Latent Space Symmetry Discovery

Jianke Yang [1] **Nima Dehmamy** [2] **Robin Walters** [3] **Rose Yu** [1]

1. Introduction

Symmetry plays an important role in the success of deep neural networks (Bronstein et al., 2021). Many different architectures have been developed to enforce various symmetries for modeling data with different structures and properties (Weiler & Cesa, 2019; Cohen et al., 2019a; Zaheer et al., 2017; Wang et al., 2021; Finzi et al., 2020; Kondor & Trivedi, 2018; Cohen et al., 2019b; Finzi et al., 2021; Bekkers, 2019).

Existing equivariant networks require knowing the symmetry explicitly before model implementation. However, for certain functions or data distributions, the underlying symmetries may not yet discovered or challenging to articulate through programming. There are different attempts towards automatic symmetry discovery from data, but most of them search only a limited portion of the space of all possible symmetries, such as subsets of known groups (Benton et al., 2020; Romero & Lohit, 2021) and finite groups (Zhou et al., 2020). LieGAN (Yang et al., 2023) can discover various types of symmetries, but its search space is still constrained to general linear groups. To achieve successful discovery, the observations have to be measured in an ideal coordinate system where linear symmetry is present. Unfortunately, this is often not true in many realistic tasks, such as high-dimensional dynamical systems (Champion et al., 2019) and vision tasks like object detection (Yu et al., 2022).

In this work, we aim to discover symmetries in latent space by applying the LieGAN framework (Yang et al., 2023) on latent representations learned by an encoder, and thus extract a latent space that both possess ideal symmetry properties and faithfully describe the original observations. This allows us to further expand the search space beyond linear symmetries. The non-linearity introduced by the encoder network may be able to capture the complicated mapping from the observations with largely arbitrary coordinate system to the intrinsic state variables. Then, it becomes possible for symmetry discovery models like LieGAN to find linear group actions on such a proper latent space. Such discovery can then be applied in equivariant learning through data augmentation procedure (Benton et al., 2020) or equivariant networks (Finzi et al., 2021; Yang et al., 2023) as downstream layers.

From a broader view, symmetry discovery with deep learning is just an example of the computational approaches to scientific discovery. Another major goal of this field is to discover the governing equations from observations. Many works have focused on identifying the governing equations in latent space where simple models with parsimonious representations are present (Champion et al., 2019; Fries et al., 2022). However, the structure of the latent space remain unconstrained in these works, so that their physical validity is not guaranteed. Latent domain symmetry can act as a regularizer in equation discovery, ensuring that the latent state variables carry some realistic physical significance. To this end, we aim to combine the symmetry discovery and equation discovery methods such as SINDy Autoencoder (Champion et al., 2019) in latent domain and see whether symmetries can also refine the discovery of equations.

2. Method

Given the observation space $\mathcal{X} \subseteq \mathbb{R}^n$ and observations $x_i \sim p_{\mathcal{X}}(x)$, we want to find a group $G \subset \mathrm{GL}(k)$, a low-dimensional latent space $\mathcal{Z} \subseteq \mathbb{R}^m$ and mapping from the observation space to the latent space, $\phi: \mathcal{X} \to \mathcal{Z}$, so that (1) ϕ maximally preserves the information in observations, i.e. there exists an approximate inverse map $\varphi: \mathcal{Z} \to \mathcal{X}$ s.t. $\varphi(\phi(x)) \approx x$; (2) $p_{\mathcal{Z}}(\phi(x))$ is invariant under some G-action on $\mathcal{Z}$, $\pi: \mathrm{GL}(k) \to \mathrm{GL}(m)$, i.e. $p_{\mathcal{Z}}(\phi(x)) \approx p_{\mathcal{Z}}(\pi(g)\phi(x))$, $\forall g, x$.

Figure 1 shows the structure of our proposed framework. We use an autoencoder to learn the mapping between the input space and the latent space, and the LieGAN framework (Yang et al., 2023) to detect latent space symmetries in the form of Lie algebra. Concretely, the model learns the Lie algebra generators $\{L_i \in \mathbb{R}^{k\times k}\}_{i=1}^c$ and generates group elements by sampling the linear combination coefficients $w_i \in \mathbb{R}$ for the Lie algebra basis:

$$w_i \sim \gamma(w),\ g = \exp\big[\sum_i w_i L_i\big] \quad (1)$$

[1]University of California San Diego [2]IBM Research [3]Northeastern University. Correspondence to: Rose Yu <roseyu@ucsd.edu>.

Presented at the 2^{nd} Annual Workshop on Topology, Algebra, and Geometry in Machine Learning (TAG-ML) at the 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA. 2023.

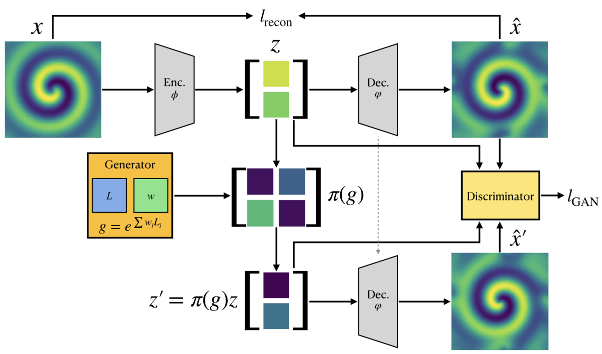

Figure 1. Structure of latent LieGAN. The encoder maps the high-dimensional input to a low-dimensional latent space. The latent representation is then transformed with the linear group action produced by LieGAN generator. The decoder reconstructs the inputs from the original and the transformed representations. Finally, the pairs of latent space representations and input space reconstructions are fed to the discriminator, whose job is to recognize the difference between the original and the transformed samples.

where γ is a chosen distribution for the coefficients and $\exp$ denotes the matrix exponential. Through adversarial training, the generator learns a group of symmetry transformations that minimizes the difference between the original and the transformed latent representations.

The latent dimension m, the group dimension k and the group action ρ are selected based on tasks. The search space of symmetry group G can also be reduced if needed.

The loss can then be formulated as

$$l_{\text{total}} = w_{\text{GAN}} \cdot l_{\text{GAN}} + w_{\text{recon}} \cdot l_{\text{recon}} \tag{2}$$

$$l_{\text{GAN}} = \mathbb{E}_{x,g}\Big[\log D(\phi(x)) + \log(1 - D(T(\phi(x))))\Big] \tag{3}$$

$$l_{\text{recon}} = \mathbb{E}_x \|\varphi(\phi(x)) - x\|^2 \tag{4}$$

where T and D denote the transformation generator and discriminator in LieGAN.

Our method is analogous to latent space equation discovery techniques (Champion et al., 2019) in terms of using an autoencoder network for nonlinear coordinate transformations. In fact, our method can be jointly trained with the objective of equation discovery simply by adding the corresponding loss terms. Concretely, if we want to find a governing equation in the latent space parameterized by θ: $\dot{z} = F_\theta(z)$, where $z = \phi(x)$, we can add the following term to the joint objective:

$$l_{\text{eq}} = \mathbb{E}_{x,\dot{x}} \|(\nabla_x z)\dot{x} - F_\theta(z)\|^2 \tag{5}$$

While equation discovery and symmetry discovery are two different objectives, we will show in the experiment that learning a symmetric latent space can simplify its structure and improve the performance of equation discovery.

3. Experiment: Reaction-Diffusion System

Many high-dimensional datasets in practical engineering and science problems derive from dynamical systems governed by partial differential equations (PDE). As an example, we consider a $\lambda - \omega$ reaction-diffusion system governed by

$$\begin{aligned} u_t &= (1 - (u^2 + v^2))u + \beta(u^2 + v^2)v + d_1(u_{xx} + u_{yy}) \\ v_t &= -\beta(u^2 + v^2)u + (1 - (u^2 + v^2))v + d_2(u_{xx} + u_{yy}) \end{aligned} \tag{6}$$

with $d_1 = d_2 = 0.1$ and $\beta = 1$. We discretize the 2D space into a 100×100 grid, which leads to an input dimension of 10^4. Figure 2a visualizes a snapshot of this system. We simulate the system up to $T = 5000$ timesteps with step size $\Delta t = 0.05$.

We train the latent LieGAN to learn low-dimensional latent representations for the high-dimensional snapshots, along with the symmetry in the latent space. We are interested in the equivariance of latent dynamics, i.e. $z_{t+1} = f(z_t) \Rightarrow gz_{t+1} = f(gz_t)$. Therefore, we take two consecutive snapshots as input, encode them to latent representations with the same encoder weights and apply the same transformations with LieGAN. We also combine our model with SINDy autoencoder (Champion et al., 2019) simply by adding the SINDy loss terms, in order to perform equation discovery in the learned latent space.

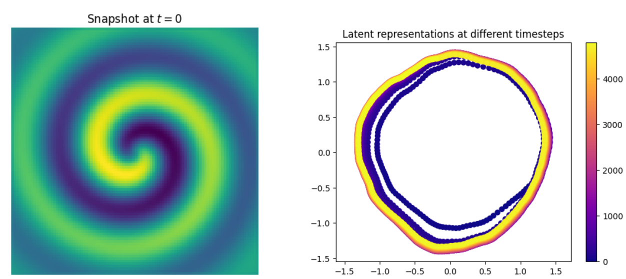

(a) (b)

Figure 2. Reaction-diffusion with 2D latent space.

We first set the latent dimension $m = 2$, which is also the setting in Champion et al. (2019). Figure 2b shows how the system evolves in the latent space throughout $T = 5000$ timesteps. The Lie algebra basis discovered by LieGAN is $L = \begin{bmatrix} 0.058 & -3.074 \\ 3.047 & -0.043 \end{bmatrix}$. This suggests $\mathrm{SO}(2)$ symmetry in the latent space, which is also evident from the visualization.

For equation discovery, we apply SINDy with up to second order polynomials as candidate functions with and without LieGAN. The results are shown in table 1. With the $\mathrm{SO}(2)$ symmetry introduced by LieGAN, the discovered equation has a more symmetric form, and is more desirable in terms of interpretation.

Model	Discovered equation
LieGAN + SINDy	$\dot{z}_1 = 0.91 z_2$ $\dot{z}_2 = -0.91 z_1$
SINDy (Champion et al., 2019)	$\dot{z}_1 = -0.85 z_2$ $\dot{z}_2 = 0.97 z_1$

Table 1. Equation discovery for reaction-diffusion system.

This system is known to have 2 intrinsic dimensions. In practice, estimating the intrinsic dimension of a high-dimensional datasets remains a challenging problem, and it is not always possible to choose the perfect latent dimension. To study the behavior of our symmetry discovery model and SINDy equation discovery method under a less ideal hyperparameter configuration, we set the latent dimension to $m = 3$ and perform the same experiments as above.

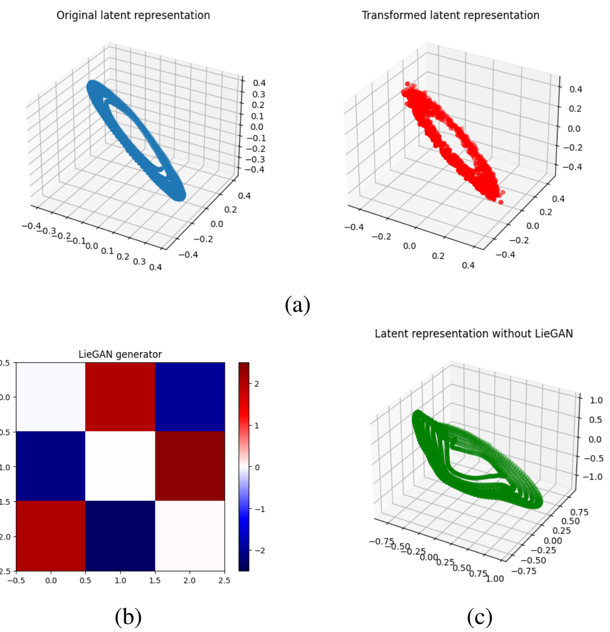

Figure 3. Reaction-diffusion with 3D latent space. (a) The latent space representations before and after LieGAN transformations. (b) LieGAN discovers the rotation symmetry around a particular axis. (c) The discovered latent space with SINDy equation discovery but without LieGAN.

Figure 3 shows the discovery results. The Lie algebra representation displayed in figure 3b is skew-symmetric, which indicates the symmetry of rotations around a particular axis. This can be easily confirmed from figure 3a, where all the latent representations mapped from the original high-dimensional inputs roughly dwell on a circular 2D subspace. In contrast, the latent space learned without LieGAN ends up with a more complicated structure, as shown in figure 3c.

The equation discovery results for this setting are listed in table 2, where we also apply SINDy with up to second order polynomials as candidate functions. The governing equation has only first-order terms in the symmetric latent space learned from LieGAN. On the other hand, applying SINDy autoencoder alone results in a highly complicated governing equation with second-order terms.

Model	Discovered equation
LieGAN +SINDy	$\dot{z}_1 = 0.43 z_2 - 0.53 z_3$ $\dot{z}_2 = -0.51 z_1 + 0.66 z_3$ $\dot{z}_3 = 0.47 z_1 - 0.52 z_2$
SINDy	$\dot{z}_1 = 0.65 z_2 - 0.16 z_3 + 0.20 z_1^2 + 0.11 z_1 z_2 + 0.29 z_1 z_3 - 0.41 z_2 z_3 - 0.16 z_3^2$ $\dot{z}_2 = -0.57 z_1 + 0.18 z_2 - 0.24 z_1 z_2 + 0.46 z_1 z_3 - 0.18 z_2^2 - 0.26 z_2 z_3 + 0.29 z_3^2$ $\dot{z}_3 = 0.45 z_1 - 0.57 z_2 - 0.27 z_1^2 + 0.18 z_2^2 - 0.19 z_2 z_3$

Table 2. Equation discovery on 3D latent space.

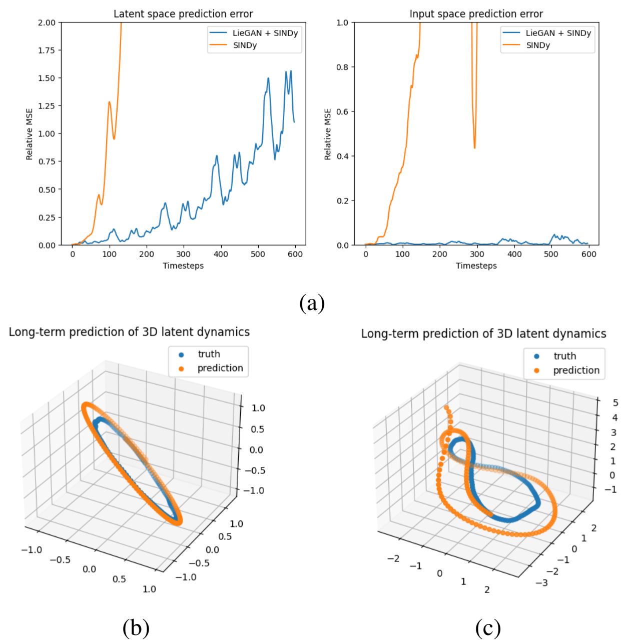

Figure 4. Long-term simulation based on the discovered equations. (a) Relative MSE losses at different timesteps. For LieGAN, the prediction error increases slowly in latent space, but remains low in input space. (b-c) Simulation trajectory in the latent space learned with and without LieGAN.

One may argue that the simplicity of an equation is not equivalent to its correctness or accuracy in terms of modeling. Indeed, there exist multiple feasible ways to encode the input to a low-dimensional latent space. To verify the accuracy of the discovered equations, we use these equations to simulate the dynamics in the latent space. Concretely, given the initial input frame x_0, we obtain its latent representation $\hat{z}_0 = \phi(x_0)$ and predict the future T timesteps by iteratively computing $\hat{z}_{t+1} = \hat{z}_t + F(\hat{z}_t) \cdot \Delta t$, where $\dot{z} = F(z)$ denotes the discovered governing equation. We then map the representations back to the input images by $\hat{x}_t = \varphi(\hat{z}_t)$. Then, we calculate the relative mean square error between the prediction and ground truth in both latent space and input space, as is shown in figure 4a. For LieGAN, the prediction error increases slowly in latent space, but remains close to zero in input space. Figure 4b shows that the simulated trajectory is approximately circular and close to ground truth. On the other hand, if SINDy autoencoder is trained without LieGAN in this case, the simulated trajectory quickly diverges from the true representations.

4. Conclusion & Future Work

In this work, we develop a method to discover nonlinear symmetries from data by applying LieGAN (Yang et al., 2023) on latent representations learned by an autoencoder. We show that such generalization from linear symmetries to nonlinear ones allows us to capture the intrinsic symmetry in high-dimensional observations. Also, our method can be jointly used with SINDy equation discovery to extract governing equations with simpler forms and better long-term prediction accuracy.

In the future, we will extend the experiment to other dynamical systems to further demonstrate the ability of our method to discover unknown symmetries. We also plan to study how the knowledge of symmetry can be better incorporated in the equation discovery process. For example, symmetry can act as a constraint to compress the search space for equations and accelerate the search. In general, our goal is to provide a general framework for detecting arbitrary symmetries in data which can also be useful in miscellaneous downstream tasks, such as data-driven prediction, discovery of other scientific properties, etc.

References

Bekkers, E. J. B-spline cnns on lie groups. *arXiv preprint arXiv:1909.12057*, 2019.

Benton, G., Finzi, M., Izmailov, P., and Wilson, A. G. Learning invariances in neural networks from training data. *Advances in neural information processing systems*, 33: 17605–17616, 2020.

Bronstein, M. M., Bruna, J., Cohen, T., and Veličković, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

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

Cohen, T., Weiler, M., Kicanaoglu, B., and Welling, M. Gauge equivariant convolutional networks and the icosahedral cnn. In *International conference on Machine learning*, pp. 1321–1330. PMLR, 2019a.

Cohen, T. S., Geiger, M., and Weiler, M. A general theory of equivariant cnns on homogeneous spaces. *Advances in neural information processing systems*, 32, 2019b.

Finzi, M., Stanton, S., Izmailov, P., and Wilson, A. G. Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. In *International Conference on Machine Learning*, pp. 3165–3176. PMLR, 2020.

Finzi, M., Welling, M., and Wilson, A. G. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. In *International Conference on Machine Learning*, pp. 3318–3328. PMLR, 2021.

Fries, W. D., He, X., and Choi, Y. LaSDI: Parametric latent space dynamics identification. *Computer Methods in Applied Mechanics and Engineering*, 399:115436, sep 2022. doi: 10.1016/j.cma.2022.115436.

Kondor, R. and Trivedi, S. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, pp. 2747–2755. PMLR, 2018.

Romero, D. W. and Lohit, S. Learning partial equivariances from data. *arXiv preprint arXiv:2110.10211*, 2021.

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

Weiler, M. and Cesa, G. General e (2)-equivariant steerable cnns. *Advances in Neural Information Processing Systems*, 32, 2019.

Yang, J., Walters, R., Dehmamy, N., and Yu, R. Generative adversarial symmetry discovery, 2023.

Yu, H.-X., Wu, J., and Yi, L. Rotationally equivariant 3d object detection, 2022.

Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. *Advances in neural information processing systems*, 30, 2017.

Zhou, A., Knowles, T., and Finn, C. Meta-learning symmetries by reparameterization. *arXiv preprint arXiv:2007.02933*, 2020.