) \times 10^{-6}$
2^7	$(3.16 \pm 0.06) \times 10^{-2}$	$(2.96 \pm 0.50) \times 10^{-2}$	$(3.71 \pm 0.09) \times 10^{-5}$	$(1.86 \pm 0.10) \times 10^{-5}$
2^5	$(4.79 \pm 0.09) \times 10^{-1}$	$(2.44 \pm 0.11) \times 10^{-1}$	$(7.90 \pm 1.21) \times 10^{-4}$	$(2.99 \pm 0.37) \times 10^{-4}$

are then averaged with weights $\omega_1, \dots, \omega_{N_s}$ which take account of the quality of each estimation. In the WENO scheme, we assume the grid points are sufficiently dense, hence the estimates are more accurate when the reconstruction polynomials p_{I_m} are smooth. The smoothness of p_{I_m} is measured by the *smoothness indicator*, defined as:

$$IS_{I_m} = \sum_{\alpha=1}^k \int_{\Delta} |\Delta|^{\alpha-1} (p_{I_m}^{(\alpha)}(x))^2 dx \quad (24)$$

where $\Delta = (x_j, x_{j+1}) \subset \mathbb{R}$ is an interval between two grid points containing x , and $|\Delta|$ is the length of Δ . The weights $\omega_1, \dots, \omega_{N_s}$ are defined as:

$$\omega_m = \frac{\alpha_m}{\sum_{m'} \alpha_{m'}}, \quad \alpha_m = \frac{\gamma_m}{(\epsilon + IS_{I_m})^b} \quad (25)$$

where γ_m is called the *linear weight*, usually chosen heuristically, b is a positive integer usually set to 2 or 4, and $\epsilon > 0$ is a small positive number preventing the denominator from being zero. The final estimate for the derivative $f^{(k)}(x)$ is computed by averaging the estimates $p_{I_1}^{(k)}(x), \dots, p_{I_{N_s}}^{(k)}(x)$ with the weights $\omega_1, \dots, \omega_{N_s}$:

$$\hat{f}^{(k)}(x) = \sum_{m=1}^{N_s} \omega_m \cdot p_{I_m}^{(k)}(x). \quad (26)$$

WENO scheme for multivariate function. The WENO scheme extends smoothly to multidimensional function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. For $\mathbf{k} = (k_1, \dots, k_n) \in \mathbb{Z}_{\geq 0}^n$, we denote $f^{(\mathbf{k})}(x) = \frac{\partial^{\mathbf{k}} f}{\partial x_1^{k_1} \dots \partial x_n^{k_n}}(x)$.

On a stencil I with an appropriate number of grid points, we compute the reconstruction polynomial $p_I(x)$ with a nonzero \mathbf{k} -degree term, and the smoothness indicator is defined similarly as:

$$IS_I = \sum_{1 \leq |\alpha| \leq k} \int_{\Delta} |\Delta|^{|\alpha|-1} (p_{I_m}^{(\alpha)}(x))^2 dx \quad (27)$$

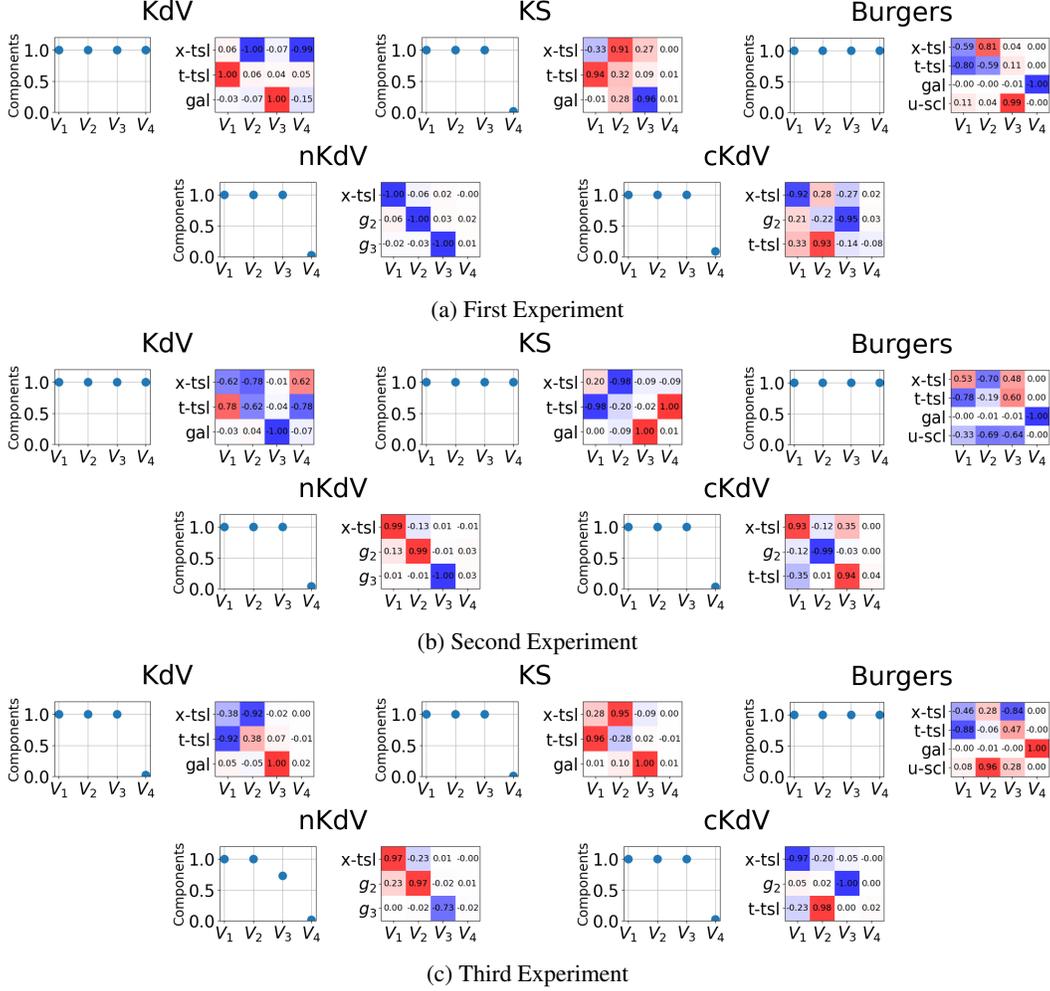


Figure 10: Inner product comparison of the learned symmetry generators with the ground truth symmetries. (a), (b), and (c) are conducted with different random initializations.

where $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_{\geq 0}^n$, $|\alpha| = \sum_{i=1}^n \alpha_i$ and the interval Δ is now a n -dimensional cell containing the point x with volume $|\Delta|$. Like the 1D case, the weights are defined using the linear weights and smoothness indicator.

Choice of parameters. We follow the choice of parameters of Dumbser & Käser [12]. The linear weight γ_m is chosen to be 100 if x is inside the center cell of I_m and 1 otherwise, reflecting the fact that the estimates are accurate as much as the point x is close to the center of I_m . The parameters b and ϵ are chosen as $b = 4$ and $\epsilon = 10^{-6}$.

D.2 Whittaker-Shannon Interpolation on a Periodic Domain

Let $\dots, f[-1], f[0], f[1], \dots$ be a discretization of a continuous signal on a real line. The Whittaker-Shannon interpolation recovers the original signal f as:

$$f(t) = \sum_{n=-\infty}^{\infty} f[n] \text{sinc}(t - n) \quad (28)$$

where sinc is the normalized sinc function $\text{sinc}(t) = \frac{\sin(\pi x)}{\pi x}$. The *Nyquist-Shannon sampling theorem* states that states that $f(t)$ is the *perfect reconstruction* of f , in a sense that if the original function does not contain any frequencies higher than a certain threshold, called the *Nyquist frequency*, then f is perfectly reconstructs the original signal.

Similarly, if $f[0], f[1], \dots, f[N-1]$ is a discretization of a continuous signal on a periodic domain $[0, N]$ for some *even* positive integer N , the Whittaker-Shannon interpolation is given as:

$$f(t) = \sum_{n=-\infty}^{\infty} f[n \bmod N] \text{sinc}(t - n) = \sum_{n=0}^{N-1} f[n] \sum_{m=-\infty}^{\infty} \text{sinc}(t - mN) = \sum_{n=0}^{N-1} f[n] D_N(t) \quad (29)$$

where $D_N(t) = \sum_{n=-\infty}^{\infty} \text{sinc}(t - mN)$ is the *Dirichlet kernel* with period N . The Dirichlet kernel $D_N(t)$ has a closed-form expression

$$D_N(t) = \frac{\sin(\pi t)}{N \tan(\pi t/N)}. \quad (30)$$

D.3 Nonlinear Time Transformation of the KdV Equation

We apply a nonlinear time transformation t defined as $t = t_0(e^{\frac{\hat{t}}{t_0}} - 1)$, where t_0 is a constant scaling factor, to the KdV equation $u_t + uu_x + u_{xxx}$. The derivative with respect to t is transformed to

$$\frac{\partial}{\partial t} = \frac{\partial \hat{t}}{\partial t} \frac{\partial}{\partial \hat{t}} = e^{\frac{\hat{t}}{t_0}} \frac{\partial}{\partial \hat{t}} \quad (31)$$

hence the KdV equation with nonlinear time transformation (nKdV) becomes

$$e^{\frac{\hat{t}}{t_0}} u_{\hat{t}} u_t + uu_x + u_{xxx} = 0. \quad (32)$$

The three symmetries of the KdV changes accordingly. The space translation is left untouched, and the other two infinitesimal generators are transformed as:

$$\frac{\partial}{\partial t} = e^{\frac{\hat{t}}{t_0}} \frac{\partial}{\partial \hat{t}} \quad (33)$$

$$t \frac{\partial}{\partial x} + \frac{\partial}{\partial u} = t_0(e^{\frac{\hat{t}}{t_0}} - 1) \frac{\partial}{\partial x} + \frac{\partial}{\partial u} \quad (34)$$

where we used the derivative notation $(1, 0, 0) = \frac{\partial}{\partial x}$, $(0, 1, 0) = \frac{\partial}{\partial t}$ and $(0, 0, 1) = \frac{\partial}{\partial u}$.

E Ablation Studies

E.1 Numerical differentiation.

When searching for symmetries of PDEs, we use the WENO scheme as a numerical differentiation method. We have found that the choice of numerical differentiator is critical, as gradients must flow through numerical differentiation during the backpropagation step. To compare different methods, we experimented with using the WENO scheme and another method: bilinear interpolation followed by discrete numerical differentiation on a regular rectangular grid. When using bilinear interpolation followed by discrete differentiation, the loss failed to converge well, even with very small learning rate 10^{-6} . The learned vector fields were not orthogonal, and they only spanned the time translation and the space translation at best, as in [Figure 11](#). These experiments were conducted using the KS equation.

We hypothesize that the bilinear interpolation distributes each transformed point into two adjacent grid points, but discrete differentiation is done by repeatedly subtracting values of two adjacent grid points, so the gradient may not flow well during the backpropagation.

E.2 Whittaker-Shannon interpolation.

In the PDE augmentation task, when the transformed PDEs are resampled into the regular rectangular grid, we use Whittaker-Shannon instead of more commonly used bilinear interpolation. This step is also crucial for training FNOs with transformed data, since FNOs are extremely sensitive to numerical error and hence any aliasing effects must be avoided. We compared the augmentation results using Whittaker-Shannon interpolation with those using bilinear interpolation, using the KS equation with 512 training data and report the result in [Table 9](#). The results clearly demonstrate that bilinear interpolation adversely affects the training FNO models.

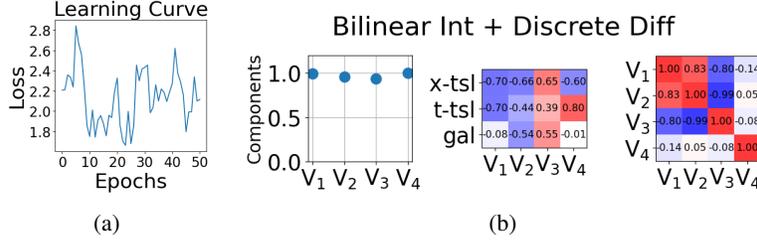


Figure 11: Experiment using using bilinear interpolation and discrete differentiation. (a) Learning curve. (b) Learned symmetries.

Table 9: Test NMSE comparison of augmentation using Whittaker-Shannon interpolation and bilinear interpolation.

	No-aug	Whittaker-Shannon	Bilinear
Test NMSE	0.00693 ± 0.00039	0.006143 ± 0.00051	0.542 ± 0.051

F Hyperparameter analysis

In this section, we analyze the roles of various hyperparameters in our learning scheme. We use experiment setting of the KS equation in § 5.2. We selected the KS equation for analysis as it proved to be the most challenging in learning symmetries, possibly due to its fourth-order derivative term u_{xxxx} . In particular, among the three symmetries (space translation, time translation, Galilean boost) the algorithm easily learns the space and time translation in a few epochs, but learning the Galilean boost takes more epochs as in Figure 12. We evaluate the results by examining whether the three symmetries are correctly found in the first three slots.

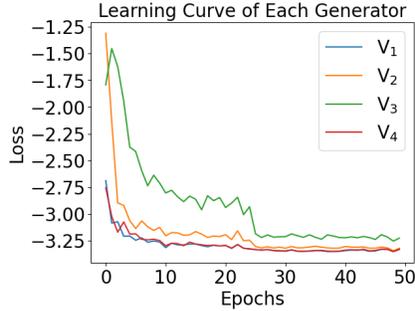


Figure 12: Learning curves of four symmetry generators.

F.1 Lipschitz Threshold

The Lipschitz threshold τ in Equation 13 is the only parameter that constrains the function space in which we search for the symmetries. The threshold τ is set to $\tau = 3$ in the main experiments. We conduct additional experiments with $\tau = 1, 2, 4, 8, 16, 32$. Surprisingly, in all experiments, regardless of the value of τ , we found the three ground truth symmetries.

F.2 Transformation Scale

The transformation scale σ controls how much we transform the data when searching for symmetries. The tuple (x, t, u) in $\mathcal{X} \times \mathcal{U}$ is scaled so that x and t form a uniform grid on $[0, 1]^2$, and u is scaled so that the standard deviation of u closely matches that of x and t , which is approximately 0.29. Hence, the default transformation scale $\sigma = 0.4$ transforms data with slightly more than its standard deviation at most. We conduct experiments with various transformation scales $\sigma = 0.1, 0.2, 0.4, 0.6, 0.8$. The ground truth symmetries were found when $\sigma = 0.1, 0.2, 0.4$. When $\sigma = 0.6$, the Galilean boost was found but allocated in the fourth slot, meaning that the learning was unstable. When $\sigma = 0.8$, the model only found space and time translation correctly.

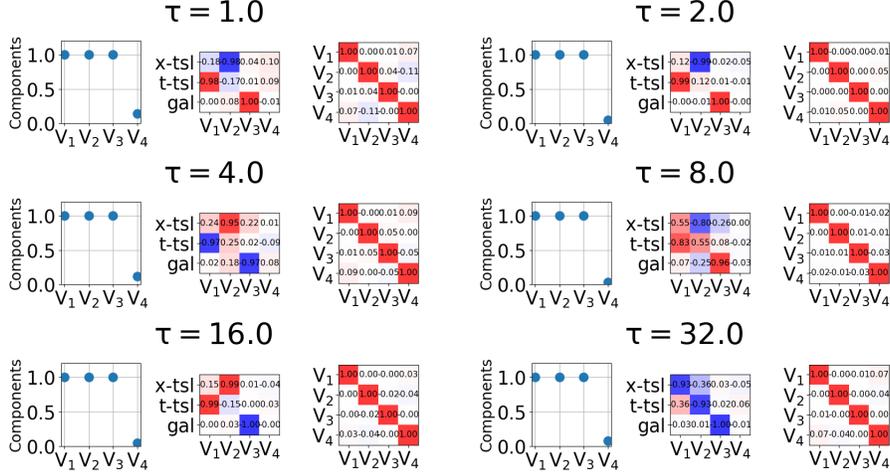


Figure 13: Experiments with various values of hyperparameter τ .

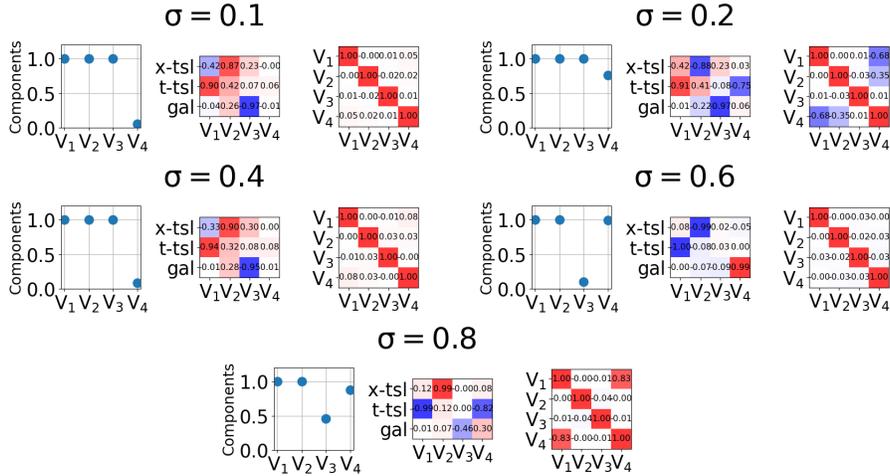


Figure 14: Experiments with various values of hyperparameter σ .

F.3 Weights of Three Losses

Although the three loss terms Equation 14 are designed to have dimensionless scales, tuning the three loss terms in appropriate ranges is inevitable. We perform a grid search over the three weights w_{sym} , w_{ortho} , $w_{\text{Lipschitz}}$ of the total loss in a fixed grid $[1, 3, 10]^3$. We report the number of correctly learned symmetries across different values of w_{sym} , w_{Lips} , and w_{ortho} in Table 10. The number 3 is the maximum number of symmetries to be discovered. We found that the ratio between w_{sym} and w_{ortho} significantly affects the results. In this case, the weight w_{ortho} should be larger than w_{sym} to prevent the learning of redundant infinitesimal generators.

G Color-space Results

Recall that in our formulation, images take the form $\mathcal{X} \rightarrow \mathcal{Y}$ where $\mathcal{X} = [-1, 1]^2$ is the pixel space and $\mathcal{Y} \in \mathbb{R}^3$ is the RGB color space. Our formulation is not limited to searching for the symmetries on the 2D plane \mathcal{X} , but is capable of learning symmetries in the color space \mathcal{Y} . In this section, we elaborate how we modeled learning scheme for color-space symmetries and present the results.

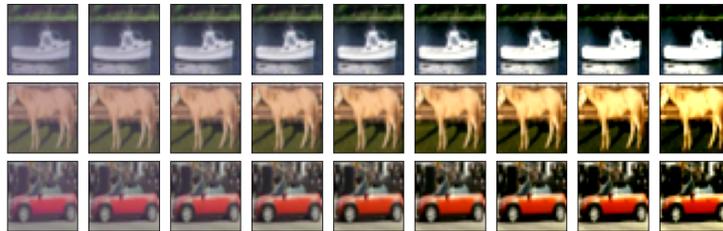
Table 10: The number of correctly learned symmetries across values of w_{sym} , w_{Lips} , and w_{ortho} .

(a) $w_{\text{sym}} = 1$				(b) $w_{\text{sym}} = 3$				(c) $w_{\text{sym}} = 10$			
w_{ortho}	w_{Lips}			w_{ortho}	w_{Lips}			w_{ortho}	w_{Lips}		
	1	3	10		1	3	10		1	3	10
1	2	3	3	1	2	2	2	1	1	2	2
3	2	3	3	3	3	3	3	3	2	2	2
10	3	3	3	10	3	3	3	10	3	3	3

Transformation of Images: V_1



Transformation of Images: V_2



Transformation of Images: V_3

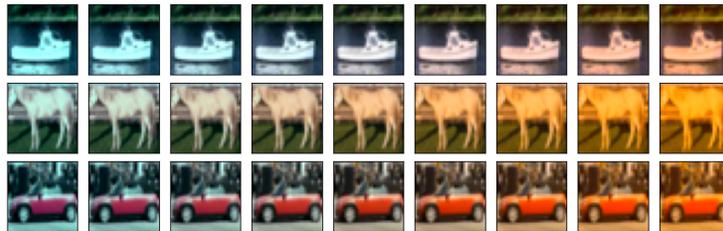


Figure 15: Learned color-space transformations.

G.1 Learning Scheme

Similar to other cases, we model one-parameter groups on \mathcal{Y} by a neural ODE, integrating over 3D vector fields $\mathcal{Y} \rightarrow \mathbb{R}^3$ parametrized by a neural network. Colors are non-uniform – for example, black, brown and gray appear much more frequently than pink or yellow in CIFAR-10 dataset. To address this, we first form a grid $\mathcal{Y}_{\text{grid}}$ on the color space \mathcal{Y} by randomly sampling 1024 colors from the CIFAR-10 dataset. Then we define the weight function $w : \mathcal{Y} \rightarrow \mathbb{R}$, like the weight function in Equation 17, by the *color sensitivity* of the neural network H_{fext} . For a grid point $\mathbf{y}_i \in \mathcal{Y}_{\text{grid}}$, the weight $\omega(\mathbf{y}_i)$ is defined as:

$$\omega(\mathbf{y}_i) = \mathbb{E}_{f \sim \mathcal{D}} \left[\left\| \frac{\partial H_{\text{fext}}(f)}{\partial \mathbf{y}_i} \right\| \right]. \quad (35)$$

To estimate the magnitude of the gradient with respect to the change of \mathbf{y}_i , we use neighboring points $\mathbf{y}_j \in \text{nbhd}(\mathbf{y}_i) \subset \mathcal{Y}_{\text{grid}}$. For a discretized image $\{\mathbf{x}_i, f(\mathbf{x}_i)\}_{\mathbf{x}_i \in \mathcal{X}_{\text{grid}}}$, we use the nearest neighbor

algorithm to collect $f(\mathbf{x}_i) \in \text{Nearest}(\mathbf{y}_i)$ close to \mathbf{y}_i , and measure how much H_{fext} changes when the color $f(\mathbf{x}_i)$ shifts towards the direction $\mathbf{y}_j - \mathbf{y}_i$:

$$\left\| \frac{\partial H_{\text{fext}}(f)}{\partial \mathbf{y}_i} \right\| = \frac{1}{|\text{nbhd}(\mathbf{y}_i)|} \sum_{\mathbf{y}_j \in \text{nbhd}(\mathbf{y}_i)} \left\| (\mathbf{y}_j - \mathbf{y}_i) \cdot \nabla_{\mathbb{1}_{\text{Nearest}(\mathbf{y}_i)}} H_{\text{fext}}(f) \right\| \quad (36)$$

where $\mathbb{1}_{\text{Nearest}(\mathbf{y}_i)}$ is a vector whose i -th entry is 1 if $f(\mathbf{x}_i) \in \text{Nearest}(\mathbf{y}_i)$ and 0 otherwise, and $\nabla_{\mathbb{1}_{\text{Nearest}(\mathbf{y}_i)}} H_{\text{fext}}(f)$ is computed by Jacobian-vector product.

Once the weight function is computed, we learn symmetries by optimizing through the validity score S , defined by cosine similarity of features of the learned neural network. Since there are (batch size) · (number of pixels) number of colors in a single batch of images, it's impractical to feed all of them directly into the MLP. Instead, we use the grid $\mathcal{Y}_{\text{grid}}$, and compute the vector field values on the grid and interpolate them using the nearest neighbor algorithm.

G.2 Results

We open three slots for color-space symmetries and train the model with $w_{\text{sym}}, w_{\text{Lips}} = 1$ and $w_{\text{ortho}} = 3$. We found the brightness control, color contrast and blue-yellow shift respectively as in [Figure 15](#).

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and precede the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

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

Answer: [Yes]

Justification: We clarified our contribution in the last two sentences in the abstract and the last paragraph in the introduction. Both sections mentioned that our scope is learning continuous symmetry with minimal prior knowledge.

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: We mentioned our limitations in the last two sentences in the conclusion.

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: Although we did not provide any theorems or propositions, which requires proofs, all formulas are numbered and cross-referenced.

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 checked the reproducibility of the reported experiments and we provided the hyperparameter analysis in the paper and we shared our source code.

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 attached our source code and we used public datasets, which are freely accessed by anyone.

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: All hyperparameters used in the experiments are provided in the appendix.

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: [Yes]

Justification: We measured the performance with 3-5 repeats. In Table 1 and Table 4, we reported the standard deviation, and in Figure 6(B), we described it using the error bar.

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: We also mentioned such information in the appendix.

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: We checked the code of ethics and our paper does not cause any hazard because it only explores theoretical symmetries.

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: [NA]

Justification: Since we explore only theoretical symmetries present in mathematical space, our work is far from the societal impacts.

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: We do not have 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: We properly credited when adopting the existing codes and properly cited the original paper that produced the code and dataset.

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 shared our code with anonymization and 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 any potential risks incurred by study participants

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.