Contact-rich SE(3)-Equivariant Robot Manipulation Task Learning via Geometric Impedance Control

Joohwan Seo[1], Nikhil P. S. Prakash[1], Xiang Zhang[1], Changhao Wang[1], Jongeun Choi[1,2], Masayoshi Tomizuka[1], and Roberto Horowitz[1]

***Abstract*—This paper presents a differential geometric control approach that leverages $SE(3)$ group invariance and equivariance to increase transferability in learning robot manipulation tasks that involve interaction with the environment. The proposed approach is based on utilizing a recently presented geometric impedance control (GIC) combined with a learning variable impedance control framework, where the gain scheduling policy is trained in a supervised learning fashion from expert demonstrations. A geometrically consistent error vector (GCEV) is fed to a neural network to achieve a gain scheduling policy. We proved that the GIC and learning representation using GCEV remain invariant under arbitrary $SE(3)$ transformations, i.e., translation and rotations. Furthermore, we show that the proposed approach is equivariant when represented relative to the spatial frame. A comparison of our proposed control and learning framework with a well-known Cartesian space learning impedance control, equipped with a Cartesian error vector-based gain scheduling policy, confirms the significantly superior learning transferability of our proposed approach.**

***Index Terms*—Geometric Impedance Control, $SE(3)$ Equivariance and Left-invariance, Variable Impedance Control, Contact-rich Manipulation Tasks**

I. Introduction

To achieve full autonomy when executing tasks that require contact-rich interactions, such as assembly and construction tasks, robots need to robustly perceive where and how a task must be performed. Vision systems produce 2D images or 3D point clouds that allow robots to perceive and understand their environment. With these vision systems, recent vision-based robotic manipulations focus on finding the desired end-effector poses [1], [2], [3]. These works for vision-based robotic manipulations are based on the philosophy that "every manipulation task is a sequence of pick-and-places" [4].

Geometric deep learning (GDL) [5], [6] has been recently adopted in the robotics society and has shown higher generalizability and sample efficiency performances. The GDL is employed to deal with vision-based inputs and outputs desired end-effector poses by exploiting group invariance and equivariance [7], [8], [9], [10], [11]. However, these vision-based robotic manipulations deal with simple tabletop manipulation tasks represented by pick-and-places that typically do not require sophisticated control laws as the contact interaction between the end-effector and the environment is minimized.

In this paper, we leverage $SE(3)$ group invariance and equivariance, extensively studied in GDL, to enhance learning transferability and sampling efficiency in contact-rich robotic manipulation tasks. A geometric learning variable impedance control is presented that utilizes the GIC in [12], [13] and incorporates a learning variable impedance control framework [14], [15], [16], where the gain scheduling policy is trained in a supervised learning fashion from expert demonstrations. A geometrically consistent error vector (GCEV) is fed to a neural network to achieve a gain scheduling policy that remains invariant under arbitrary $SE(3)$ transformations. Our results demonstrate that the SE(3) equivariance at the force (or control) level is the key to achieving high sample efficiency and learning transferability in peg-in-hole (PiH) tasks undergoing $SE(3)$ transformation.

Looking ahead to the next step in vision-based robotic manipulation, let's assume that the desired end-effector pose is obtained through $SE(3)$ equivariant vision modules, such as those described in [9], [10]. These modules make the vision system both sample-efficient and robust to $SE(3)$ transformations. This leads to a crucial question: **Does the control law underlying these tasks also need to be equivariant to be effective?** Our comparison of geometric and Cartesian space-based formulations reveals that even with an $SE(3)$ equivariant vision module, the entire pipeline cannot accomplish the contact-rich and direction-sensitive tasks that demand precise force interaction without the force-level equivariance.

This workshop paper is a short version of [17], where a more detailed exploration of these concepts can be found.

II. Preliminaries

We are interested in the control problem of the manipulator's end-effector. The configuration of the manipulator's end-effector can be defined by its position and orientation, and the configuration manifold lies in the Special Euclidean group $SE(3)$. Thus, we rely on the differential geometric formulation of $SE(3)$ manifold to control and formulate the problem. For the details of the formulation and their notations, the readers are referred to [17], [12], [13]. As our control and benchmark control law formulations play crucial roles in the paper, we will briefly introduce those in this section.

A. Geometric Impedance Control Law

We employ the geometric impedance control (GIC) law proposed in [12]. We first note that g denotes the current configuration matrix, with p and R representing the current position and rotation matrix, i.e., $g = (R, p)$. In a similar

[1]Joohwan Seo, Nikhil P.S. Prakash, Xiang Zhang, Changhao Wang, Jongeun Choi, Masayoshi Tomizuka, and Roberto Horowitz are with UC Berkeley, Department of Mechanical Engineering. {joohwan_seo, nikhilps, xiang_zhang_98, changhaowang, tomizuka, horowitz} @berkeley.edu

[2]Jongeun Choi is with Yonsei University, School of Mechanical Engineering. jongeunchoi@yonsei.ac.kr

way, $g_d = (R_d, p_d)$ denotes the desired current configuration matrix. In a nutshell, a GIC control law $\tilde{T} \in se^*(3)$ in the wrench for a fixed PiH task, is given by

$$\tilde{T} = -f_{\scriptscriptstyle G} - K_d V^b + \tilde{G}, \tag{1}$$

where $\tilde{C}$, and $\tilde{G}$ are matrices in the operational space formulation, K_d is a symmetric positive definite damping matrix, $f_{\scriptscriptstyle G}$ is the elastic generalized force in $se^*(3)$ and V^b is a body-frame velocity, e.g., $V^b = J_b(q)\dot{q}$ with q is joint coordinate and $J_b(q)$ is body-frame Jacobian matrix, both of them described on the body-frame, which will be described subsequently. The elastic geometric wrench (force) $f_{\scriptscriptstyle G}(g, g_d)$ is given by

$$f_{\scriptscriptstyle G}(g, g_d) = \begin{bmatrix} f_p \\ f_R \end{bmatrix} = \begin{bmatrix} R^T R_d K_p R_d^T (p - p_d) \\ (K_R R_d^T R - R^T R_d K_R)^\vee \end{bmatrix}, \tag{2}$$

where K_p and K_R denote symmetric positive definite stiffness matrices in translation and rotation, respectively. Note that the control law (1) can be interpreted as a PD control together with gravity compensation. For more details on the GIC, such as derivation and stability properties, we refer to [12], [13] and its references [18], [19]. Here, we define a geometrically consistent error vector (GCEV) $e_{\scriptscriptstyle G}$, which will be utilized in the learning impedance gains later in this paper. The GCEV $e_{\scriptscriptstyle G}(g, g_d)$ is defined as follows:

$$e_{\scriptscriptstyle G}(g, g_d) = \begin{bmatrix} e_p \\ e_R \end{bmatrix} = \begin{bmatrix} R^T (p - p_d) \\ (R_d^T R - R^T R_d)^\vee \end{bmatrix} \in \mathbb{R}^6. \tag{3}$$

Finally, we note that all the vectors/wrench, such as $e_{\scriptscriptstyle G}$, V^b, and $f_{\scriptscriptstyle G}$, are described in the body-frame coordinate unless otherwise specified.

B. Cartesian space Impedance Control

As a benchmark approach for the proposed GIC, we also briefly introduce a Cartesian space Impedance Controller (CIC), which is a currently standard method for impedance control. In the operational space formulation, correctly representing the rotational dynamics has always received significant interest [20]. In particular, a positional Cartesian error vector (CEV) $e_{\scriptscriptstyle C}$ widely utilized in CIC can be defined in the following way [21], [22], [23].

$$e_{\scriptscriptstyle C} = [(e_{\scriptscriptstyle C,p})^T,\ (e_{\scriptscriptstyle C,R})^T]^T, \quad \text{where} \tag{4}$$
$$e_{\scriptscriptstyle C,p} = p - p_d, \quad e_{\scriptscriptstyle C,R} = (r_{d_1} \times r_1 + r_{d_2} \times r_2 + r_{d_3} \times r_3),$$

with $R = [r_1, r_2, r_3]$ and $R_d = [r_{d_1}, r_{d_2}, r_{d_3}]$. Utilizing the positional error vector in Cartesian space and considering a fair comparison with the GIC formulation, we will utilize the following CIC formulation for the PiH task.

$$\tilde{T}_{\scriptscriptstyle C} = -K_{\scriptscriptstyle C} e_{\scriptscriptstyle C} - K_{d_C} V^s + \tilde{G}_{\scriptscriptstyle C}, \tag{5}$$

where K_{d_C} is a damping matrix for the CIC, and $K_{\scriptscriptstyle C} = \text{blkdiag}(K_{\scriptscriptstyle C,p}, K_{\scriptscriptstyle C,R})$ with $K_{\scriptscriptstyle C,p}$ and $K_{\scriptscriptstyle C,R}$ are translational and rotational stiffness matrices, respectively. To implement control law (1) and (5), the wrenches should first be converted to joint torque T by multiplying corresponding Jacobian matrices, i.e., $T = J_b^T \tilde{T}$ and $T = J_s^T \tilde{T}_{\scriptscriptstyle C}$, respectively.

III. Problem Definition and Solution Approach

1) Overview: Our ultimate goal is to provide a learning-based solution to solve a peg-in-hole task, a classic representative of contact-rich force-based robotic manipulation tasks. We will address this problem in the framework of learning variable impedance control, where the gain scheduling policy of the impedance control laws is trained using learning algorithms. In particular, behavior cloning (BC) from expert demonstrations is utilized in a supervised learning fashion to obtain the gain scheduling policy.

To achieve this, we introduce a gain scheduling policy parameterized by a simple neural network. This neural network takes as input the positional signals representing the current end-effector pose and the desired goal pose, such as GCEV ($e_{\scriptscriptstyle G}$) and CEV ($e_{\scriptscriptstyle C}$). Formally,

$$(K_p, K_R)_t = \pi_\theta(s_t), \tag{6}$$

K_p and K_R are impedance gains defined in (2) or (5), and θ denotes parameters of the neural network-based policy π_θ. In our case, s_t denotes input to the neural network, either $e_{\scriptscriptstyle G}$ or $e_{\scriptscriptstyle C}$. We employ a standard multi-layer perceptron (MLP) as a neural network. We will call π_θ as a gain-scheduling policy and drop the subscript t for the compactness of notation. To show the effectiveness of the geometric formulation, we propose and compare two different approaches to learning variable impedance control: 1. Selection of the control rule (GIC vs CIC), 2. Selection of the states s_t ($e_{\scriptscriptstyle G}$ and $e_{\scriptscriptstyle C}$).

The performances are evaluated on four main scenarios in Fig. 1. The gain scheduling policy π_θ is trained only in the default scenario (Fig. 1(a)). The trained policy is then tested in the other scenarios (Fig. 1s(b)-(d)) to evaluate its zero-shot transferability and robustness to OOD data. The detailed environmental setup and descriptions regarding data collection can be found in [17].

A. Solution Approach

In this subsection, we will first introduce our proposed approach. In what follows, the behavior cloning (BC) to obtain the gain-scheduling policy is introduced.

1) Proposed Approach: Our proposed approach utilizes GIC with a learning impedance gain scheduling policy, where the input to the neural network s is the GCEV $e_g(g, g_d)$ defined in (3). The action from the gain scheduling policy then becomes $(K_p, K_R) = \pi_\theta(e_{\scriptscriptstyle G})$. The GIC control law (2) equipped with gain scheduling policy $\pi_\theta(e_{\scriptscriptstyle G})$ has a crucial property for learning transferability as shown in the following lemma.

Lemma 1. ***Left-invariance of the GCEV*** **(3)** ***and the elastic wrench*** **(2)**
$e_{\scriptscriptstyle G}(g, g_d)$ and $f_{\scriptscriptstyle G}(g, g_d)$ are left-invariant to the arbitrary left-transformation g_l in $SE(3)$, i.e. $\forall g_l \in SE(3)$,

$$e_{\scriptscriptstyle G}(g, g_d) = e_{\scriptscriptstyle G}(g_l g, g_l g_d), \quad f_{\scriptscriptstyle G}(g, g_d) = f_{\scriptscriptstyle G}(g_l g, g_l g_d) \tag{7}$$

Proof. See [17]. □

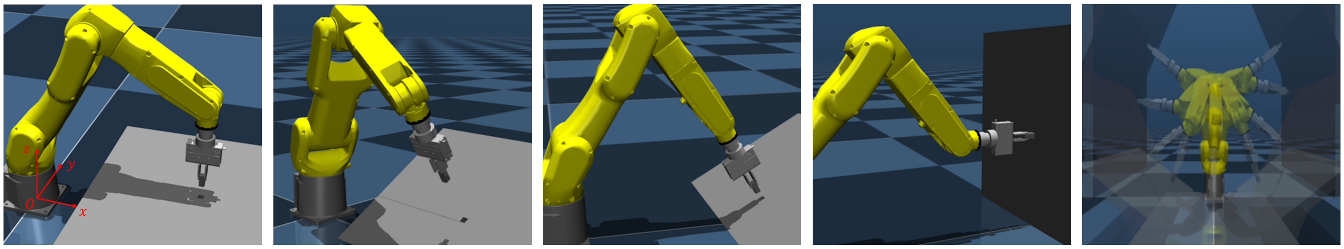

(a) Default Case (b) Case 1 (c) Case 2 (d) Case 3 (e) Case 4

Fig. 1: Robot performing a peg into a hole insertion task in different scenarios for testing learning transferability. The data is collected, and the policy is trained only in performing the task shown in (a). The trained policy is then tested on the tasks shown in (b)-(e), where the insertion hole is translated in different orientations. (b): tilted in $+x$ direction for $30°$ (c): tilted in $-y$ direction for $30°$ (d): tilted in $-y$ direction for $90°$ (e): tilted in $+x$ direction in arbitrary angle $\phi \in [-135°, 135°]$. The coordinate frame is attached to the first figure.

Domain Randomization: Domain randomization is a crucial technique in both learning manipulation tasks to enhance robustness by allowing the neural network to explore a broader range of state space [24]. In conventional impedance gains learning problem [16], domain randomization is typically applied to both the initial and goal end-effector poses. However, under the proposed GIC framework, domain randomization is only necessary for the initial pose of the end-effector, as demonstrated in the following proposition.

Proposition 1. *For the learning variable impedance control problem based on GIC law* (1)*, the following equation holds true:*

$$f_G(g, g_l g_d) = f_G(g_l^{-1} g, g_d) \tag{8}$$

Proof. See [17]. □

The effects of Proposition 1 on learning strategies are as follows. First, since the main driving force of (1) with gain scheduling policy is $f_G(g, g_d)$ in (2), we focus on the properties of $f_G(g, g_d)$. The domain randomization on the target pose can be represented by $f_G(g, g_l g_d)$ where $g_l \in SE(3)$ is arbitrary. Then, the result of Proposition 1 reads that

$$f_G(g, g_l g_d) = f_G(g_l^{-1} g, g_d)$$

Note that following the axioms of groups (Chap 2.1, [25]), g_l^{-1} can be denoted by another group element g_l' since g_l is arbitrary. Finally, $f_G(g_l' g, g_d)$ means that the domain randomization on the target pose of the end-effector is identical to randomization on the initial pose. Therefore, independent domain randomizations on both the target pose and the initial pose are not necessary during the training process. Only the initial pose relative to the target pose needs to be randomized, or vice versa, but not both. This result is crucial in robotics applications where the cost of collecting data in the real world is substantial.

IV. Experiments and Discussions

A. Behavior Cloning (BC) result

The results of the BC experiments are presented in Table. I. Each task was tested 100 times, and the success cases were counted. The BC policy was trained with the GCEV and executed with the GIC (GIC+GCEV), and the trained policy was successfully transferred to the other tasks without a significant drop in the success rate. However, the BC policy trained with a CEV and executed with CIC (CIC+CEV) failed to transfer the trained policy, resulting in a dramatic decrease in the success rate. The reason for this difference in transferability can be attributed to the error vector representation. The relationship between left invariance and transferability is further explained in Remark 1.

TABLE I: Success rates of the BC policies for the proposed and the benchmark approaches (Tested 100 times each, Values in Percentage %)

Method	Default	Case 1	Case 2	Case 3
BC Proposed (GIC+GCEV)	100	99	95	100
BC Benchmark (CIC+CEV)	100	0	0	1
BC Mixed 1 (GIC+CEV)	99	54	49	27
BC Mixed 2 (CIC+GCEV)	100	0	0	0

Remark 1. ***Why does left-invariance matter?***
The left-invariance of the error vector implies that the chosen error vector is invariant to the selection of the global coordinate system. In this paper, we interpret left-invariance in a slightly different manner. Consider the situation where the desired and current configurations are transformed through a left action of the SE(3) group, which corresponds to a change in the spatial coordinate frame – See Fig. 2. Due to the left-invariance of the GCEV (3), the error vector remains unchanged in cases (a) and (b) in Fig. 2. Therefore, from the perspective of the proposed approach, the task remains invariant to translational/rotational perturbations. In this perspective, the use of the left-invariant error vector e_G can help address distributional shift or out-of-distribution issues, as the trained policy will consistently encounter the same input e_G.

B. Left-Invariance is not enough

The question arises whether training a policy based on GCEV or equivalent left-invariant features enables transferability to translational/rotational perturbations. To answer this question, we test different combinations of trained gain scheduling policies and control methods: BC policy with

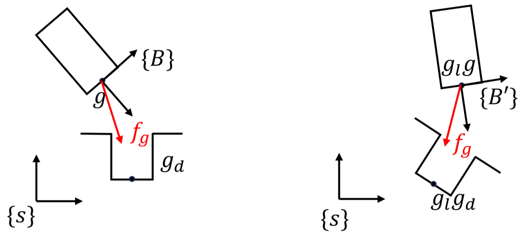

Fig. 2: (Left) Peg and hole configurations are represented by g and g_d, respectively. (Right) Peg and hole undergo left transformation via an action $g_l \in SE(3)$. The GCEV e_G (3) and elastic force f_G (2) are invariant to the left transformation of $SE(3)$. $\{s\}$ represents spatial coordinate frame, while $\{B\}$ and $\{B'\}$ represent body coordinate frames.

Cartesian error vector executed with GIC (GIC+CEV) and BC policy with a GCEV executed with CIC (CIC+GCEV). If left-invariance of the feature were the only factor for transferability, we would expect the BC policy with the Cartesian error vector executed with GIC (GIC+CEV) to not be transferable due to the lack of left-invariance in the Cartesian error vector. Conversely, the BC policy using a GCEV executed with CIC would be transferable.

However, the experimental results presented in 3rd and 4th rows of Table I contradict the hypothesis. The GIC with a gain scheduling policy trained with the Cartesian error vector (GIC + CEV) showed some transferability in Case 1 and Case 2, with success rates near 50%. However, in Case 3, with a tilt of 90 degrees, the success rate drops to 27% due to encountering an unexperienced state distribution. This suggests that a left-invariant gain scheduling policy alone is not enough for transferability.

On the other hand, the CIC using a GCEV produced results consistent with our hypothesis. Even though the gain scheduling policy for CIC was trained with a GCEV, the direction of the gains and resulting force direction are still represented in the Cartesian frame. Thus, the policy outputs the same gains it was trained on, but the resulting force direction is not suitable for tilted cases.

Consequently, it is concluded that the left-invariant gain scheduling policy alone is not enough, and a more fundamental factor is needed to address transferability - the direction of forces. Unlike CIC, the forces in GIC are defined in the body frame, resulting in the automatic change in the direction of forces (See Fig. 2, f_G on each case) – which implies an ***equivariance*** property. We state that the key to transferability lies in a control law represented in the body-frame coordinate and the left-invariant gain scheduling policy. The invariant gain scheduling policy is obtained from the neural network with a GCEV as input, while the left-invariant feedback control law is inherited from the structure of GIC. To establish a connection between our statement and the equivariance property, we first present the definition of equivariance.

Definition 1. *Consider a function $f : \mathcal{X} \to \mathcal{Y}$, i.e., $y = f(x)$ with $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. The function f is equivariant to the group g if the following condition is satisfied [6]:*

$$f(\rho^{\mathcal{X}} x) = \rho^{\mathcal{Y}} f(x) \tag{9}$$

where $\rho^{\mathcal{X}}$ represents the action of group g in the domain $\mathcal{X}$ and $\rho^{\mathcal{Y}}$ represents the action of group g in the codomain $\mathcal{Y}$.

Based on Definition 1, we propose the following proposition.

Proposition 2. *The feedback terms in GIC law* (1) *described in the body frame are equivariant if it is described in the spatial frame, i.e., $\forall g_l \in SE(3)$*

$$f^s_G(g_l g, g_l g_d) = Ad^T_{g_l} f_G(g, g_d) \tag{10}$$

Here, superscript s denotes that the given vector is represented to the spatial frame.

Proof. See [17]. □

Note that we only consider feedback terms in (1) since the feedforward terms are just employed to cancel the manipulator dynamics, not affecting the closed-loop dynamics. A remark on the Proposition 2 is provided.

Remark 2. ***Extensions to general force-based policy***
It is worth mentioning that our concept proposed in Proposition 2 can also be extended to general force-based policies. If a force-based policy is left-invariant in the body-frame, e.g., implemented using a neural network with left-invariant features as input, and described in the end-effector body frame, it will be guaranteed to be left-equivariant in the spatial frame. The GCEV (3) introduced in this paper is an example of such a left-invariant feature.

V. Conclusions and Future Works

In this paper, a geometric approach leveraging $SE(3)$ group invariance and equivariance for contact-rich robotic manipulation task learning is presented. To solve the Peg-in-Hole (PiH) task, the proposed approach builds on top of the geometric impedance control (GIC), where its impedance gains are changed via a left-invariant gain scheduling policy. Expert behavior cloning is chosen to train the gain scheduling policy. Through theoretical analysis, we prove that the proposed GIC and the geometrically consistent error vector (GCEV) used for learning are left-invariant relative to $SE(3)$ group transformations when represented in the end-effector's body frame system, enabling learning transferability. Furthermore, we show that left invariance in the body-frame representation leads to $SE(3)$ equivariance of the proposed approach when described in a spatial frame. A PiH simulation experiment confirms the learning transferability of our proposed method, which is not exhibited by the well-known Cartesian space-based benchmark approach.

For future work, we will address more realistic scenarios where the exact goal poses are unknown and need to be estimated via sensors, e.g., images or point cloud-based inputs. By combining the proposed approach with widely studied $SE(3)$ equivariant vision modules, we anticipate achieving fully end-to-end, vision-to-force $SE(3)$ equivariance.

REFERENCES

[1] P. R. Florence, L. Manuelli, and R. Tedrake, “Dense object nets: Learning dense visual object descriptors by and for robotic manipulation,” *arXiv preprint arXiv:1806.08756*, 2018.

[2] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, “kpam: Keypoint affordances for category-level robotic manipulation,” in *The International Symposium of Robotics Research*. Springer, 2019, pp. 132–157.

[3] W. Gao and R. Tedrake, “kpam 2.0: Feedback control for category-level robotic manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2962–2969, 2021.

[4] M. Shridhar, L. Manuelli, and D. Fox, “Cliport: What and where pathways for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2022, pp. 894–906.

[5] E. J. Bekkers *et al.*, “Roto-translation covariant convolutional networks for medical image analysis,” in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I*. Springer, 2018, pp. 440–448.

[6] T. Cohen and M. Welling, “Group equivariant convolutional networks,” in *International conference on machine learning*. PMLR, 2016, pp. 2990–2999.

[7] A. Zeng *et al.*, “Transporter networks: Rearranging the visual world for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2021, pp. 726–747.

[8] A. Simeonov *et al.*, “Neural descriptor fields: SE(3)-equivariant object representations for manipulation,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6394–6400.

[9] H. Ryu *et al.*, “Equivariant descriptor fields: SE(3)-equivariant energy-based models for end-to-end visual robotic manipulation learning,” in *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.

[10] ——, “Diffusion-edfs: Bi-equivariant denoising generative modeling on se (3) for visual robotic manipulation,” *arXiv preprint arXiv:2309.02685*, 2023.

[11] J. Kim *et al.*, “Robotic manipulation learning with equivariant descriptor fields: Generative modeling, bi-equivariance, steerability, and locality,” in *RSS 2023 Workshop on Symmetries in Robot Learning*, 2023.

[12] J. Seo *et al.*, “Geometric impedance control on SE(3) for robotic manipulators,” *IFAC World Congress 2023, Yokohama, Japan*, 2023.

[13] J. Seo, N. P. S. Prakash, J. Choi, and R. Horowitz, “A comparison between lie group-and lie algebra-based potential functions for geometric impedance control,” *arXiv preprint arXiv:2401.13190*, 2024.

[14] X. Zhang *et al.*, “Learning variable impedance control via inverse reinforcement learning for force-related tasks,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2225–2232, 2021.

[15] S. Kozlovsky, E. Newman, and M. Zacksenhouse, “Reinforcement learning of impedance policies for peg-in-hole tasks: Role of asymmetric matrices,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 898–10 905, 2022.

[16] C. C. Beltran-Hernandez *et al.*, “Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach,” *Applied Sciences*, vol. 10, no. 19, p. 6923, 2020.

[17] J. Seo, N. P. Prakash, X. Zhang, C. Wang, J. Choi, M. Tomizuka, and R. Horowitz, “Contact-rich se (3)-equivariant robot manipulation task learning via geometric impedance control,” *IEEE Robotics and Automation Letters*, 2023.

[18] T. Lee *et al.*, “Geometric tracking control of a quadrotor uav on SE(3),” in *49th IEEE conference on decision and control (CDC)*. IEEE, 2010, pp. 5420–5425.

[19] F. Bullo and R. M. Murray, “Tracking for fully actuated mechanical systems: a geometric framework,” *Automatica*, vol. 35, no. 1, pp. 17–34, 1999.

[20] F. Caccavale *et al.*, “Six-dof impedance control based on angle/axis representations,” *IEEE Transactions on Robotics and Automation*, vol. 15, no. 2, pp. 289–300, 1999.

[21] Y. Zhu *et al.*, “robosuite: A modular simulation framework and benchmark for robot learning,” in *arXiv preprint arXiv:2009.12293*, 2020.

[22] H. Ochoa and R. Cortesão, “Impedance control architecture for robotic-assisted mold polishing based on human demonstration,” *IEEE Transactions on Industrial Electronics*, vol. 69, no. 4, pp. 3822–3830, 2021.

[23] S. Shaw, B. Abbatematteo, and G. Konidaris, “RMPs for safe impedance control in contact-rich manipulation,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2707–2713.

[24] J. Tobin *et al.*, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.

[25] R. M. Murray, Z. Li, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 1994.