The Surprising Effectiveness of Equivariant Models in Domains with Latent Symmetry

Dian Wang WANG.DIAN@NORTHEASTERN.EDU
Jung Yeon Park PARK.JUNGY@NORTHEASTERN.EDU
Neel Sortur SORTUR.N@NORTHEASTERN.EDU
Lawson L.S. Wong L.WONG@NORTHEASTERN.EDU
Robin Walters* R.WALTERS@NORTHEASTERN.EDU
Robert Platt* R.PLATT@NORTHEASTERN.EDU
Northeastern University

Editors: List of editors' names

Abstract

Extensive work has demonstrated that equivariant neural networks can significantly improve sample efficiency and generalization by enforcing an inductive bias in the network architecture. These applications typically assume that the domain symmetry is fully described by explicit transformations of the model inputs and outputs. However, many real-life applications contain only latent or partial symmetries which cannot be easily described by simple transformations of the input. In these cases, it is necessary to *learn* symmetry in the environment instead of imposing it mathematically on the network architecture. We discover, surprisingly, that imposing equivariance constraints that do not exactly match the domain symmetry is very helpful in learning the true symmetry in the environment. We differentiate between *extrinsic* and *incorrect* symmetry constraints and show that while imposing incorrect symmetry can impede the model's performance, imposing extrinsic symmetry can actually improve performance. We demonstrate that an equivariant model can significantly outperform non-equivariant methods on domains with latent symmetries.
Keywords: Equivariant Learning, Reinforcement Learning, Robotics

1. Introduction

Recently, equivariant learning has shown great success in various machine learning domains like trajectory prediction (Walters et al., 2020), robotics (Simeonov et al., 2022), and reinforcement learning (Wang et al., 2022c). Equivariant networks (Cohen and Welling, 2016, 2017) can improve generalization and sample efficiency during learning by encoding task symmetries directly into the model structure. However, this requires problem symmetries to be perfectly known and modeled at design time – something that is sometimes problematic. It is often the case that the designer knows that a latent symmetry is present in the problem but cannot easily express how that symmetry acts in the input space. For example, Figure 1b is a rotation of Figure 1a. However, this is not a rotation of the image – it is a rotation of the objects present in the image when they are viewed from an oblique angle. In this situation, the conventional wisdom would be to discard the model structure altogether since it is not fully known and to use an unconstrained model. Instead, we explore whether

* Equal Advising

it is possible to benefit from equivariant models even when the way a symmetry acts on the problem input is not precisely known. We show empirically that this is indeed the case and that an inaccurate equivariant model is often better than a completely unstructured model.

This paper makes two contributions. First, we define three different relationships between problem symmetry and model symmetry: *correct equivariance*, *incorrect equivariance*, and *extrinsic equivariance*. Correct equivariance means the model correctly models the problem symmetry; incorrect equivariance is when the model symmetry interferes with the problem symmetry; and extrinsic equivariance is when the model symmetry transforms the input data to out-of-distribution data. Second, we empirically compare extrinsic and incorrect equivariance in a supervised learning task and show that a model with extrinsic equivariance can improve performance compared with an unconstrained model. Finally, we explore this idea in a reinforcement learning context and show that an extrinsically constrained model can outperform state-of-the-art conventional CNN baselines.

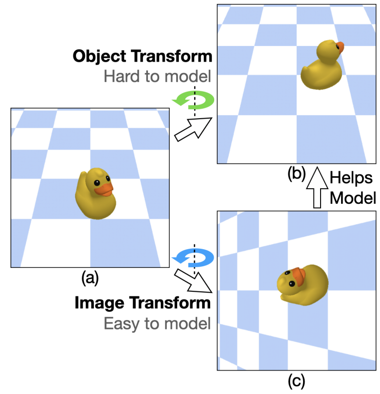

Figure 1: Object transform rotates the object itself (b), while image transform rotates the image (c).

2. Learning Symmetry Using Other Symmetries

2.1. Model Symmetry Versus True Symmetry

This paper focuses on tasks where the way in which the symmetry group operates on the input space is unknown. In this case the ground truth function $f : X \to Y$ is equivariant with respect to a group G which acts on X and Y by ρ_x and ρ_y respectively. However, *the action ρ_x on the input space is not known* and may not be a simple or explicit map. Since ρ_x is unknown, we cannot pursue the strategy of learning f using an equivariant model class f_ϕ constrained by ρ_x. As an alternative, we propose restricting to a model class f_ϕ which satisfies equivariance *with respect to a different group action* $\hat{\rho}_x$, i.e., $f_\phi(\hat{\rho}_x(g)x) = \rho_y(g)f_\phi(x)$. This paper tests the hypothesis that if the model is constrained to a symmetry class $\hat{\rho}_x$ which is related to the true symmetry ρ_x, then it may help learn a model satisfying the true symmetry. For example, if x is an image viewed from an oblique angle and ρ_x is the rotation of the objects in the image, $\hat{\rho}_x$ can be the rotation of the whole image (which is different from ρ_x because of the tilted view angle). Section 2.3 will describe this example in detail.

2.2. Correct, Incorrect, and Extrinsic Equivariance

Our findings show that the success of this strategy depends on how $\hat{\rho}_x$ relates to the ground truth function f and its symmetry. We classify the model symmetry as *correct equivariance*, *incorrect equivariance*, or *extrinsic equivariance* with respect to f. Correct symmetry means that the model symmetry correctly reflects a symmetry present in the ground truth function f. An extrinsic symmetry may still aid learning whereas an incorrect symmetry is necessarily detrimental to learning. We illustrate the distinction with a classification example shown in Figure 2a. Let $D \subseteq X$ be the support of the input distribution for f.

Definition 1 *The action $\hat{\rho}_x$ has* correct equivariance *with respect to f if $\hat{\rho}_x(g)x \in D$ for all $x \in D, g \in G$ and $f(\hat{\rho}_x(g)x) = \rho_y(g)f(x)$.*

The action $\hat{\rho}_x$ has incorrect equivariance *with respect to f if there exist $x \in D$ and $g \in G$ such that $\hat{\rho}_x(g)x \in D$ but $f(\hat{\rho}_x(g)x) \neq \rho_y(g)f(x)$.*

The action $\hat{\rho}_x$ has extrinsic equivariance *with respect to f if for $x \in D$, $\hat{\rho}_x(g)x \notin D$.*

When $\hat{\rho}_x$ has correct equivariance, the model symmetry preserves the input space D and f is equivariant with respect to it. For example, consider the action $\hat{\rho}_x$ of the group $G_1 = C_2$ acting on $\mathbb{R}^2$ by reflection across the horizontal axis and $\rho_y = 1$, the trivial action fixing labels. Figure 2b shows the untransformed data $x \in D$ as circles along the unit circle. The transformed data $\hat{\rho}_x(g)x$ (shown as crosses) also lie on the unit circle, and hence the support D is reflection invariant. Moreover, the ground truth labels $f(x)$ (shown as orange or blue) are preserved by this action.

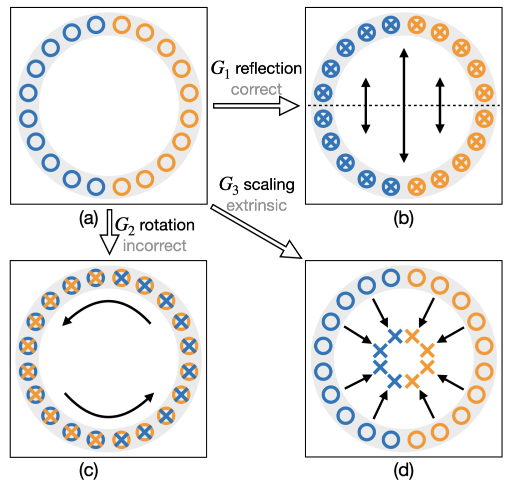

Figure 2: An example for correct, incorrect, and extrinsic equivariance. Grey is the input distribution. Circles are the data in distribution where the color shows the ground truth label. Crosses show the group transformed data.

When $\hat{\rho}_x$ has incorrect equivariance, the model symmetry partially preserves the input distribution, but does not correctly preserve labels. In Figure 2c, the rotation group $G_2 = \langle \mathrm{Rot}_\pi \rangle$ maps the unit circle to itself, but the transformed data does not have the correct label. Thus, constraining the model f_ϕ by $f_\phi(\hat{\rho}_x(g)x) = f_\phi(x)$ will force f_ϕ to mislabel data. In this example, for $a = \sqrt{2}/2$, $f(a, a) = \textsc{orange}$ and $f(-a, -a) = \textsc{blue}$, however, $f_\phi(a, a) = f_\phi(\mathrm{Rot}_\pi(a, a)) = f_\phi(-a, -a)$.

Extrinsic equivariance is when the equivariant constraint in the equivariant network f_ϕ enforces equivariance to out-of-distribution data. Since $\hat{\rho}_x(g)x \notin D$, the ground truth $f(\hat{\rho}_x(g)x)$ is undefined. An example of extrinsic equivariance is given by the scaling group G_3 shown in Figure 2d. For the data $x \in D$, enforcing scaling invariance $f_\phi(\hat{\rho}_x(g)x) = f_\phi(x)$ where $g \in G_3$ will not increase error, because the group transformed data (in crosses) are out of the distribution D of the input data shown in the grey ring. In fact, we hypothesize that such extrinsic equivariance may even be helpful for the network to learn the ground truth function.

2.3. Object Transformation and Image Transformation

In tasks with visual inputs ($X = \mathbb{R}^{c \times h \times w}$), incorrect or extrinsic equivariance will exist when the transformation of the image does not match the transformation of the latent state of the task. In such case, we call ρ_x the *object transform* and $\hat{\rho}_x$ the *image transform*. For an image input $x \in X$, the image transform $\hat{\rho}_x(g)x$ is defined as a simple transformation

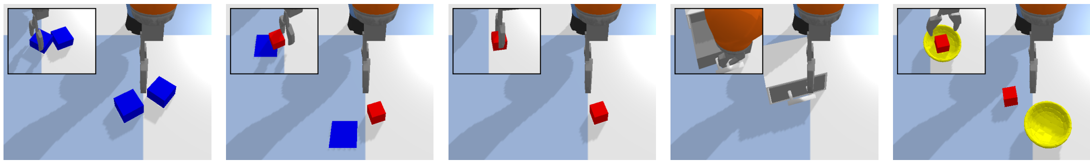

(a) Block Pulling (b) Block Pushing (c) Block Picking (d) Drawer Opening (e) Block in Bowl

Figure 3: The manipulation environments from BulletArm benchmark Wang et al. (2022b) implemented in PyBullet. The top-left shows the goal for each task.

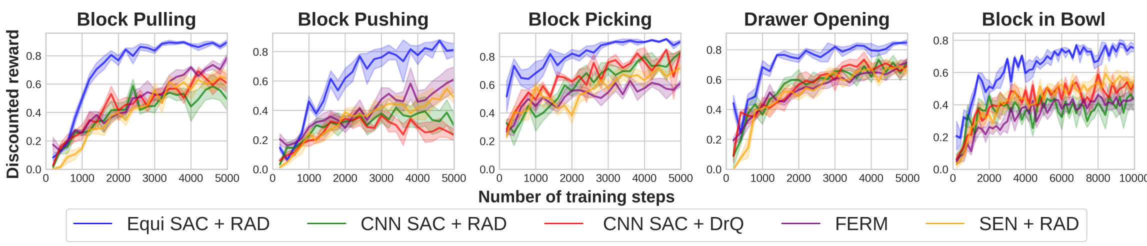

Figure 4: Comparison of Equivariant SAC (blue) with baselines in evaluation. The evaluation is performed every 200 training steps.

of pixel locations (e.g., Figure 1a-c), while the object transform $\rho_x(g)x$ is an implicit map transforming the objects in the image (e.g., Figure 1a-b).

3. Extrinsic Equivariance in Robotic Reinforcement Learning

In this section, we demonstrate that extrinsic equivariance can significantly improve sample efficiency in reinforcement learning.

We experiment in five robotic manipulation environments shown in Figure 3. The state space $S = \mathbb{R}^{4\times h\times w}$ is a 4-channel RGBD image captured from a fixed camera pointed at the workspace. The action space $A = \mathbb{R}^5$ is the change in gripper pose (x, y, z, θ), where θ is the rotation along the z-axis, and the gripper open width λ. The task has latent O(2) symmetry: when a rotation or reflection is applied to the poses of the gripper and the objects, the action should rotate and reflect accordingly. However, such symmetry does not exist in image space because the image perspective is skewed instead of top-down. We enforce such extrinsic symmetry (group D_4) using Equivariant SAC (Wang et al., 2022c,a) equipped with random crop augmentation using RAD (Laskin et al., 2020) (Equi SAC + RAD) and compare it with a number of sample-efficient baselines using unconstrained CNN. All methods use Prioritized Experience Replay (PER) (Schaul et al., 2015) with pre-loaded expert demonstrations. We also add an L2 loss towards the expert action in the actor to encourage expert actions.

Figure 4 shows that Equivariant SAC (blue) outperforms all baselines, suggesting that an equivariant model is still powerful when the symmetry is extrinsic.

4. Conclusion

This paper defines correct equivariance, incorrect equivariance, and extrinsic equivariance, and demonstrates experimentally that extrinsic equivariance can provide significant performance improvements in reinforcement learning. A limitation of this paper is that we focus on planar equivariant networks. In future work, we are interested in evaluating extrinsic equivariance in network architectures that process different types of data.

Acknowledgments

This work is supported in part by NSF 1724257, NSF 1724191, NSF 1763878, NSF 1750649, NSF 2107256, and NASA 80NSSC19K1474. R. Walters is supported by the Roux Institute and the Harold Alfond Foundation and NSF 2134178.

References

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016.

Taco S. Cohen and Max Welling. Steerable CNNs. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=rJQKYt5ll.

Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *Advances in neural information processing systems*, 33:19884–19895, 2020.

Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.

Anthony Simeonov, Yilun Du, Andrea Tagliasacchi, Joshua B Tenenbaum, Alberto Rodriguez, Pulkit Agrawal, and Vincent Sitzmann. Neural descriptor fields: Se (3)-equivariant object representations for manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6394–6400. IEEE, 2022.

Robin Walters, Jinxi Li, and Rose Yu. Trajectory prediction using equivariant continuous convolution. *arXiv preprint arXiv:2010.11344*, 2020.

Dian Wang, Mingxi Jia, Xupeng Zhu, Robin Walters, and Robert Platt. On-robot learning with equivariant models. In *6th Annual Conference on Robot Learning*, 2022a. URL https://openreview.net/forum?id=K8W6ObPZQyh.

Dian Wang, Colin Kohler, Xupeng Zhu, Mingxi Jia, and Robert Platt. Bulletarm: An open-source robotic manipulation benchmark and learning framework. *arXiv preprint arXiv:2205.14292*, 2022b.

Dian Wang, Robin Walters, and Robert Platt. SO(2)-equivariant reinforcement learning. In *International Conference on Learning Representations*, 2022c. URL https://openreview.net/forum?id=7F9cOhdvfk_.