Cross-space Symmetry Composition in Robotics

Loizos Hadjiloizou, Rodrigo Pérez-Dattari, and Noémie Jaquier

***Abstract*— Robots exhibit a rich variety of symmetries arising from their mechanical properties and their assigned tasks. Although many manipulation tasks exhibit several symmetries simultaneously, most works exploit them in isolation, failing to exploit their combined potential. This paper proposes a cross-space symmetry composition framework to learn robot policies that are simultaneously equivariant to several symmetries. Our framework considers symmetries arising in the robot configuration and task spaces and builds on the geometric properties of forward kinematics map to transfer them into a common space, where they are subsequently composed. We validate our framework on a simulated dual-arm letter-writing task, demonstrating that jointly leveraging multiple symmetries yields improved generalization.**

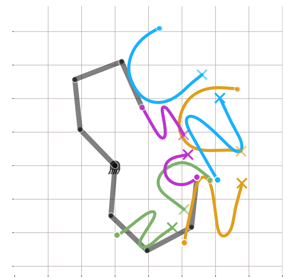

Fig. 1: Dual-arm robot tracing letters (—) and their symmetric transformations via rotation (—), rotation and scaling (—), rotation, scaling and morphological reflection (—).

I. Introduction

Incorporating inductive bias is a key avenue to improve sample efficiency and generalization in robot policy learning. Among possible inductive biases, *symmetries* are particularly powerful as they allow a policy to generalize from each training sample to all its symmetric counterparts [1]. Intuitively, a symmetry is a transformation that preserves some property of interest. They can be grounded on physics [2] or engineered from invariances detected in the data [3].

Robots exhibit a rich variety of symmetries rooted in their mechanical structure and assigned tasks. In configuration space, the momenta of floating-base robots are preserved under rigid body transformations, allowing policies to be designed on lower-dimensional spaces [2], [4]. Morphological symmetries [5] capture the interchangeability between replicated kinematic chains, e.g., the arms of a humanoid. In task space, Euclidean isometries have been widely used to learn equivariant manipulation policies that generalize across unseen object poses [6], [7], [8]. Task-induced symmetries [9] emerge directly from the problem itself, where all states achieving the same objective are considered symmetric and treated equivalently.

In many manipulation tasks, robots exhibit several of the aforementioned symmetries simultaneously, see Fig. 1. Yet, most works consider them in isolation [4], [5], [7], [9], thereby overlooking relevant inductive bias. Compositions of symmetries have been explored by the machine learning community. Kim *et al.* [10] proposed to weakly enforce multiple symmetries by adding multiple regularizers to the objective function. In a similar line, Li *et al.* propose to learn dynamic models via latent flows that preserve several symmetries simultaneously. In both works, the considered symmetries arise in a single space. Instead, the key challenge in robotics is that symmetries emerge in different spaces: Floating-base and morphological symmetries naturally arise in the robot's configuration space, while task isometries and task-induced symmetries are expressed in task space.

In this paper, we contribute: *(1)* A symmetry transfer framework that maps symmetries between configuration and task spaces along with feasibility conditions (Sec. III-A); *(2)* A framework for composing symmetries within a given space including a characterization of the necessary compatibility conditions (Sec. III-B). We validate our approach in simulated robotics experiments, showcasing that jointly leveraging multiple symmetries leads to improved generalizability.

II. Preliminaries

A. Groups, Actions, and Equivariance

Given a manifold $\mathcal{M}$, a symmetry can be described as the intrinsic property of some quantity $\mathcal{C}:\mathcal{M}\to\mathbb{R}$ to remain constant after applying a map $\Phi:\mathcal{M}\to\mathcal{M}$, i.e., $\mathcal{C}(\boldsymbol{p})=\mathcal{C}(\Phi(\boldsymbol{p}))$ $\forall\boldsymbol{p}\in\mathcal{M}$. In practice, the mathematical structure of *groups* provides a powerful framework for representing symmetries via their *group action* Φ. A group $\mathbb{G}$ is a set equipped with an associative binary operation $\circ:\mathbb{G}\times\mathbb{G}\to\mathbb{G}$, an identity element $e\in\mathbb{G}$, and an inverse $g^{-1}\in\mathbb{G}$ for each element g [11]. Groups transform elements of arbitrary manifolds $\mathcal{M}$ via their action Φ. We focus on the *left action*, which is an associative map $\Phi^{\mathcal{M}}:\mathbb{G}\times\mathcal{M}\to\mathcal{M}$ that respects the identity transformation $\Phi^{\mathcal{M}}(e,\boldsymbol{p})=\boldsymbol{p}$ $\forall\boldsymbol{p}\in\mathcal{M}$. A map $h:\mathcal{M}\to\mathcal{N}$ across two manifolds $\mathcal{M},\mathcal{N}$ is $\mathbb{G}$-*equivariant* if

$$h\big(\Phi^{\mathcal{M}}(g,\boldsymbol{p})\big)=\Phi^{\mathcal{N}}\big(g,h(\boldsymbol{p})\big)\quad\forall\boldsymbol{p}\in\mathcal{M}. \tag{1}$$

Intuitively, a map is $\mathbb{G}$-equivariant when transforming the input before applying the map is equivalent to first applying the map before transforming the output. The *equivariance condition* (1) underpins equivariant (symmetric) learning models

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program funded by the Knut and Alice Wallenberg Foundation.

The authors are with the Department of Robotics, Perception and Learning at KTH Royal Institute of Technology, Stockholm, Sweden.

in the literature, used to perform data augmentation [12], penalize symmetry violations via regularization [13], or enforce a symmetry directly in the network architecture [14].

Lie Groups. Robotics symmetries involve both finite groups, arising from discrete symmetries, e.g., morphological symmetries and *Lie groups* associated with continuous symmetries, e.g., floating-base symmetries. A Lie group $\mathbb{G}$ is a group which is also a smooth manifold. Its tangent space at the identity $\mathfrak{g} = \mathcal{T}_e\mathbb{G}$, known as the Lie algebra, provides an infinitesimal description of $\mathbb{G}$. The Lie group and its Lie algebra are linked by the exponential map $\exp : \mathfrak{g} \to \mathbb{G}$. When working with robotic systems, it is often required to operate at the velocity level, i.e., on the tangent spaces of $\mathcal{M}$. A Lie group acts on the tangent spaces of $\mathcal{M}$ via the differential $\mathrm{d}\Phi_g^{\mathcal{M}}|_{\boldsymbol{p}} : \mathcal{T}_{\boldsymbol{p}}\mathcal{M} \to \mathcal{T}_{\Phi^{\mathcal{M}}(g,\boldsymbol{p})}\mathcal{M}$ with $g \in \mathbb{G}$, $\boldsymbol{p} \in \mathcal{M}$.

Infinitesimal description of the action. Lie groups symmetries can be analyzed from an infinitesimal perspective via vector fields on $\mathcal{M}$. Here, the Lie algebra $\mathfrak{g}$ acts on $\mathcal{M}$ via *infinitesimal generators*. For each $\boldsymbol{\xi} \in \mathfrak{g}$, such generators correspond to the vector fields

$$X_{\boldsymbol{\xi}}^{\mathcal{M}}(\boldsymbol{p}) = \frac{d}{dt}\Big|_{t=0} \Phi^{\mathcal{M}}(\exp(t\boldsymbol{\xi}), \boldsymbol{p}). \tag{2}$$

Integrating along $X_{\boldsymbol{\xi}}^{\mathcal{M}}$ is equivalent to applying the actions $\Phi^{\mathcal{M}}(\exp(t\boldsymbol{\xi}), \boldsymbol{p})$. For instance, SO(2) acting on $\mathbb{R}^2$ induces rotational vector fields tangent to circles around the origin.

B. Riemannian Manifolds and Maps

A Riemannian manifold is a smooth manifold $\mathcal{M}$ equipped with an inner product $m : \mathcal{T}_{\boldsymbol{p}}\mathcal{M} \times \mathcal{T}_{\boldsymbol{p}}\mathcal{M} \to \mathbb{R}$ varying smoothly with $\boldsymbol{p} \in \mathcal{M}$, called a Riemannian metric. A map $f : \mathcal{M} \to \mathcal{N}$ between two Riemannian manifolds is a *Riemannian submersion* at point $\boldsymbol{p} \in \mathcal{M}$ if its differential $\mathrm{d}f|_{\boldsymbol{p}} : \mathcal{T}_{\boldsymbol{p}}\mathcal{M} \to \mathcal{T}_{f(\boldsymbol{p})}\mathcal{N}$ is surjective. The Riemannian submersion induces an orthogonal decomposition of each tangent space $\mathcal{T}_{\boldsymbol{p}}\mathcal{M}$ into a vertical subspace $\mathcal{V}_{\boldsymbol{p}} = \mathrm{Ker}(\mathrm{d}f)|_{\boldsymbol{p}} = \{\boldsymbol{v} \in \mathcal{T}_{\boldsymbol{q}}\mathcal{M} : \mathrm{d}f|_{\boldsymbol{p}}(\boldsymbol{v}) = 0\}$ and a horizontal subspace $\mathcal{H}_{\boldsymbol{p}} = (\mathcal{V}_{\boldsymbol{p}})^{\perp}$ [15].

The robot configuration space $(\mathcal{Q}, m_{\mathcal{Q}})$ and task space $(\mathcal{X}, m_{\mathcal{X}})$ are Riemannian manifolds whose elements represent the joint configurations and end-effector poses, respectively. Both spaces are related through the forward kinematics map $f : \mathcal{Q} \to \mathcal{X}$. In coordinates, its differential $\mathrm{d}f|_{\boldsymbol{q}} : \mathcal{T}_{\boldsymbol{q}}\mathcal{Q} \to \mathcal{T}_{f(\boldsymbol{q})}\mathcal{N}$ corresponds to the robot Jacobian matrix $\boldsymbol{J}(\boldsymbol{q})$. As shown in [16], the forward kinematics map of redundant robots is a Riemannian submersion. The vertical subspace of each $\mathcal{T}_{\boldsymbol{q}}\mathcal{Q}$ is the nullspace of the Jacobian and contains the joint motions that do not move the end-effector.

III. Cross-Space Symmetry Composition

In this section, we consider robot manipulation tasks exhibiting several symmetries both in the configuration space $\mathcal{Q}$ and the task space $\mathcal{X}$. We aim to learn an equivariant policy with respect to all symmetries at hand. In other words, the policy must satisfy the equivariance condition (1) for a set of group actions $[\Phi_1^{\mathcal{Q}}, \ldots, \Phi_M^{\mathcal{Q}}, \Phi_1^{\mathcal{X}}, \ldots, \Phi_N^{\mathcal{X}}]$. To do so, we propose to transfer all symmetries in a common space where they are subsequently composed.

A. Transferring Symmetries across Spaces

Transferring a symmetry across spaces entails preserving the effect of its action in all spaces in which it is applied. Next, we describe the transfer of symmetries from $\mathcal{Q}$ to $\mathcal{X}$ and vice versa via the forward kinematics map $f : \mathcal{Q} \to \mathcal{X}$. For redundant robots (our use case) $\dim(\mathcal{Q}) > \dim(\mathcal{X})$.

Descending Symmetries. We consider a symmetry associated with a group $\mathbb{G}$ acting on the configuration space $\mathcal{Q}$ via $\Phi^{\mathcal{Q}} : \mathbb{G} \times \mathcal{Q} \to \mathcal{Q}$. The symmetry can be *descended* from the configuration space onto the task space, if the forward kinematics map f is $\mathbb{G}$-equivariant, i.e., if f satisfies (1) for all group actions $\Phi^{\mathcal{Q}}$. Then, the task space action $\Phi^{\mathcal{X}}$ is uniquely determined by $\Phi^{\mathcal{Q}}$ and f through (1), such that

$$\Phi^{\mathcal{X}}(g, \boldsymbol{x}) = \Phi^{\mathcal{X}}(g, f(\boldsymbol{q})) = f(\Phi^{\mathcal{Q}}(g, \boldsymbol{q})). \tag{3}$$

An example of symmetries that can be descended are morphological symmetries $\mathbb{G}_{\mathrm{M}}$, which permute the state of kinematically-identical chains (e.g., the left and right arms of a humanoid) in the robot configuration space [5]. For $g_{\mathrm{M}} \in \mathbb{G}_{\mathrm{M}}$, the action $\Phi_M^{\mathcal{Q}}$ on the configuration manifold is $\Phi_M^{\mathcal{Q}}(g_{\mathrm{M}}, \boldsymbol{q}) = \rho_{\mathcal{Q}_J}(g_{\mathrm{M}})\boldsymbol{q}$, where $\rho_{\mathcal{Q}_J}(g_{\mathrm{M}})$ denotes the permutation of kinematic chains induced by g_{M}. As the forward kinematics map f is known to be $\mathbb{G}_{\mathrm{M}}-$equivariant [5], morphological symmetries descend naturally and the descended group action $\Phi_{\mathrm{M}}^{\mathcal{X}} : \mathbb{G}_{\mathrm{M}} \times \mathcal{X} \to \mathcal{X}$ takes the form

$$\Phi_{\mathrm{M}}^{\mathcal{X}}(g_{\mathrm{M}}, \boldsymbol{x}) = \rho_{\mathcal{X}}(g_{\mathrm{M}})\boldsymbol{x}, \tag{4}$$

where $\rho_{\mathcal{X}}(g_{\mathrm{M}})$ is the permutation of the end-effector poses.

Lifting Symmetries. We now proceed to study the *lifting* of symmetries, i.e., the transfer of task-space symmetries to the configuration space. We consider a symmetry associated with the group $\mathbb{G}$, acting on the task space $\mathcal{X}$ as $\Phi^{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \to \mathcal{X}$ and seek a corresponding action in $\mathcal{Q}$ as $\Phi^{\mathcal{Q}} : \mathbb{G} \times \mathcal{Q} \to \mathcal{Q}$ such that a given symmetry in $\mathcal{X}$ is preserved. For redundant robots, $\Phi^{\mathcal{Q}}$ cannot be constructed directly from (1) due to the non-uniqueness of the inverse kinematics map f^{-1}. Instead, we propose to leverage the geometric properties of f to lift the symmetries via their infinitesimal generators.

Specifically, as f is a Riemannian submersion [16], any smooth vector field $X^{\mathcal{X}}$ on $\mathcal{X}$ has a unique smooth horizontal lift to $\mathcal{Q}$ [17, Prop. 2.25], i.e., an associated vector field $X^{\mathcal{Q}}$ lying on the horizontal space $\mathcal{H}_{\boldsymbol{q}}$. As shown in [16, Thm. III.2], the horizontal lift is given by

$$X^{\mathcal{Q}}(\boldsymbol{q}) = (\mathrm{d}f|_{\mathcal{H}_{\boldsymbol{q}}})^{-1}(X^{\mathcal{X}}(f(\boldsymbol{q}))), \tag{5}$$

with $(\mathrm{d}f|_{\mathcal{H}_{\boldsymbol{q}}})^{-1}$ corresponding, in coordinates, to the Moore-Penrose pseudo-inverse of the Jacobian matrix $\boldsymbol{J}(\boldsymbol{q})$.

Using (5), we obtain the infinitesimal generator $X_{\boldsymbol{\xi}}^{\mathcal{Q}}$ of $\mathbb{G}$ acting on $\mathcal{Q}$ as the horizontal lift of the infinitesimal generator $X_{\boldsymbol{\xi}}^{\mathcal{X}}$. Next, we aim to construct the group action $\Phi_{\boldsymbol{\xi}}^{\mathcal{Q}}$ from $X_{\boldsymbol{\xi}}^{\mathcal{Q}}$. To do so, we need to ensure that moving along $X_{\boldsymbol{\xi}}^{\mathcal{Q}}$ yields the same task-space motion as flowing along $X_{\boldsymbol{\xi}}^{\mathcal{X}}$. By the naturality of flows [15, Proposition 9.13], the vector field correspondence extends to their flows, i.e., to their group actions $\Phi_{\boldsymbol{\xi}}^{\mathcal{Q}}, \Phi_{\boldsymbol{\xi}}^{\mathcal{X}}$. Therefore, we obtain the unique group action of $\mathbb{G}$ on $\mathcal{Q}$ by using (2) on $X_{\boldsymbol{\xi}}^{\mathcal{Q}}$.

As an example, we consider rotational invariance in task space. The symmetry is described by the actions $\Phi_{\mathrm{R}}^{\mathcal{X}}$ of $\mathbb{G}_{\mathrm{R}} = \mathrm{SO}(3)$ on $\mathcal{X}$ expressed via rotation matrices as $\Phi_{\mathrm{R}}^{\mathcal{X}}(g_{\mathrm{R}}, \boldsymbol{x}) = \boldsymbol{R}(g_{\mathrm{R}})\boldsymbol{x}$. The corresponding infinitesimal generator $X_{\mathrm{R}}^{\mathcal{X}}$ is constructed using (2). It is then lifted to $\mathcal{Q}$ via (5), yielding $X_{\mathrm{R}}^{\mathcal{Q}}$, and subsequently integrated to obtain the action $\Phi_{\mathrm{R}}^{\mathcal{Q}}(t)$ of $\mathbb{G}_{\mathrm{R}}$ on $\mathcal{Q}$. Importantly, the integration horizon t of $X_{\mathrm{R}}^{\mathcal{Q}}$ determines the resulting action. In this case, increasing the integration horizon leads to a larger rotation.

B. Composing Symmetries

In this section, we propose to combine the symmetries in the common space (either $\mathcal{Q}$ or $\mathcal{X}$) and discuss the necessary conditions to do so. Two groups $\mathbb{G}_1, \mathbb{G}_2$ acting on $\mathcal{Q}$ via Φ_1 and Φ_2 can be composed into the new group $\mathbb{G} = \mathbb{G}_1 \times \mathbb{G}_2 = \{g = (g_1, g_2) | g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2\}$ via the direct product $\times$ if their actions commute, i.e.,

$$[\Phi_1(g_1, \boldsymbol{q}), \Phi_2(g_1, \boldsymbol{q})] = 0, \forall g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2, \boldsymbol{q} \in \mathcal{Q}. \quad (6)$$

The composed group action Φ is given as

$$\Phi(g, \boldsymbol{q}) = \Phi_1(g_1, \Phi_2(g_2, \boldsymbol{q})) = \Phi_2(g_2, \Phi_1(g_1, \boldsymbol{q})). \quad (7)$$

Intuitively, it does not depend on the order in which the actions are applied. In the case where the group actions do not commute, i.e., (6) does not hold but $\mathbb{G}2$ acts on $\mathbb{G}1$ via automorphisms, the groups $\mathbb{G}_1, \mathbb{G}_2$ can be composed via the semi-direct product $\mathbb{G} = \mathbb{G}_1 \ltimes \mathbb{G}_2$ [11], leading to the composed action

$$\Phi(g, \boldsymbol{q}) = \Phi_1(g_1, \Phi_2(\rho(g_1)(g_2), \boldsymbol{q})). \quad (8)$$

In this case, the order of the actions matters as one group action influences the other.

IV. Experimental Validation

We validate our framework on a dual-arm letter-tracing task in simulation. We consider a fixed-base planar dual-arm robot with two 4-DoF arms. The robot's configuration space is $\mathcal{Q} = \mathcal{Q}_{\mathrm{l}} \times \mathcal{Q}_{\mathrm{r}} = \mathbb{T}^4 \times \mathbb{T}^4$ and its task space is $\mathcal{X} = \mathcal{X}_{\mathrm{l}} \times \mathcal{X}_{\mathrm{r}} = \mathbb{R}^2 \times \mathbb{R}^2$. The task of the robot is to simultaneously trace the letters C and N with its left and right arms. We aim to learn a policy that follows these letters and symmetric variations thereof.

We train a goal-conditioned policy $\pi : \mathcal{Q} \times \mathcal{X} \to \mathcal{TQ}$ that takes $\boldsymbol{q} \in \mathcal{Q}$ and the last position along the desired letter trajectory $\boldsymbol{x}^* \in \mathcal{X}$ as inputs, and outputs a velocity $\dot{\boldsymbol{q}} \in \mathcal{T}_{\boldsymbol{q}}\mathcal{Q}$. The joint-space representation of the demonstrations is obtained via inverse kinematics and the network is trained via imitation learning as in [18] using 7 demonstrations.

We enforce one configuration- and two task-space symmetries in π: *(1)* A bilateral *morphological* symmetry described by the cyclic-two group $\mathbb{G}_{\mathrm{M}}^{\mathcal{Q}} = \mathbb{C}_2$; *(2)* A *rotational* symmetry under the planar rotation group $\mathbb{G}_{\mathrm{R}}^{\mathcal{X}} = \mathrm{SO}(2)$; and *(3)* A *task-induced* scaling symmetry $\mathbb{G}_{\mathrm{S}}^{\mathcal{X}} = \mathbb{S}_2$. Following Sec. III, all symmetries are transferred to and composed in $\mathcal{Q}$. We incorporate the symmetries into the policy via data augmentation procedure where the original demonstrations are transformed using the composed group actions, leading to trajectories that are scaled, rotated, and reflected between the robot's arms. We ablate the composition of symmetries by considering four policies: *(1)* π, vanilla without any symmetry; *(2)* $\pi_{\mathbb{G}_{\mathrm{R}}}$, with rotational symmetry; *(3)* $\pi_{\mathbb{G}_{\mathrm{RT}}}$, with rotational and task-induced symmetry; and *(4)* $\pi_{\mathbb{G}_{\mathrm{MRT}}}$, with all three symmetries.

TABLE I: RMSE of trained models on test datasets augmented with an increasing number of composed symmetries.

	Test datasets			
Policy	Original	$\mathbb{G}_{\mathrm{R}}$	$\mathbb{G}_{\mathrm{RT}}$	$\mathbb{G}_{\mathrm{MRT}}$
π	$\mathbf{0.026}_{\pm \mathbf{0.027}}$	$20.927_{\pm 11.822}$	$25.754_{\pm 8.326}$	$11.656_{\pm 9.593}$
$\pi_{\mathbb{G}_{\mathrm{R}}}$	$0.044_{\pm 0.033}$	$\mathbf{0.048}_{\pm \mathbf{0.037}}$	$16.622_{\pm 17.780}$	$19.771_{\pm 11.575}$
$\pi_{\mathbb{G}_{\mathrm{RT}}}$	$0.050_{\pm 0.067}$	$0.075_{\pm 0.108}$	$0.036_{\pm 0.073}$	$15.690_{\pm 10.160}$
$\pi_{\mathbb{G}_{\mathrm{MRT}}}$	$0.073_{\pm 0.06}$	$0.064_{\pm 0.049}$	$\mathbf{0.035}_{\pm \mathbf{0.028}}$	$\mathbf{0.046}_{\pm \mathbf{0.037}}$

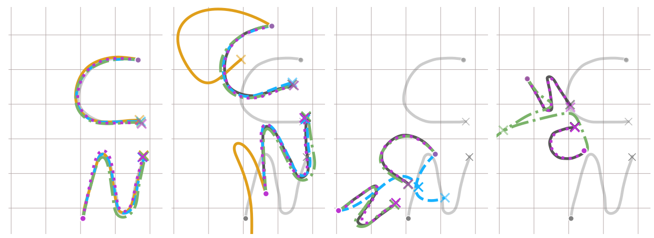

(a) Original (b) $\mathbb{G}_{\mathrm{R}}$ (c) $\mathbb{G}_{\mathrm{RT}}$ (d) $\mathbb{G}_{\mathrm{MRT}}$

Fig. 2: Policies π (—), $\pi_{\mathbb{G}_{\mathrm{R}}}$ (—), $\pi_{\mathbb{G}_{\mathrm{RT}}}$ (—), $\pi_{\mathbb{G}_{\mathrm{MRT}}}$ (—) evaluated on test trajectories (—) obtained as symmetric transformations of the original demonstrations (—). Trajectories' start and end are depicted as $\cdot$ and $\times$. Diverging models are excluded from the visualization of the next composition.

Table I shows the RMSE between the nominal and predicted trajectories starting from the same initial condition. Nominal trajectories correspond to both the original demonstrations and their symmetric transformations. We observe that, in every case, policies benefit from incorporating the symmetries present in the dataset: When a policy did not encode a symmetry contained in the data, its performance degrades. Moreover, $\pi_{\mathbb{G}_{\mathrm{MRT}}}$ performs well across all datasets.

Fig. 2 depicts trajectories of the learned models under multiple symmetric transformations. We observe that leveraging multiple symmetries yields increased generalizability. In particular, π only reproduces the original demonstrations, π_{R} generalizes to rotated letters but fails on scaling and morphological reflections, and π_{RT} generalizes to simultaneously rotated and scaled trajectories. Finally, π_{MRT} successfully generalizes across all four conditions.

V. Conclusion

We presented a framework for transferring and composing configuration- and task-space symmetries of redundant robots by leveraging the Riemannian submersion structure of the forward kinematics map. Our experiments on a simulated letter-writing task showcased the generalization benefits of robot policies considering multiple symmetries simultaneously via data augmentation. Future work will explore alternative approaches for enforcing composed symmetries.

References

[1] S. E. Otto, N. Zolman, J. N. Kutz, and S. L. Brunton, "A Unified Framework to Enforce, Discover, and Promote Symmetry in Machine Learning," *Journal of Machine Learning Research*, vol. 26, no. 248, pp. 1–83, 2025.

[2] J. E. Marsden and T. S. Ratiu, *Introduction to Mechanics and Symmetry*. Springer, 1999, vol. 2.

[3] K. Desai, B. Nachman, and J. Thaler, "Symmetry Discovery with Deep Learning," *Physical Review D*, vol. 105, no. 9, p. 096031, 2022.

[4] H. Mishra, G. Garofalo, A. M. Giordano, M. De Stefano, C. Ott, and A. Kugi, "Reduced Euler-Lagrange Equations of Floating-Base Robots: Computation, Properties, & Applications," *IEEE Trans. on Robotics*, vol. 39, no. 2, pp. 1439–1457, 2023.

[5] D. O. Apraez, G. Turrisi, V. Kostic, M. Martin, A. Agudo, F. Moreno-Noguer, M. Pontil, C. Semini, and C. Mastalli, "Morphological Symmetries in Robotics," *Intl. Journal of Robotics Research*, vol. 44, no. 10-11, pp. 1743–1766, 2025.

[6] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani *et al.*, "Transporter Networks: Rearranging the Visual World for Robotic Manipulation," in *Conference on Robot Learning (CoRL)*, 2021, pp. 726–747.

[7] Z. Su, X. Huang, D. Ordoñez-Apraez, Y. Li, Z. Li, Q. Liao, G. Turrisi, M. Pontil, C. Semini, Y. Wu *et al.*, "Leveraging Symmetry in RL-based Legged Locomotion Control," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2024, pp. 6899–6906.

[8] R. Hyunwoo, L. Hong-in, L. Jeong-Hoon, and C. Jongeun, "Equivariant Descriptor Fields: SE(3)-Equivariant Energy-Based Models for End-to-End Visual Robotic Manipulation Learning," in *Intl. Conf. on Learning Representations (ICLR)*, 2023.

[9] V. Kallem, D. E. Chang, and N. J. Cowan, "Task-Induced Symmetry and Reduction with Application to Needle Steering," *IEEE Transactions on Automatic Control*, vol. 55, no. 3, pp. 664–673, 2010.

[10] H. Kim, H. Lee, H. Yang, and J. Lee, "Regularizing Towards Soft Equivariance Under Mixed Symmetries," in *Intl. Conf. on Machine Learning (ICML)*, 2023.

[11] B. C. Hall, *Lie Groups, Lie Algebras, and Representations*, 2nd ed. Springer, 2015, vol. 2.

[12] J. Luo, D. Chen, and Q. Zhang, "Reinforcement Learning with Euclidean Data Augmentation for State-Based Continuous Control," in *Neural Information Processing Systems (NeurIPS)*, 2024, pp. 90 253–90 276.

[13] R. Wang, R. Walters, and R. Yu, "Approximately Equivariant Networks for Imperfectly Symmetric Dynamics," in *Intl. Conf. on Machine Learning (ICML)*, 2022, pp. 23 078–23 091.

[14] J. Seo, S. Yoo, J. Chang, H. An, H. Ryu, S. Lee, A. Kruthiventy, J. Choi, and R. Horowitz, "SE(3)-Equivariant Robot Learning and Control: A Tutorial Survey," *Intl. Journal of Control, Automation and Systems*, vol. 23, no. 5, pp. 1271–1306, 2025.

[15] J. M. Lee, *Introduction to Smooth Manifolds*. Springer, 2012, vol. 218.

[16] C. Altafini, "Redundant Robotic Chains on Riemannian Submersions," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 335–340, 2004.

[17] J. M. Lee, *Introduction to Riemannian Manifolds*. Springer, 2018, vol. 2.

[18] R. Pérez-Dattari and J. Kober, "Stable Motion Primitives via Imitation and Contrastive Learning," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3909–3928, 2023.