Morphologically Equivariant Flow Matching for Bimanual Mobile Manipulation

Max Siebenborn*[1], Daniel Ordoñez Apraez*[2], Sophie Lueth[1],
Giulio Turrisi[2], Massimiliano Pontil[2], Claudio Semini[2], Georgia Chalvatzaki[1,3,4]

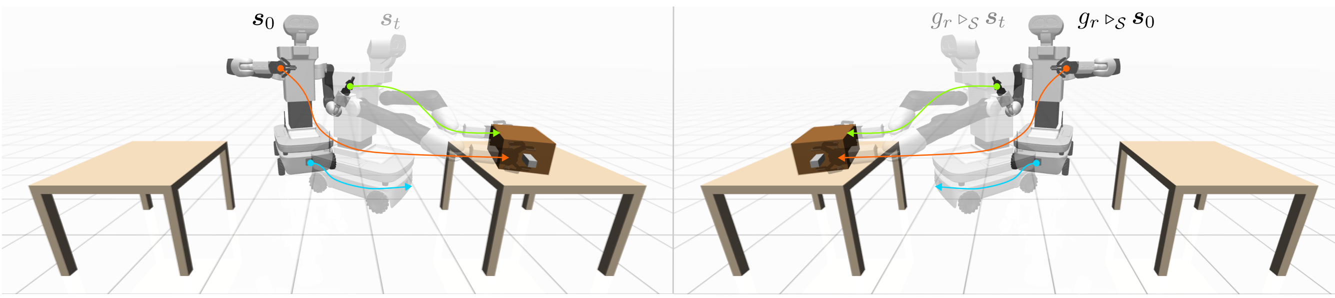

Fig. 1: Simulated bimanual *box lifting* task, illustrating reflection morphological symmetry of mobile manipulators. A successful trajectory $(\boldsymbol{s}_t, \boldsymbol{a}_t)_{t=0..T}$ (*left*) transfers zero-shot to the mirrored setting: starting from the reflected initial state $g_r \triangleright_{\mathcal{S}} \boldsymbol{s}_0$, executing the reflected action sequence $(g_r \triangleright_{\mathcal{A}} \boldsymbol{a}_t)_{t=0..T}$ produces the trajectory $(g_r \triangleright_{\mathcal{S}} \boldsymbol{s}_t, g_r \triangleright_{\mathcal{A}} \boldsymbol{a}_t)_{t=0..T}$ that solves the task in the mirrored setting (*right*). In this paper, we leverage this reflection symmetry prior for behavior cloning with flow matching, demonstrating improved sample efficiency, generalization, and optimality in bimanual mobile manipulation.

***Abstract*— Mobile Manipulation requires coordinated control of high-dimensional, bimanual robots. Imitation Learning methods have been broadly used to solve these robotic tasks, yet typically ignore the bilateral morphological symmetry inherent in such systems. We argue that morphological symmetry is an underexplored but crucial inductive bias for learning in bimanual Mobile Manipulation: knowing how to solve a task in one configuration directly determines how to solve its mirrored counterpart. In this paper, we formalize this symmetry prior and show that it constrains optimal bimanual policies to be ambidextrous and equivariant under reflections across the robot's sagittal plane. We introduce a $\mathbb{C}_2$-equivariant Flow Matching policy that enforces reflective symmetry either via a regularized training loss or an equivariant velocity network. Across planar and 6-DoF Mobile Manipulation tasks, symmetry-informed policies consistently improve sample efficiency and achieve zero-shot generalization to mirrored configurations absent from the training distribution. We further validate this zero-shot generalization capability on a real-world manipulation task with a TIAGo++ robot. Together, our findings establish morphological symmetry as an effective, generalizable, and scalable inductive bias for ambidextrous generative policy learning.**

I. Introduction

Over the last decade, deep learning has driven remarkable progress in robotics and promises to solve long-standing challenges in service robotics, including household assistance, healthcare support, and collaborative assembly. Such applications require robots to solve Mobile Manipulation (MoMa) tasks involving dextrous bimanual manipulation. Current approaches often rely on Imitation Learning (IL), using datasets of expert human demonstrations to train stochastic policies via Behavior Cloning (BC) [1]–[4]. However, these methods typically require vast amounts of data and compute due to task complexity and low sample efficiency of generative IL paradigms. Given the cost and difficulty of collecting expert demonstrations, and the risks of malfunction, methods with improved sample efficiency, generalization, and optimality are critically needed.

This paper leverages a core physics-informed inductive bias in bimanual MoMa frequently ignored in IL: the robot's bilateral morphological symmetry [5], rooted in the reflection symmetry of the left and right sides of the body morphologies of humans and most bimanual manipulators (see Fig. 1).

Intuitively, this symmetry prior implies that observations and actions at any time are related by known reflection transformations to those of the mirrored manipulation task (see Fig. 1). Hence, any optimal policy solving one task transfers directly to the reflected task, yielding an *ambidextrous* manipulation policy [6]. Figure 2 illustrates an analogous control symmetry in the planar Push-T environment [1]: any expert demonstration can be reflected across the vertical symmetry axis of the target-T to produce an expert rollout that solves the task from the reflected initial state.

In this paper, we formalize this symmetry prior for bimanual MoMa, reformulating IL as a symmetry-constrained optimization problem enforceable explicitly or via regular-

[1]Technical University of Darmstadt, [2]Istituto Italiano di Tecnologia, [3]Hessian.AI, [4]Robotics Institute Germany

*Equal contribution.

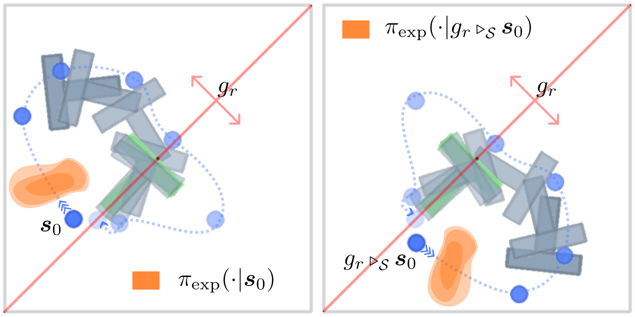

Fig. 2: Reflection symmetry of the Push-T environment [1], illustrating a rollout starting from state $\boldsymbol{s}_0$ (left) and the corresponding reflected rollout starting from the reflected state $g_r \triangleright_{\mathcal{S}} \boldsymbol{s}_0$ (right). The expert policy distribution (orange) is reflection-invariant, motivating the use of this symmetry as an inductive bias for policy learning.

ization within Flow Matching (FM). Our contributions are:

(i) we formalize reflection symmetry-induced equivariance in FM for IL and show two strategies to enforce these constraints: explicitly via equivariant Neural Networks (NNs), or softly via regularization;

(ii) we introduce a $\mathbb{G}$-equivariant transformer for arbitrary finite groups, enabling ambidextrous bimanual control;

(iii) we show that exploiting symmetry improves sample efficiency, generalization, and enables zero-shot transfer to mirrored tasks outside the training distribution.

II. Problem Formulation

We model each bimanual manipulation task as a Partially Observable Markov Decision Process (POMDP) defined by $(\mathcal{S}, \mathcal{A}, \tau, r, \mathbb{P}_0, \mathcal{O}, \phi)$, where $\boldsymbol{s} \in \mathcal{S}$ is the world state, $\boldsymbol{a} \in \mathcal{A}$ the agent's action, $\tau : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}_+$ the transition kernel, $r : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ the reward function, $\mathbb{P}_0 : \mathcal{S} \to \mathbb{R}_+$ the initial state distribution, and $o \in \mathcal{O}$ an observation with $\phi : \mathcal{O} \times \mathcal{S} \to \mathbb{R}_+$ the observation model. We aim to learn a control policy $\pi_{\boldsymbol{\theta}}$ from a dataset of N expert trajectories, $\mathbb{D}^{\text{exp}} = \{(\boldsymbol{o}_t^{(n)}, \boldsymbol{a}_t^{(n)})_{t=0}^{T_n}\}_{n=1}^{N}$, consisting of observation–action sequences generated by an expert policy π_{exp}. We thus frame policy learning as a Behavior Cloning problem to learn a parametric stochastic policy $\pi_{\boldsymbol{\theta}}(\cdot|\boldsymbol{x})$ with parameters $\boldsymbol{\theta}$ that approximates the expert policy $\pi_{\text{exp}}(\cdot|\boldsymbol{x})$, where $\boldsymbol{x}$ a history of H consecutive observations, $\boldsymbol{x}_t := [\boldsymbol{o}_{t-H+1}, \ldots, \boldsymbol{o}_t] \in \mathcal{X} \subseteq \mathcal{O}^H$. Formally, this objective minimizes the Kullback–Leibler divergence between the expert and learned policy:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathbb{D}_H^{\text{exp}}} D_{\text{KL}} \left(\pi_{\text{exp}}(\cdot|\boldsymbol{x}) \,\|\, \pi_{\boldsymbol{\theta}}(\cdot|\boldsymbol{x})\right). \tag{1}$$

Crucially, we observe that in bimanual manipulation the POMDP admits a unique reflection symmetry, denoted g_r, induced by the robot's bilateral morphological symmetry [5] (see Fig. 1). This symmetry acts *linearly* on the state, action, and observation spaces via group actions $(\triangleright_{\mathcal{S}}) : \mathbb{C}_2 \times \mathcal{S} \to \mathcal{S}$, $(\triangleright_{\mathcal{A}}) : \mathbb{C}_2 \times \mathcal{A} \to \mathcal{A}$, and $(\triangleright_{\mathcal{O}}) : \mathbb{C}_2 \times \mathcal{O} \to \mathcal{O}$, such that for any tuple $(\boldsymbol{s}, \boldsymbol{a}, \boldsymbol{o}) \in \mathcal{S} \times \mathcal{A} \times \mathcal{O}$ the reflected tuple is given by $(g_r \triangleright_{\mathcal{S}} \boldsymbol{s}, g_r \triangleright_{\mathcal{A}} \boldsymbol{a}, g_r \triangleright_{\mathcal{O}} \boldsymbol{o}) \in \mathcal{S} \times \mathcal{A} \times \mathcal{O}$. Here, $\mathbb{C}_2 := \{e, g_r \mid g_r^2 = e\}$ denotes the reflection symmetry group, consisting of the identity e and reflection transformation g_r.

Intuitively, the reflective symmetry of the POMDP implies that under reflected observations, the respective optimal control action should also be reflected. Formally, the symmetric POMDP imposes a $\mathbb{C}_2$-invariance constraint on the learned stochastic policy [6]–[8]:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathbb{D}_H^{\text{exp}}} D_{\text{KL}} \left(\pi_{\text{exp}}(\cdot|\boldsymbol{x}) \;\|\; \pi_{\boldsymbol{\theta}}(\cdot|\boldsymbol{x})\right), \tag{2a}$$

$$\text{s.t. } \pi_{\boldsymbol{\theta}}(\boldsymbol{a}|\boldsymbol{x}) = \pi_{\boldsymbol{\theta}}(g_r \triangleright_{\mathcal{A}} \boldsymbol{a} | g_r \triangleright_{\mathcal{X}} \boldsymbol{x}), \;\; \forall\, \boldsymbol{a} \in \mathcal{A}, \boldsymbol{x} \in \mathcal{X}. \tag{2b}$$

Here, $g \triangleright_{\mathcal{X}} \boldsymbol{x} := [g \triangleright_{\mathcal{O}} \boldsymbol{o}_t]_{t=1}^{H}$ denotes the symmetry transformation of an observation-history sample.

III. Methodology

We address the constrained Behavior Cloning objective in Eq. (2) through *Flow Matching* [9], where sampling from the policy is modeled as integrating a continuous-time ODE:

$$\boldsymbol{a} = \boldsymbol{a}^{(0)} + \int_0^1 \boldsymbol{v}_*(\boldsymbol{a}^{(k)}, \boldsymbol{x}, k)\, \mathrm{d}k, \;\; \text{with } \boldsymbol{a}^{(0)} \sim \bar{\pi}_0(\cdot|\boldsymbol{x}). \tag{3}$$

Here, $\boldsymbol{v}_*$ is the velocity field that governs the evolution of the probability path from initial Gaussian noise $\bar{\pi}_0 = \mathcal{N}(0, I_d)$ to the target expert action distribution $\bar{\pi}_1 = \pi_{\text{exp}}$.

Previous work [10], [11] demonstrates that $\bar{\pi}_1$ is $\mathbb{G}$-invariant when both a $\mathbb{G}$-invariant prior (e.g., Gaussian noise) and a $\mathbb{G}$-equivariant velocity field are employed. Therefore, to address the symmetry-constrained BC objective in Eq. (2), we solve a FM problem with equivariance constraints:

$$\arg\min_{\boldsymbol{\theta}} \mathop{\mathbb{E}}_{\substack{\boldsymbol{a}^{(k)} \sim \bar{\pi}_k(\cdot|\boldsymbol{x}) \\ \boldsymbol{x} \sim \mathbb{D}_H^{\text{exp}}, k \sim U[0,1]}} \left\| \boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{a}^{(k)}, \boldsymbol{x}, k) - \boldsymbol{v}_*(\boldsymbol{a}^{(k)}, \boldsymbol{x}, k) \right\|_2^2, \tag{4a}$$

$$\begin{aligned} \text{s.t.} \quad & g_r \triangleright_{\mathcal{A}} \boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{a}, \boldsymbol{x}, k) = \boldsymbol{v}_{\boldsymbol{\theta}}(g_r \triangleright_{\mathcal{A}} \boldsymbol{a}, g_r \triangleright_{\mathcal{X}} \boldsymbol{x}, k), \\ & \forall\, \boldsymbol{a} \in \mathcal{A},\; \boldsymbol{x} \in \mathcal{X},\; k \in [0, 1]. \end{aligned} \tag{4b}$$

While the velocity field error in Eq. (4a) effectively minimizes the BC objective in Eq. (1) [12], enforcing equivariance of the velocity field $\boldsymbol{v}_{\boldsymbol{\theta}}$ in Eq. (4b) ensures that the policy $\pi_{\boldsymbol{\theta}}$ satisfies the invariance constraint in Eq. (2b) [11]. In this paper, we consider three ways to leverage the described symmetry prior for policy learning with FM:

(i) **Data Augmentation** (*SymAug*, Eq. (5)) adds reflected observation-action pairs to the dataset [5], [13]:

$$\mathbb{D}_H^{\text{aug}} := \{(g \triangleright_{\mathcal{X}} \boldsymbol{x}, g \triangleright_{\mathcal{A}} \boldsymbol{a}) \mid (\boldsymbol{x}, \boldsymbol{a}) \in \mathbb{D}_H^{\text{exp}}, g \in \mathbb{C}_2\}. \tag{5}$$

(ii) The **Equivariance Consistency Loss** (*EquivReg*, Eq. (6)) encourages reflection symmetry of the velocity field through a regularized training loss, as in [14]:

$$\arg\min_{\boldsymbol{\theta}} \mathcal{L}_{\text{CFM}}(\boldsymbol{\theta}) + \lambda \underbrace{\left\| \begin{matrix} g_r \triangleright_{\mathcal{A}} \boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{a}^{(k)}, \boldsymbol{x}, k) \\ - \boldsymbol{v}_{\boldsymbol{\theta}}(g_r \triangleright_{\mathcal{A}} \boldsymbol{a}^{(k)}, g_r \triangleright_{\mathcal{X}} \boldsymbol{x}, k) \end{matrix} \right\|^2}_{=:\mathcal{R}(\boldsymbol{\theta})}. \tag{6}$$

(iii) The **Explicit Equivariance Constraint** (*EquivNet*) parameterizes $\boldsymbol{v}_{\boldsymbol{\theta}}$ with a $\mathbb{C}_2$-equivariant NN, guaranteeing equivariance also for out-of-distribution inputs. We introduce a novel $\mathbb{G}$-equivariant transformer in which all modules are equivariant to an arbitrary compact group $\mathbb{G}$; details and code will be released with the final paper.

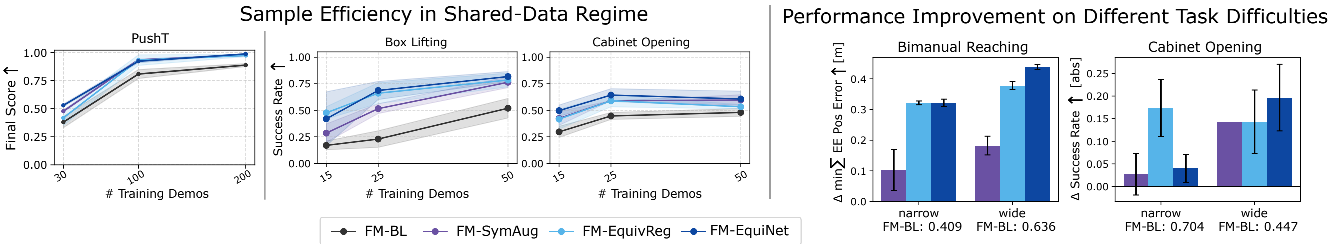

Fig. 3: Simulated mobile manipulation results in *shared-data* regime. *Left*: Sample efficiency across simulated tasks — symmetry-informed policies (*SymAug*, *EquivReg*, *EquivNet*) yield on average a 2× gain over the symmetry-agnostic baseline (*BL*). *Right*: Performance improvements of symmetry-informed policies over baseline increase with task difficulty.

Method	**Box Lifting Zero-shot e / g_r /total ↑**	**Cabinet Opening Zero-shot e / g_r /total ↑**
FM-Baseline	0.59 / 0.00 / 0.29	0.89 / 0.00 / 0.44
FM-SymAug	0.63 / 0.67 / 0.65	0.90 / 0.88 / 0.89
FM-EquivReg	0.70 / 0.72 / 0.71	0.85 / 0.87 / 0.86
FM-EquivNet	0.67 / 0.66 / 0.67	0.87 / 0.83 / 0.85

TABLE I: *Zero-shot* evaluation in simulated mobile manipulation: policies are trained on the original configuration (e) and tested on both e and its reflection (g_r). We report mean success rates over 50 rollouts and three training seeds.

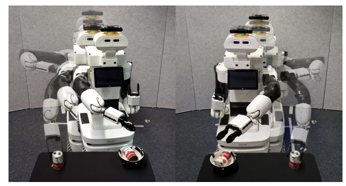

Mirrored (g_r) Original (e)

Method	e	g_r
Baseline	10/10	0/10
SymAug	10/10	10/10
EquivReg	10/10	10/10
EquivNet	10/10	10/10

Fig. 4: Real-world Tiago++ results: symmetry-informed policies generalize *zero-shot* to mirrored can-in-bowl task (g_r).

IV. EXPERIMENTS

Our experiments evaluate the reflective symmetry prior for IL in terms of sample efficiency, generalization, and optimality. We compare all symmetry-aware methods from Sec. III against a symmetry-agnostic FM baseline (*BL*) policy. We consider three settings: (i) the *Push-T* benchmark (Fig. 2); (ii) simulated MoMa tasks with the bimanual Tiago++ mobile manipulator (Fig. 1); and (iii) a real-world Tiago++ task. We use state-based observations of end-effector and object poses, and end-effector pose control. For the Tiago++ tasks, all quantities are defined in the robot base frame, and pose targets are tracked by a whole-body controller [15], [16].

Symmetric Zero-Shot Generalization Table I demonstrates successful zero-shot transfer of all symmetry-aware methods to mirrored MoMa configurations absent from the training data, matching performance in the original setting.

Sample Efficiency In a *shared-data* regime, where demonstrations from both the original (e) and reflected (g_r) tasks are available during training, all symmetry-informed policies consistently outperform the baseline across all dataset size regimes (Fig. 3-*left*). For the challenging bimanual *box lifting* task we observe additional optimality gains of enforcing the policy invariance optimality constraint Eq. (2b) (*EquivReg*, *EquivNet*) over data augmentation (*SymAug*).

Spatial Generalization We evaluate two difficulty levels (*narrow* and *wide*) for the *bimanual reaching* and *cabinet opening* tasks, corresponding to narrower and broader distributions of target and cabinet poses. Figure 3-*right* shows that symmetry-informed policies yield larger performance gains at higher task difficulty, indicating that morphological symmetry priors scale well to more complex tasks.

Real-World Results Finally, we validate the morphological symmetry prior on a real static-base unimanual Tiago++ task in *zero-shot* setting: the policy is trained on a pick-and-place task with the left arm, and evaluated both in this original (e) and the mirrored setting (g_r, right arm) (Fig. 4-*left*). Figure 4-*right* confirms repeatable zero-shot transfer under real-world conditions, validating the morphological symmetry prior beyond idealized simulation.

V. RELATED WORK

Previous work shows successful IL in MoMa [3], [17]–[21]. Symmetry priors in generative IL have been used to improve sample efficiency and generalization [22]–[30], mostly focusing on $\mathbb{SE}(3)$-equivariance. Morphological symmetry priors have been studied for locomotion [31]–[35] and ambidextrous manipulation [6], [14]. The work most closely related to ours is the concurrent study of [14], which exploits morphological symmetry for IL in bimanual manipulation via regularization. In contrast, our work addresses MoMa, focusing on generative policy learning with FM, with symmetry-priors enforced via equivariant NNs, loss regularization, and data augmentation.

VI. CONCLUSIONS

We introduced a FM framework for IL that, by construction, respects the bilateral morphological symmetry of mobile manipulators. Our results show improved sample efficiency, generalization, and optimality, and indicate that these benefits scale with increasing task complexity. Further, our analysis highlights the advantages of imposing the symmetry prior as an optimality constraint (Eq. (4b); *EquivReg*, *EquivNet*) over data augmentation. In future work, we will integrate visual observations and scale our method to temporal and rotational symmetries in the environment.

References

[1] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.

[2] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv preprint arXiv:2304.13705*, 2023.

[3] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, "Rt-1: Robotics transformer for real-world control at scale," *arXiv preprint arXiv:2212.06817*, 2022.

[4] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, "π_0: A vision-language-action flow model for general robot control," *arXiv preprint arXiv:2410.24164*, 2024.

[5] D. Ordoñez-Apraez, G. Turrisi, V. Kostic, M. Martin, A. Agudo, F. Moreno-Noguer, M. Pontil, C. Semini, and C. Mastalli, "Morphological symmetries in robotics," *The International Journal of Robotics Research*, vol. 44, no. 10-11, p. 02783649241282422, 2025. [Online]. Available: https://doi.org/10.1177/02783649241282422

[6] Z. Li, Y. Jin, D. O. Apraez, C. Semini, P. Liu, and G. Chalvatzaki, "Morphologically symmetric reinforcement learning for ambidextrous bimanual manipulation," *arXiv preprint arXiv:2505.05287*, 2025.

[7] M. Zinkevich and T. R. Balch, "Symmetry in markov decision processes and its implications for single agent and multiagent learning," in *Proceedings of the Eighteenth International Conference on Machine Learning*, ser. ICML '01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, p. 632.

[8] J. Brehmer, J. Bose, P. De Haan, and T. S. Cohen, "Edgi: Equivariant diffusion for planning with embodied agents," *Advances in Neural Information Processing Systems*, vol. 36, pp. 63 818–63 834, 2023.

[9] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," *arXiv preprint arXiv:2210.02747*, 2022.

[10] J. Köhler, L. Klein, and F. Noé, "Equivariant flows: exact likelihood generative learning for symmetric densities," in *International conference on machine learning*. PMLR, 2020, pp. 5361–5370.

[11] L. Klein, A. Krämer, and F. Noé, "Equivariant flow matching," *Advances in Neural Information Processing Systems*, vol. 36, pp. 59 886–59 910, 2023.

[12] M. Su, J. Y.-C. Hu, S. Pi, and H. Liu, "On flow matching kl divergence," *arXiv preprint arXiv:2511.05480*, 2025.

[13] I. Higgins, S. Racanière, and D. Rezende, "Symmetry-based representations for artificial and biological general intelligence," *Frontiers in Computational Neuroscience*, vol. 16, p. 836498, 2022.

[14] Z. Zhang, A. Mohan, S. Han, W. Shou, D. Wang, and Y. She, "Equibim: Learning symmetry-equivariant policy for bimanual manipulation," *arXiv preprint arXiv:2603.08541*, 2026.

[15] K. Zakka, "Mink: Python inverse kinematics based on MuJoCo," May 2025. [Online]. Available: https://github.com/kevinzakka/mink

[16] O. E. R. O. S. F. N. Andrea Del Prete, Nicolas Mansard, "Implementing torque control with high-ratio gear boxes and without joint-torque sensors," in *Int. Journal of Humanoid Robotics*, 2016, p. 1550044. [Online]. Available: https://hal.archives-ouvertes.fr/hal-01136936/document

[17] Z. Fu, T. Z. Zhao, and C. Finn, "Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation," in *Conference on Robot Learning (CoRL)*, 2024.

[18] J. Yang, I. Huang, B. Vu, M. Bajracharya, R. Antonova, and J. Bohg, "Mobi-π: Mobilizing your robot learning policy," *arXiv preprint arXiv:2505.23692*, 2025.

[19] N. M. M. Shafiullah, A. Rai, H. Etukuru, Y. Liu, I. Misra, S. Chintala, and L. Pinto, "On bringing robots home," *arXiv preprint arXiv:2311.16098*, 2023.

[20] P. Sundaresan, R. Malhotra, P. Miao, J. Yang, J. Wu, H. Hu, R. Antonova, F. Engelmann, D. Sadigh, and J. Bohg, "Homer: Learning in-the-wild mobile manipulation via hybrid imitation and whole-body control," *arXiv preprint arXiv:2506.01185*, 2025.

[21] A. Bahety, A. Balaji, B. Abbatematteo, and R. Martín-Martín, "Safemimic: Towards safe and autonomous human-to-robot imitation for mobile manipulation," *arXiv preprint arXiv:2506.15847*, 2025.

[22] D. Wang, S. Hart, D. Surovik, T. Kelestemur, H. Huang, H. Zhao, M. Yeatman, J. Wang, R. Walters, and R. Platt, "Equivariant diffusion policy," in *8th Annual Conference on Robot Learning*, 2024. [Online]. Available: https://openreview.net/forum?id=wD2kUVLT1g

[23] D. Wang, B. Hu, S. Song, R. Walters, and R. Platt, "A practical guide for incorporating symmetry in diffusion policy," *arXiv preprint arXiv:2505.13431*, 2025.

[24] C. Tie, Y. Chen, R. Wu, B. Dong, Z. Li, C. Gao, and H. Dong, "Et-seed: Efficient trajectory-level se(3) equivariant diffusion policy," 2024. [Online]. Available: https://arxiv.org/abs/2411.03990

[25] J. Yang, Z.-a. Cao, C. Deng, R. Antonova, S. Song, and J. Bohg, "Equibot: Sim(3)-equivariant diffusion policy for generalizable and data efficient learning," in *8th Annual Conference on Robot Learning*, 2024.

[26] J. Yang, C. Deng, J. Wu, R. Antonova, L. Guibas, and J. Bohg, "Equivact: Sim(3)-equivariant visuomotor policies beyond rigid object manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 9249–9255.

[27] J. Jankowski, P. Klink, I. Posner, E. Gundogdu, K. Park, and C. Erdogan, "Guaranteed se(3)-equivariant control via hand-centric behavior cloning," 2025. [Online]. Available: https://corl25-genpriors.github.io/Papers/15_Guaranteed_SE_3_Equivariant.pdf

[28] M. Jia, D. Wang, G. Su, D. Klee, X. Zhu, R. Walters, and R. Platt, "Seil: Simulation-augmented equivariant imitation learning," *arXiv preprint arXiv:2211.00194*, 2022.

[29] N. Funk, J. Urain, J. Carvalho, V. Prasad, G. Chalvatzaki, and J. Peters, "Actionflow: Equivariant, accurate, and efficient policies with spatially symmetric flow matching," *arXiv preprint arXiv:2409.04576*, 2024.

[30] J. Chang, R. Mei, W. Ke, and X. Xu, "Efficientflow: Efficient equivariant flow policy learning for embodied ai," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 40, no. 24, 2026, pp. 19 496–19 504.

[31] M. Mittal, N. Rudin, V. Klemm, A. Allshire, and M. Hutter, "Symmetry considerations for learning task symmetric robot policies," *arXiv preprint arXiv:2403.04359*, 2024.

[32] F. Abdolhosseini, H. Y. Ling, Z. Xie, X. B. Peng, and M. van de Panne, "On learning symmetric locomotion," in *Motion, Interaction and Games*, ser. MIG '19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: https://doi.org/10.1145/3359566.3360070

[33] 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 *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 6899–6906.

[34] S. Wei, X. Chen, F. Xie, G. E. Katz, Z. Gan, and L. Gan, "Ms-ppo: Morphological-symmetry-equivariant policy for legged robot locomotion," *arXiv preprint arXiv:2512.00727*, 2025.

[35] F. Xie, S. Wei, Y. Song, Y. Yue, and L. Gan, "Morphological-symmetry-equivariant heterogeneous graph neural network for robotic dynamics learning," *arXiv preprint arXiv:2412.01297*, 2024.