



**Università
di Genova**

PHD PROGRAM IN BIOENGINEERING AND ROBOTICS

**Morphological Symmetries in Robot Learning
A Hitchhiker's Guide to Symmetry-aware Robot
Modelling, Control, and Estimation**

by

Daniel Felipe Ordoñez Apraez

Thesis submitted for the degree of *Doctor of Philosophy* (38° cycle)

Wednesday 13th May, 2026

Prof. Dr. Massimiliano Pontil
Dr. Claudio Semini
Prof. Dr. Carlos Mastalli

Supervisor
Supervisor
Supervisor

Prof. Dr. Vladimir S. Kostić
Dr. Giulio Turrisi

Postdoctoral supervisor
Postdoctoral supervisor

Prof. Dr. Paolo Massobrio

Head of the PhD program

Thesis Jury:

Prof. Dr. Noémie Jaquier, *KTH - Royal Institute of Technology*
Prof. Dr. Justin Carpentier, *Inria & École normale supérieure*
Prof. Dr. Carlo Ciliberto, *University College London*

External examiner
External examiner
External examiner

Dibris

Department of Informatics, Bioengineering, Robotics and Systems Engineering

A mi Mamá, a mi Hermano y a mi Papá.
Por su apoyo incondicional e influencia que me han permitido perseguir este sueño.
También a toda la familia Apraez y Ordoñez.
Que sea yo el primero, pero no el último, académico de estos bellos lineajes Colombianos.
Y por ultimo a mi Alma Mater, la Universidad Nacional de Colombia.
Por ser el ideal de Colombia que aspiro mantener y ayudar a construir.

"Mathematics is the art of giving the same name to different things"

Henri Poincaré

"Machine learning is the art of giving different names to the same thing"

Yi Ma

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Daniel Felipe Ordoñez Apraez
May 2026

Acknowledgements

My entire path towards this PhD and this ongoing pursuit of an academic and research career has been blessed with incredible support from family, loved ones, friends, mentors, and colleagues, which makes me feel incredibly lucky and grateful.

To my family: sin ustedes, no estaría finalizando un doctorado, ni quizás me habría atrevido a intentarlo. A mi Papá, le agradezco por inculcar y presionar, a mi hermano y a mí, a alcanzar la excelencia académica, profesional y personal. De ti aprendí la ambición que me empuja día a día a seguir aprendiendo y apuntando siempre más alto. A mi Mamá, mi soporte y fuente inagotable de sabiduría emocional, le agradezco por su constante ejemplo que es una guía, aún hoy en día, de cómo puedo mejorar como persona. Por inculcarme el deseo de dirigir mi vida profesional hacia una contribución significativa a la sociedad. A mi hermano, por ser el constante ejemplo de excelencia que me empujó a este camino. No me hubiera esforzado tanto en el colegio, ni desarrollado el amor al conocimiento, si no tuviese que competir siempre con el fantasma de su nerdismo. No me hubiera enamorado de la robótica, si no lo hubiese visto a usted primero enamorarse de ella. No me hubiera atrevido a intentar ser profesor e investigador, si usted no hubiera perseguido ese camino primero. No habría Daniel sin Nicolás.

To the loved one: A causa de vivir a más de 9000 km de distancia de casa, el apoyo más presente, constante e importante que recibí durante el doctorado fue el tuyo, Katya. Es difícil imaginar que hubiera cumplido esta meta sin tu apoyo y amor. Por acompañarme en las traspasadas y el estrés antes de los deadlines. Por ser mi verificadora de la realidad cuando la tristeza y el síndrome del impostor me inmovilizaban. Por no dejarme rendirme las muchas veces que quise tirar la toalla. Por tu inmensa curiosidad y paciencia cuando te explicaba mis tontas ideas. Por ayudarme a creer en mí, y por ser mi compañera de vida durante esta etapa. Te estaré siempre eternamente agradecido.

I chose this PhD program because it offered me the opportunity to be part of a group of applied mathematicians doing machine learning and a bunch of roboticists doing legged locomotion.

To the roboticists: Dr. Semini, Dr. Mastalli, and Dr. Turrise, I am deeply grateful for their mentorship and for encouraging me to explore obscure corners of robotics. Claudio has been an invaluable guide but also the greatest support during the many bureaucratic nightmares I put myself through during this period. It is not easy to run such a big research team, but it is even harder to ensure that every member of the team feels taken care of and supported. For this, my hat goes off to you, Claudio. To Dr. Mastalli, I thank you for the many recurrent weekly meetings that ended in prolonged hours of brainstorming and discussions, which shaped this thesis and my research in general. To Dr. Turrise, I thank you for providing your constant supervision, for believing in my ideas, and for letting me explore freely.

To the mathematicians: Dr. Pontil, Dr. Kostić, and Dr. Lounici. I cannot thank you enough for adopting me into your group, despite my lack of mathematical background, and making me feel like I

belong. Nobody would have guessed at the beginning of this PhD that I, as a roboticist, would end up spending most of the time with the math guys. To Dr. Pontil, I am deeply grateful for your support during the highs and lows of this period, for telling me the first day that the most important thing of the PhD was to do something that I love, for the flexibility to explore and the gentle pressures to reach milestones. To Vladi: I said it before and I will say it again, your supervision and your guidance have been the most influential professional relationship I have had thus far. Thank you for the patience you had with my lack of formal background, for shaping how I approach a new problem and how I do research, and for the many times in which you behaved more like a very wise friend than a supervisor. To Dr. Lounici, although our collaboration was brief, I am grateful for the many times I found you at 3AM still doing corrections to the manuscript and for the kindness of your mentorship.

To the friends and colleagues I gained during this period: Andrea Patrizi, Soha Satvati, Mattia Risiglione, Alek Fröhlich, Matia Bjović, Erfan Mirzaei, Riccardo Grazi, and Isak Falk. Thank you for making my time in Italy more beautiful.

To the students with whom I had the pleasure of collaborating with: Zhi Su, Zechu (Steven) Li, and Maximilien Siebenborn. I thank you for believing in my ideas and for the time you invest in making symmetry exploitation in robotics more real.

Daniel Felipe Ordoñez Apraez
Wednesday 13th May, 2026

*Wise men in their bad hours have envied
The little people making merry like grasshoppers
In spots of sunlight, hardly thinking
Backward but never forward, and if they somehow
Take hold upon the future they do it
Half asleep, with the tools of generation
Foolishly reduplicating
Folly in thirty-year periods; they eat and laugh too,
Groan against labors, wars and partings,
Dance, talk, dress and undress; wise men have pretended
The summer insects enviable;
One must indulge the wise in moments of mockery.
Strength and desire possess the future,
The breed of the grasshopper shrills,
“What does the future matter, we shall be dead?”
Ah, grasshoppers,
Death’s a fierce meadowlark: but to die having made
Something more equal to the centuries
Than muscle and bone, is mostly to shed weakness.
The mountains are dead stone, the people
Admire or hate their stature, their insolent quietness,
The mountains are not softened nor troubled
And a few dead men’s thoughts have the same temper.*

— Robinson Jeffers

Abstract

This thesis presents a comprehensive theoretical study of symmetries in robot dynamics, characterizing them as foundational, physics-informed geometric priors that can significantly enhance the efficiency, generalization, and optimality of learning methods for dynamics modeling, optimal control, and state estimation in robotics. Our main focus is on morphological symmetries—a previously underexplored type of symmetry arising from regularities in a robot’s morphology, associated with the replication of kinematic structures and the balanced distribution of mass. In analytical methods, we characterize how these symmetries extend to the robot’s equations of motion, generalized mass matrix, Jacobians, configuration space, and the observation space of both proprioceptive and exteroceptive sensor measurements. For data-driven methods, we characterize the conditions under which these symmetries transform the problems of dynamics modeling, optimal control, and state estimation into symmetry-constrained learning problems, where optimal models must satisfy strict invariance and equivariance properties. Thus providing a clear theoretical justification for exploiting symmetries in these applications. To substantiate our claims, we present extensive empirical evidence in locomotion and bimanual manipulation control, deterministic and probabilistic state estimation, and dynamics modeling via transfer/Koopman operators, demonstrating that leveraging symmetry in robot learning leads to significant improvements in sample efficiency, generalization, and robustness. Lastly, to facilitate the practical use of the theory and applications outlined in this work, we introduce two open-access repositories, [morpho_symm](#) and [symm_learning](#), which provide a comprehensive collection of tools and resources for symmetry exploitation in robot learning.

Table of contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Which symmetries arise in robot dynamics? | 2 |
| 1.2 | Motivation: why exploiting symmetries matter in robot learning? | 4 |
| 1.3 | Related work | 5 |
| 1.4 | Thesis structure | 7 |
| 1.5 | Thesis deliverables | 7 |
| 1.5.1 | Open-access software | 8 |
| 1.6 | Notation | 10 |
| 2 | Preliminaries | 11 |
| 2.1 | Analytical models of robot dynamics | 12 |
| 2.2 | Data-driven stochastic dynamics models | 17 |
| 2.3 | Group and representation theory for mortals | 20 |
| 2.3.1 | Decomposition of symmetric vector spaces | 28 |
| 2.3.2 | Maps between symmetric vector spaces | 29 |
| 2.4 | \mathbb{G} -invariant random variables | 31 |
| 3 | Symmetries of robot dynamics | 37 |
| 3.1 | Symmetries as energy-preserving transformations | 37 |
| 3.2 | Taxonomy of symmetries in robotics | 42 |
| 3.2.1 | Temporal symmetries | 42 |
| 3.2.2 | Spatial symmetries | 43 |
| 4 | Morphological Symmetries (MSs) | 46 |
| 4.1 | Modularity and symmetry in the robot's morphology | 50 |
| 4.2 | Morphology constraints induced by symmetries | 54 |

| | | |
|-----------|---|------------|
| 4.3 | Dynamics' harmonic analysis | 57 |
| 5 | Leveraging symmetries in robot learning | 62 |
| 5.1 | Approximating \mathbb{G} -invariant conditional probability distributions | 63 |
| 5.2 | When to leverage symmetry priors in robot learning? | 65 |
| 5.3 | How to leverage symmetry priors in robot learning? | 67 |
| 6 | Morphological symmetries in dynamics modelling | 70 |
| 6.1 | Symmetric stochastic dynamics | 70 |
| 6.2 | Modelling symmetric dynamics under partially observability | 74 |
| 7 | Morphological symmetries in state estimation | 79 |
| 7.1 | Symmetry-priors in optimal state estimation | 79 |
| 7.2 | Deterministic inference | 83 |
| 7.2.1 | Supervised regression | 83 |
| 7.2.2 | Supervised classification | 86 |
| 7.3 | Probabilistic inference | 88 |
| 7.3.1 | Conditional probability modelling and uncertainty quantification | 89 |
| 8 | Morphological symmetries in data-driven optimal control | 96 |
| 8.1 | Symmetry priors in optimal control | 96 |
| 8.2 | Symmetry priors in optimal control under partial observability | 103 |
| 8.3 | Leveraging symmetries in data-driven optimal control | 105 |
| 8.3.1 | Imitation learning | 105 |
| 8.3.2 | Model-free reinforcement learning | 114 |
| 9 | \mathbb{G}-invariant conditional probability estimation | 123 |
| 9.1 | Operator-theoretic modelling of conditional distributions | 123 |
| 9.1.1 | Symmetry priors of the conditional expectation operator | 125 |
| 9.2 | Contrastive spectral representation learning | 127 |
| 9.2.1 | Equivariant contrastive spectral representation learning | 129 |
| 9.3 | Koopman/Transfer dynamics operators | 135 |
| 9.3.1 | Closed-loop transfer/Koopman operators | 138 |
| 9.3.2 | Modelling robot dynamics via closed-loop Koopman/Transfer operators . . . | 143 |
| 10 | Discussion and Future Work | 149 |

| | |
|------------------------|------------|
| Table of contents | ix |
| References | 151 |
| I Glossary | 161 |
| List of figures | 164 |
| List of tables | 167 |

Chapter 1

Introduction

In the last decade, robotics has witnessed an experimental revolution driven by the adoption of deep learning methods to solve long-standing challenges in robot modelling, control, and estimation. Indeed, many applications once thought to be decades away from resolution—such as legged locomotion, dexterous manipulation, and autonomous driving—have proven attainable given vast amounts of data and computational resources (Grigorescu et al., 2020; Kumar et al., 2016; Rudin et al., 2022).

However, expanding robotics into applications such as elderly care, household assistance, and last-mile delivery demands stringent safety and generalization, neither of which can be achieved merely by scaling compute and data (Roy et al., 2021; Sünderhauf et al., 2018). The high costs and risks of data collection, limited computational budgets, and the gap between simulated and real-world data remain key challenges, underscoring the need for learning approaches with greater sample efficiency, robustness, and generalizability (Kaelbling, 2020; Roy et al., 2021; Sünderhauf et al., 2018).

Geometric Deep Learning (GDL) (Bronstein et al., 2021) has emerged as a unifying theoretical paradigm showing that deep learning’s success stems from the use of Neural Network (NN) architectures designed to leverage the *symmetries* and structure of the data they process (Cohen et al., 2019; Kondor and Trivedi, 2018; Weiler et al., 2021). From the GDL’s perspective, the robustness and sample efficiency of convolutional neural networks in computer vision, transformers in natural language processing, and graph neural networks in particle physics arises from the architectures’ ability to exploit the grid structure of images, the sequential nature of language, and the graph structure of particle interactions.

This thesis provides a comprehensive characterization of the GDL paradigm in robotics by identifying foundational, “bitter-pilled” symmetry priors that are present in virtually every robot control task and operational environment. These priors can therefore be leveraged across the core learning problems in robotics: dynamics modelling, control, and estimation. The primary goal of this thesis is to present a compelling theoretical and empirical argument for symmetry exploitation in robot learning, while also providing practical tools and methods that enable practitioners to identify and incorporate these symmetries in their own work.

To this end, the thesis theoretically characterizes how symmetries of Newtonian/classical mechanics can be used to improve the data efficiency, generalization, and optimality of learning algorithms in robotics, and provides substantial empirical evidence supporting these claims. Crucially, the analysis covers the rotational and translational symmetries commonly exploited in mechanical-system




Figure 1.1 Illustration of the spatial symmetries of the Atlas humanoid robot. **Left:** Depiction of the robot’s bilateral reflection morphological symmetry. This symmetry enables a feasible reconfiguration of the robot’s configuration state $(\mathbf{q}, \dot{\mathbf{q}})$ to a reflected state $(g_s \triangleright_Q \mathbf{q}, g_s \triangleright_{T_q Q} \dot{\mathbf{q}})$ that exhibits equivalent dynamics up to reflection. **Right:** Invariance of classical mechanics under environmental symmetry transformations, corresponding to rigid transformations of the chosen inertial reference frame leaving the underlying dynamics invariant.

modelling and control (Bloch et al., 1996; Wieber, 2006), while extending beyond them to a broader class of symmetries arising from regularities in robot morphology. We refer to these as *morphological symmetries* and constitute the main focus of this thesis.

To facilitate the practical adoption of symmetry exploitation in robotics, this work is accompanied by two software libraries, `morpho_symm` and `symm_learning`, which provide tools for identifying and leveraging symmetry priors in robotic systems.

1.1 Which symmetries arise in robot dynamics?

To no surprise, the symmetries of robot dynamics are exactly those of Newtonian (classical) mechanics, which are responsible for the invariance of Newton’s equations of motion to changes in the chosen inertial reference frame and the direction of time. Specifically, these symmetries—characterized algebraically by the seminal works of Emmy Noether and Henri Poincaré (Noether, 1983; Poincaré, 1905)—consist of *spatial symmetries*: rigid translations, rotations, and reflections in 3-dimensional space; and *temporal symmetries*: time translations and time reflection/reversal.




Figure 1.2 Taxonomy of symmetries in robot dynamics.

When studying the dynamics of a specific robotic system, the interplay between a robot’s morphology and classical mechanics enables the manifestation of spatial symmetries in two distinct ways: *environmental/base symmetries* and *morphological symmetries*. Environmental/base symmetries—often referred to as base symmetries in the context of floating-base robots (Bloch et al., 1996), and more generally interpretable as reference-frame symmetries—represent the standard interpretation of spatial

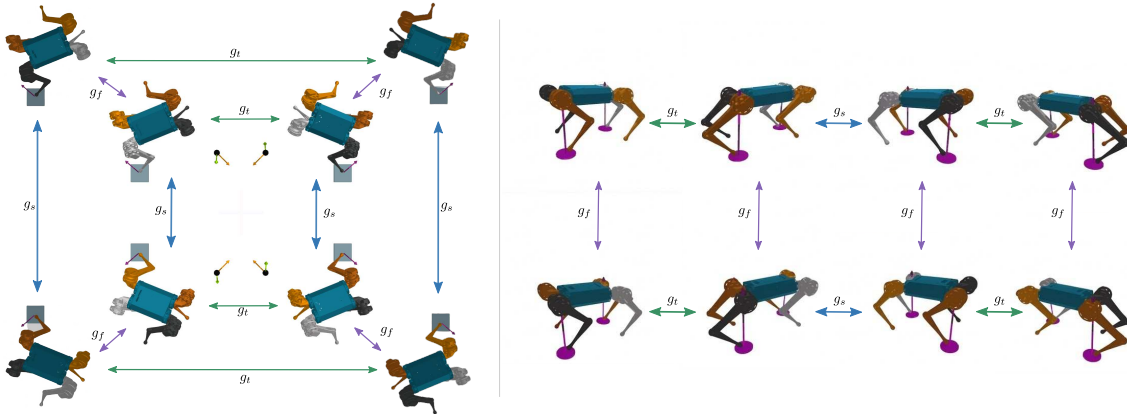


Figure 1.3 Morphological symmetry group of the Mini-Cheetah quadruped robot (Katz et al., 2019), defined by the cuboid symmetry group $\mathbb{G} = \{e, g_s \mid g_s^2 = e\} \times \{e, g_t \mid g_t^2 = e\} \times \{e, g_f \mid g_f^2 = e\}$ of order 8, generated by reflections with respect to the robot’s sagittal (g_s), transversal (g_t), and frontal (g_f) planes. In practice, these 8 symmetries imply that for any given robot state, there exist 7 distinct symmetric states with equivalent dynamics. Consequently, the robot can locomote with its torso oriented forward, backward, left, right, or even upside down (see animation). The output of any dynamics model, control policy, or state-estimation model at a particular state can be transferred to the other 7 symmetric states by applying the corresponding symmetry transformation. For this robot, any locomotion control policy is optimal only if it performs equivalently across all 8 symmetric states.

symmetries as rigid transformations of the environment, or equivalently of the chosen inertial reference frame, that leave the underlying dynamics invariant (see fig. 1.1-right). Morphological symmetries, on the other hand, are reconfigurations of the robot’s body that allow the robot to mimic some of these spatial rigid transformations and consequently also represent transformations that leave the underlying dynamics invariant (see fig. 1.1-left). These symmetries are enabled by intrinsic regularities of the robot’s morphology, associated with replications of kinematic chains and balanced distributions of mass.

The simplest and most common example of a morphological symmetry is the bilateral/sagittal symmetry in humans, and the majority of robotic systems, where each side of the body is a mirror image of the other. This symmetry allows individuals to easily mimic a mirrored body pose by simply reorienting the torso and permuting the roles of their arms and legs (see fig. 1.1-left). While this bilateral symmetry is ubiquitous in both biological and robotic systems (Gupta et al., 2021; Holló, 2017), robots can exhibit a broader number of such symmetries. For instance, certain quadruped robots have the ability to mimic up to eight distinct spatial transformations, as illustrated in this animation and fig. 1.3-left.

The literature of symmetry exploitation in robotics has largely focused on environmental symmetries, with the primary goal being constructing dynamics models (Brehmer et al., 2024; Kim et al., 2023), control policies (Brehmer et al., 2024; Wang et al., 2023a, 2024), perception modules (Lin et al., 2024; Zhang et al., 2024), and state-estimation pipelines (Barrau and Bonnabel, 2018; Hartley et al., 2020) that are invariant to changes in the chosen inertial reference frame (Ghaffari et al., 2022), that is, invariant to environmental symmetry transformations (see fig. 1.1-right). In contrast, there is a clear gap in the robotics literature covering the formal characterization and exploitation of morphological symmetries. As a consequence, this thesis aims to tackle two goals: (i) to satisfy the requirements of a PhD, by writing a thesis describing our original research; and (ii) to fill the literature gap by proving a comprehensive and self-contained characterization of spatial symmetries in robotics, aiming to be both theoretically rigorous and accessible and intuitive for practitioners.

1.2 Motivation: why exploiting symmetries matter in robot learning?

The primary motivation for studying and leveraging these symmetries is that they provide a powerful geometric prior for modelling and controlling the temporal evolution of a robot. Specifically, for any given robot state, the optimal predictions of a dynamics model, control policy, or state-estimation model can be directly transferred to all symmetry-related states; see [figs. 1.1](#) and [1.3](#).

To build intuition, consider the following scenario: you are walking toward a target and encounter an obstacle with your right leg, causing you to stumble. The optimal visuomotor control response—needed to react, regain balance, and replan your trajectory—is essentially a reflected version of the response you would execute if the same event occurred with your left leg. This is because the dynamics of both mirrored body poses at the moment of contact are equivalent (up to a reflection) ([fig. 1.1-left](#)). As a result, the processes of state estimation, planning, and control that your brain undertakes in one event are related by a symmetry transformation to those in the mirrored event. This same reasoning applies to any humanoid, bimanual manipulator, or quadruped robot with bilateral symmetry, as well as to robots with a greater number of morphological symmetries (see [fig. 1.3](#)).

To ground this intuition theoretically, in [chapters 6 to 8](#) we examine the implications of symmetry priors for the three core learning problems in robotics: dynamics modelling, optimal control, and state estimation; which we formulate in their most general form as approximations of conditional probability distributions. We find that for a robotic system with an environmental or morphological symmetry group \mathbb{G} —such as the reflection group of Atlas in [fig. 1.1](#) or the cuboid symmetry group of Mini Cheetah in [fig. 1.3](#)—the following properties hold:

- ✱ **Optimal modelling:** Robot dynamics modelling should be treated as a symmetry-constrained learning problem, since the set of optimal stochastic dynamics models is restricted to a subset of \mathbb{G} -invariant models (see [prop. 6.1](#)).
- ✱ **Optimal estimation:** Probabilistic state estimation should likewise be treated as a symmetry-constrained learning problem, as the set of optimal estimation models is restricted to a subset of \mathbb{G} -invariant state-estimation models (see [prop. 7.1](#)). In the case of deterministic inference, optimal models are constrained to be \mathbb{G} -equivariant.
- ✱ **Optimal control:** Optimal control should be formulated as a symmetry-constrained learning problem, since any optimal stochastic control policy must be \mathbb{G} -invariant ([corollary 8.2](#)), and any optimal deterministic control policy must be \mathbb{G} -equivariant or a combination of \mathbb{G} -equivariant policies (see [corollary 8.3](#)).

We argue that these theoretical results are satisfied by virtually any control task and operational environment for a robot with a morphological or environmental symmetry group, and provide substantial empirical evidence of the benefits of symmetry exploitation in the applications of robot locomotion and bimanual manipulation.

A practitioner’s intuition To further motivate the reader we provide grounded intuitions for the two core applications in robotics: locomotion and manipulation.

- ✱ **Symmetry aware locomotion:** If you aim to learn a locomotion policy that is capable of overcoming inertial and moving disturbances—e.g., carrying a grocery bag with one arm, hitting an obstacle with your right shoulder, a malfunctioning actuator in the left leg—and you

believe the symmetry transformed disturbance is equally likely to occur in practice (fig. 3.3)—e.g., carrying the bag with the other arm, hitting the obstacle with the left shoulder, an actuator malfunction on the right leg—then you should leverage spatial symmetries.

- * **Symmetry aware manipulation:** The dynamics of in-hand manipulation with a left hand are equivalent to the dynamics of in-hand manipulation with the right hand, up to reflection for any object. If you desire your bimanual manipulation policy to be ambidextrous, you should leverage spatial symmetries (see sections 8.3.1 and 8.3.2).

1.3 Related work

Structural symmetries in physics

The morphological symmetries studied here closely parallel structural (point-group) symmetries in particle physics (Dresselhaus et al., 2007, 8.3), which arise from the replication of identical particles or atoms in molecules and crystals. The principles and methods for exploiting these symmetries in atomistic and molecular dynamics (Cornwell, 1997; Jumper et al., 2021; Klein et al., 2023; Noé et al., 2020) are transferable to robotics. However, their implications in robotics diverge from those in physics for which symmetries primarily inform modeling and estimation, while in robotics, they serve to inform both modelling and control. Furthermore, a robot morphology is typically stable and engineered, making these symmetries a persistent inductive bias—unlike in physics, where symmetries can change with particle gain or loss.

Symmetry exploitation in robotics

Environmental symmetries Most modern robotics literature leveraging symmetries in robotics learning focuses on the full group of Euclidean isometries of Newtonian/classical mechanics, \mathbb{E}_3 , or on one of its subgroups. These works aim to learn control policies (Brehmer et al., 2024; Kim et al., 2023; Wang et al., 2023a, 2024; Zhu et al., 2023), perception modules (Lin et al., 2024; Zhang et al., 2024), and state-estimation pipelines (Barrau and Bonnabel, 2018; Hartley et al., 2020) that are equivariant or invariant to rotations, translations, and reflections in 2D or 3D operational spaces (Rezaei-Shoshtari et al., 2022; Wang et al., 2023a). The main goal of these approaches is to respect the manifold structure of \mathbb{E}_3 and to design data-driven models that are invariant to the choice of inertial reference frame (Ghaffari et al., 2022) (fig. 3.2-a).

In the context of floating-base robotic systems—e.g., humanoids, quadrupeds, aerial vehicles, and space vehicles—environmental symmetries are often referred to as *base symmetries* (Bloch et al., 1996; Mishra, 2025; Murray et al., 2017). They have been extensively studied for the design of analytic dynamics models, state estimation, and optimal control policies that are \mathbb{E}_3 -equivariant (Ghaffari et al., 2022; Mishra, 2025), in accordance with the Euclidean isometries of classical mechanics. These works present an analysis, similar in spirit to that of chapter 3, of the structural properties of the Lagrangian formulation of robot dynamics in the presence of symmetries. However, they rely heavily on a differential-geometric formulation of mechanics and focus exclusively on the implications of environmental/base symmetries. In contrast, this thesis presents a group-theoretic characterization of the implications of symmetries in Lagrangian mechanics, covering both environmental/base and morphological symmetries. The latter being the core focus of this work.

Morphological symmetries Prior to the formalization of morphological symmetries in (Ordoñez-Apraez et al., 2022; Ordoñez-Apraez et al., 2025), most research focused exclusively on humanoid bilateral symmetry. Early work by Yeh et al. (2019) exploited human reflection symmetry to improve pose estimation performance and sample efficiency. The SMPL and SMPL-X body models subsequently incorporated reflection symmetry into their parametrization, enabling its use in pose estimation, body modeling, and character control (Loper et al., 2015; Pavlakos et al., 2019). In control, (Finzi et al., 2021; van der Pol et al., 2020) explored equivariant NNs for symmetry-aware control, while (Abdolhosseini et al., 2019; Finzi et al., 2021; Ordoñez-Apraez et al., 2022) provided the first empirical qualitative and quantitative evidence of the benefits of bilateral symmetry exploitation in humanoid locomotion learning.

Recent work has expanded the study of morphological symmetries to more complex robot morphologies and larger symmetry groups. In locomotion, (Ding and Gan, 2024) uses harmonic analysis to reveal diverse quadrupedal gait patterns, echoing the approach of (Ordoñez-Apraez et al., 2024). (Su et al., 2024) leverages reflection symmetry in quadruped robots for model-free reinforcement learning, leading to Bao et al. (2025); Cathomen et al. (2025), who exploit the Klein-four/rectangle symmetry group, \mathbb{K}_4 (see fig. 2.1), of the Anymal quadruped to learn approximately \mathbb{K}_4 -equivariant policies, significantly improving sample efficiency and optimality. In estimation, Martinez et al. (2026) utilize equivariance constraints on the robot’s generalized mass matrix to enhance receding-horizon estimation of inertial parameters and disturbances, showing that the morphological symmetry prior is crucial for optimal system identification. (Xie et al., 2024) introduces a symmetry-aware graph NN that exploits both morphological symmetry and kinematic structure, yielding state estimation performance that consistently surpasses standard baselines, highlighting the relevance of topological deep learning (Hajij et al., 2022) in robotics. In manipulation, Brehmer et al. (2024); Li et al. (2025); Zhang et al. (2026) leverage the reflection symmetry of bimanual manipulators in model-free reinforcement and imitation learning, resulting in approximately and exactly ambidextrous manipulation control policies that demonstrate consistent gains in sample efficiency, and control optimality.

Symmetries as local regularization priors A considerable body of recent work uses symmetry groups as regularization priors to improve the performance of learned control policies for robotic manipulation. For instance, Wang et al. (2022a) promotes two-dimensional rotational equivariance of a manipulator policy on a workspace plane, i.e., \mathbb{SO}_2 -equivariance; Wang et al. (2024, 2026) promote rotation and translation equivariance, i.e., \mathbb{SE}_3 -equivariance; and Yang et al. (2024) promotes three-dimensional rotation, translation, and scale equivariance, i.e., \mathbb{SIM}_3 -equivariance. These regularization strategies have been shown to be effective in specific manipulation tasks and constrained operational settings. However, the way these works leverage local symmetries as regularization priors differs substantially from the concept of symmetry exploitation considered in this thesis.

In this thesis, symmetries of robot dynamics are understood as transformations of the robot state space that relate states with equivalent dynamics (see chapter 3). Crucially, these transformations are assumed to hold over the entire state and action spaces, and therefore induce global invariances of the robot Lagrangian, following the interpretation of symmetries in (Bloch et al., 1996; Murray et al., 2017; Wieber, 2006). From this perspective, the works of Wang et al. (2022a, 2024, 2026); Yang et al. (2024), and related approaches, can be interpreted as leveraging task- or workspace-level symmetry priors that are valid on a restricted subset of the robot state-action space. Such priors can be beneficial within that subset, but they need not correspond to global symmetries of the underlying

robot dynamics, since reachability, morphology, joint limits, and other configuration-space constraints may break the assumed equivariances.

Following the taxonomy of incorrect, correct, and extrinsic symmetry priors proposed by Wang et al. (2023b), this thesis focuses on *correct* symmetry priors: symmetries that hold globally across the state-action space and arise from the structure of Newtonian mechanics and the robot morphology.

1.4 Thesis structure

This thesis comprises eight chapters (excluding introduction and conclusion), organized into background material and the study of symmetries in both analytical models of robot dynamics and the core robot learning applications:

- ✱ **Background:** Chapter 2 provides background on deterministic and stochastic models of robot dynamics, core learning problems in robotics (modelling, control, estimation), and foundational concepts in group/representation theory and probability theory in the presence of symmetries.
- ✱ **Symmetries in analytical models of robot dynamics:** Chapter 3 formally introduces spatial symmetries in Lagrangian mechanics and presents a taxonomy of symmetries in robotics. Chapter 4 is entirely dedicated to morphological symmetries and their implications, with a special focus on floating-base robotic systems.
- ✱ **Leveraging symmetries in robot learning:** Chapter 5 studies the learning problem of approximating conditional distributions under the symmetry priors arising from spatial symmetries. The results of this chapter are used to motivate the use of symmetry exploitation in robot modeling, control, and estimation in chapters 6 to 8.
- ✱ **Symmetries in robot learning:** In chapters 6 to 8 we study the problem of dynamics modelling, state estimation, and optimal control in the presence of spatial symmetry priors. Finally, chapter 9 develops a framework for approximating \mathbb{G} -invariant conditional probabilities via conditional expectation operators, applied to uncertainty quantification in section 7.3.1 and Koopman/transfer operator dynamics modelling in section 9.3.

1.5 Thesis deliverables

Below is a list of publications included in this thesis composed of seven conference papers and one journal article. For each publication where I am a co-author, I provide a detailed description of my contributions. Unless explicitly stated in editorial notes, the contents of the thesis describe original work by myself and supervisors.

- ✱ Max Siebenborn, Daniel Ordoñez-Apraez, Sophie Lueth, Giulio Turrisi, Massimiliano Pontil, Claudio Semini, and Georgia Chalvatzaki. Morphologically equivariant flow matching for bimanual mobile manipulation. *arXiv preprint arXiv:2605.12228*, 2026. doi: 10.48550/arXiv.2605.12228. URL <https://arxiv.org/abs/2605.12228>
First and second author have equal contribution.

- ✱ [Daniel Ordoñez-Apraez](#), Vladimir Kostić, Alek Fröhlich, Vivien Brandt, Karim Lounici, and Massimiliano Pontil. Representation learning for equivariant inference with guarantees. In *The International Conference on Learning Representations*, 2026
- ✱ [Daniel Ordoñez-Apraez](#), Giulio Turrisi, Vladimir Kostic, Mario Martin, Antonio Agudo, Francesc Moreno-Noguer, Massimiliano Pontil, Claudio Semini, and Carlos Mastalli. Morphological symmetries in robotics. *The International Journal of Robotics Research (IJRR)*, 2025
- ✱ Zechu Li, Yufeng Jin, [Daniel Ordoñez-Apraez](#), Claudio Semini, Puze Liu, and Georgia Chaltatzaki. Morphologically symmetric reinforcement learning for ambidextrous bimanual manipulation. In *The Conference on Robot Learning (CoRL)*. IEEE, 2025
Contribution: Substantial contributions to writing; theoretical analysis of reflection symmetries in model-free reinforcement learning for multi-task and multi-agent scenarios; assisted in identifying symmetry transformations and provided code for parameterizing data augmentation and invariant policies.
- ✱ [Daniel Ordoñez-Apraez](#), Vladimir Kostic, Giulio Turrisi, Pietro Novelli, Carlos Mastalli, Claudio Semini, and Massimiliano Pontil. Dynamics harmonic analysis of robotic systems: Application in data-driven koopman modelling. In *6th Annual Learning for Dynamics & Control Conference*, pages 1318–1329. PMLR, 2024
- ✱ Zhi Su, Xiaoyu Huang, [Daniel Ordoñez-Apraez](#), Yunfei Li, Zhongyu Li, Qiayuan Liao, Giulio Turrisi, Massimiliano Pontil, Claudio Semini, Yi Wu, et al. Leveraging symmetry in rl-based legged locomotion control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024
Contribution: Substantial contributions to writing; theoretical analysis of reflection symmetries in model-free reinforcement learning for locomotion; assisted in identifying symmetry transformations and provided code for parameterizing data augmentation and invariant policies.
- ✱ [Daniel Ordoñez-Apraez](#), Mario Martin, Antonio Agudo, and Francesc Moreno-Noguer. On discrete symmetries of robotics systems: A group-theoretic and data-driven analysis. In *Proceedings of the 19th Robotics: Science and Systems Conference (RSS)*, Daegu, Republic of Korea, 2023. URL <https://www.roboticsproceedings.org/rss19/index.html>

1.5.1 Open-access software



Morphological Symmetries

`morpho_symm` is an open-access Python package offering a diverse library of robotic systems with their corresponding morphological symmetry groups identified, along with tutorials and examples for data augmentation, invariance/equivariance regularization, and the use of invariant/equivariant NN architectures.



Symmetric Learning

`symm_learning` is a torch-based library for machine learning with symmetry priors, specifically targeting robotics and dynamical systems. It provides utilities for symmetry-aware representation learning, statistics, linear algebra, optimization, and NN architectures.

1.6 Notation

Linear algebra

| | |
|------------------------------------|--|
| $x, \mathbf{x}, \mathbf{X}$ | A scalar, a vector, and a matrix |
| $\mathbf{x}_1 \oplus \mathbf{x}_2$ | Direct sum (stacking) of vectors, such that $\mathbf{x}_1 \oplus \mathbf{x}_2 := \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}$ |
| $\mathbf{A} \oplus \mathbf{B}$ | Direct sum of matrices, equivalent to $\mathbf{A} \oplus \mathbf{B} := \begin{bmatrix} \mathbf{A} & \mathbf{O} \\ \mathbf{O} & \mathbf{B} \end{bmatrix}$ |
| A, B, E | A linear operator |
| Id_n | Identity matrix in $\mathbb{R}^{n \times n}$ |
| $\delta_{i,j}$ | The Kronecker function, equal to 1 when $i = j$, and 0 when $i \neq j$ |

Sets, Vector Spaces, and Function Spaces

| | |
|---|---|
| $\mathcal{X}, \mathcal{Z}, \mathcal{H}, \mathcal{F}$ | A vector or Hilbert space |
| $\bar{\mathcal{X}}, \bar{\mathcal{Z}}, \bar{\mathcal{V}}$ | An irreducible \mathbb{G} -stable space (definition 2.9) |
| $\mathcal{X} \oplus \mathcal{Y}$ | Direct sum of vector spaces \mathcal{X} and \mathcal{Y} , such that if $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$, then $\mathbf{x} \oplus \mathbf{y} \in \mathcal{X} \oplus \mathcal{Y}$ |
| $\mathcal{L}_{\mathbb{P}_{\mathbf{x}}}^2 := \mathcal{L}_{\mathbb{P}_{\mathbf{x}}}^2(\mathcal{X}, \mathbb{R})$ | Space of square-integrable functions on \mathcal{X} . That is $\{f \mid \int_{\mathcal{X}} f(\mathbf{x}) ^2 \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) < \infty, f : \mathcal{X} \rightarrow \mathbb{R}\}$ |
| $\langle f, f' \rangle_{\mathbb{P}_{\mathbf{x}}}$ | Inner product $\langle f, f' \rangle_{\mathbb{P}_{\mathbf{x}}} := \int_{\mathcal{X}} f(\mathbf{x}) f'(\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x})$ |

Group and Representation Theory

| | |
|--|--|
| \mathbb{G} | A symmetry group |
| \mathbb{E}, SE | The Euclidean group and the Special Euclidean group |
| $\text{U}(\mathcal{X}), \text{GL}(\mathcal{X})$ | Unitary and General Linear groups on the vector space \mathcal{X} |
| \mathbb{C}_n | Cyclic group of order n |
| \mathbb{K}_4 | Klein four-group |
| g | A symmetry group element |
| $g \triangleright_{\mathcal{X}} \mathbf{x}$ | The left group action of g on \mathcal{X} (definition 2.2) |
| $\mathbf{x} \triangleleft_{\mathcal{X}} g$ | The right group action of g on \mathcal{X} (definition 2.3) |
| $\rho_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$ | A representation of the group \mathbb{G} on the vector space \mathcal{X} (definition 2.5) |
| $\bar{\rho}_k : \mathbb{G} \times \bar{\mathcal{H}} \rightarrow \bar{\mathcal{H}}$ | An irreducible representation of the group \mathbb{G} (definition 2.11) |
| $d_k := \bar{\rho}_k $ | Dimensionality of the irreducible representation $\bar{\rho}_k$ |
| m_k | Multiplicity of the irrep $\bar{\rho}_k$ in a given larger representation |
| n_{iso} | Number of distinct irreps present in a given larger representation. |
| $\rho_{\mathcal{X}} \oplus \rho_{\mathcal{Y}}$ | Direct sum of representations, equivalent to $\begin{bmatrix} \rho_{\mathcal{X}} & \\ & \rho_{\mathcal{Y}} \end{bmatrix}$ |
| $\mathbb{G}\mathbf{x}$ | The group orbit of \mathbf{x} , defined as $\mathbb{G}\mathbf{x} := \{g \triangleright \mathbf{x} \mid g \in \mathbb{G}\}$ |
| $\mathbb{G}_a \times \mathbb{G}_b$ | Direct product of groups \mathbb{G}_a and \mathbb{G}_b |

Probability Theory

| | |
|--|---|
| $(\mathcal{X}, \Sigma_{\mathcal{X}}, \mathbb{P}_{\mathbf{x}})$ | A probability space, where \mathcal{X} is the sample space, $\Sigma_{\mathcal{X}}$ is a σ -algebra of subsets of \mathcal{X} , and $\mathbb{P}_{\mathbf{x}}$ is a marginal probability measure on $\Sigma_{\mathcal{X}}$ |
| $\mathbf{x}, \mathbf{x} \in \mathcal{X}$ | A vector-valued random variable \mathbf{x} and a realization \mathbf{x} |
| $\mathbb{E}_{\mathbf{x}}[f(\mathbf{x})]$ | Expectation of $f(\mathbf{x})$ with respect to $\mathbb{P}_{\mathbf{x}}$ |
| $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ | Gaussian distribution over \mathbf{x} with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ |

Chapter 2

Preliminaries

This chapter presents the essential background needed to understand and contextualize the majority of results in this thesis. It is organized as follows:

- * [section 2.1](#): reviews deterministic analytic models of robot dynamics—used in [chapters 3](#) and [4](#) to define symmetries and the concept of dynamics equivalence. Also provides a friendly robotics introduction for the machine learning audience.
- * [section 2.2](#): reviews stochastic modeling of robot dynamics for contact-rich scenarios. Formulates robot dynamics modeling and state estimation as conditional distribution approximation problems, used in [chapters 6](#) to [8](#).
- * [section 2.3](#): provides a friendly introduction to the fundamental concepts and results from group and representation theory leveraged in the thesis.
- * [section 2.4](#): surveys key results from probability theory and optimal transport under symmetry priors.

The dynamics of a robot—like those of a falling apple or a moving car—are governed by classical mechanics, that is, by Newton’s *deterministic* second law of motion:

$$\underbrace{\text{Mass} \times \text{Acceleration}}_{\text{Inertial Forces}} = \underbrace{\text{Force}}_{\text{Moving forces}}$$


If both the mass distribution and all instantaneous forces are known—or can be inferred from measurements—the robot’s dynamics can, in principle, be described exactly by an ideal deterministic model, formalized as an [ordinary differential equation \(ODE\)](#) (see [section 2.1](#)). In reality, however, the robot’s state is only partially observable: some state variables cannot be directly measured, sensors are noisy, and neither the mass distribution nor the full set of forces acting on the robot are fully accessible. This partial observability introduces epistemic and measurement uncertainty in the robot’s state and transitions, motivating the use of stochastic data-driven models introduced in [section 2.2](#).

The following section, which reviews the analytical mechanics principles underlying robot dynamics, is deeply influenced by—and necessarily falls short of—the seminal works of [Lanczos \(2020\)](#) and [Wieber \(2006\)](#), which I highly recommend to readers seeking a deeper treatment.

2.1 Analytical models of robot dynamics

The dynamics of a robotic system are commonly modeled within the framework of *analytical mechanics*, a branch of classical mechanics that derives the equations of motion of any mechanical system from two scalar quantities: the *kinetic energy* and the *work function*.


Within this framework, the equations of motion of any mechanical system—recoverable as a generalized form of Newton’s second law—are derived from the unifying *principle of least action*. This principle states that the physically realized trajectories of a mechanical system are those that render stationary (and, in many cases, minimize) a scalar functional known as the *action*, defined as the time integral of the excess of kinetic energy over the work function, which often reduces to the integral of the difference between kinetic and potential energy (Lanczos, 2020). Before defining the action, we must first choose a set of coordinates to model the system’s state. The first hallmark of analytical mechanics is to describe a system’s configuration using the minimum number of independent coordinates required to uniquely specify it. These coordinates are known as *generalized coordinates*.




Stationary points of a function f satisfy $df(x)/dx = 0$ and can be minima, maxima, or inflection points.

Generalized coordinates and the robot’s configuration space

The state of a robot is determined by the positions and velocities of its $n_b \in \mathbb{N}$ links, or body parts. However, because these links are typically connected through joints that constrain their relative motion, the robot state is most naturally parameterized using the positions and velocities of the system’s **degrees of freedom (DoF)**, i.e., the minimum number of independent coordinates required to uniquely describe the configuration of the n_b links.



The Mini-Cheetah quadruped



Q a smooth manifold, and $\mathcal{T}_q Q$ its tangent space at q , carrying the velocity vector \dot{q} .

For example, although the Mini-Cheetah quadruped is composed of $n_b = 13$ rigid bodies—requiring $13 \times 6 = 78$ coordinates to describe the position and orientation of each body—its configuration can be fully described using only $n_q = 6 + 12 = 18$ DoF: six DoF describing the position and orientation of the robot base in three-dimensional space, and twelve DoF corresponding to the relative angles of the three revolute joints defining the shoulder abduction, shoulder flexion, and elbow flexion of each leg.

Formally, for a robot with $n_q \in \mathbb{N}$ DoF, the state is described by generalized position coordinates $q \in Q \subset \mathbb{R}^{n_q}$ and generalized velocity coordinates $\dot{q} \in \mathcal{T}_q Q \subset \mathbb{R}^{n_q}$. Here, Q denotes the constrained configuration manifold—a smooth manifold encoding all admissible position configurations of the system together with its kinematic constraints—while $\mathcal{T}_q Q$ denotes the space of generalized velocities, corresponding to the local tangent space of the manifold at q (Lanczos, 2020, I.IV). Hence, the real-world robot state evolution can be numerically represented as the time evolution of a point $(q, \dot{q}) \in \mathcal{T}Q \subseteq \mathbb{R}^{2n_q}$ in the tangent bundle, defined as the disjoint union of all tangent spaces over the configuration manifold, $\mathcal{T}Q := \bigsqcup_{q \in Q} \mathcal{T}_q Q$.

To build some intuition, let us return to the Mini-Cheetah example. In this case, the robot’s configuration manifold is defined as $\mathcal{Q} := \mathbb{SE}_3 \times \mathbb{SO}_2^{12} \subseteq \mathbb{R}^{18}$, where \mathbb{SE}_3 denotes the special Euclidean group encoding the position and orientation of the robot base in three-dimensional space, and each \mathbb{SO}_2 encodes the angle of one of the twelve unbounded revolute joints connecting the legs to the base¹.

A second hallmark of analytical mechanics is the observation that the inertial forces acting on the system—corresponding to the *mass times acceleration* term in Newton’s second law—can be derived from a single scalar quantity: the system’s *kinetic energy*. Likewise, (some type of) generalized forces—corresponding to the *moving force* term in Newton’s formulation—can be derived from another scalar quantity: the system’s *work function*. In order to define both quantities in terms of the generalized coordinates, we first need to introduce the *forward kinematics* and *Jacobian matrix* functions.

Forward kinematics and the Jacobian matrix function Given the generalized position coordinates $\mathbf{q} \in \mathcal{Q}$, the forward kinematics function $\mathbf{f}_k : \mathcal{Q} \rightarrow \mathbb{SE}_3^{n_b}$ determines the position and orientation of each of the robot’s bodies:

$$\mathbf{X}_{1:n_b} := \begin{bmatrix} \mathbf{X}_1 \\ \vdots \\ \mathbf{X}_{n_b} \end{bmatrix} = \mathbf{f}_k(\mathbf{q}). \quad \text{with} \quad \mathbf{X}_i := \begin{bmatrix} \mathbf{R}_i & \mathbf{r}_i \\ \mathbf{0} & 1 \end{bmatrix} \in \mathbb{SE}_3, \quad (2.1)$$

where $\mathbf{r}_i \in \mathbb{R}^3$ and $\mathbf{R}_i \in \mathbb{SO}_3$ denote the position and orientation of the i -th body, respectively. Differentiating this relationship with respect to time yields the robot’s *Jacobian matrix* function $\mathbf{J} : \mathcal{Q} \rightarrow \mathbb{R}^{(6 \cdot n_b) \times n_q}$ —a configuration-dependent linear relation between $\dot{\mathbf{q}} \in \mathcal{T}_{\mathbf{q}}\mathcal{Q}$ and the velocities of the robot’s rigid bodies, that is:

$$\begin{aligned} \frac{d\mathbf{X}_{1:n_b}}{dt} &= \frac{d}{dt}(\mathbf{f}_k(\mathbf{q})) = \frac{\partial \mathbf{f}_k(\mathbf{q})}{\partial \mathbf{q}} \frac{d\mathbf{q}}{dt} \\ \dot{\mathbf{x}}_{1:n_b} &= \begin{bmatrix} \dot{\mathbf{x}}_1 \\ \vdots \\ \dot{\mathbf{x}}_{n_b} \end{bmatrix} = \mathbf{J}(\mathbf{q}) \dot{\mathbf{q}} = \begin{bmatrix} \mathbf{J}_1(\mathbf{q}) \\ \vdots \\ \mathbf{J}_{n_b}(\mathbf{q}) \end{bmatrix} \dot{\mathbf{q}}. \quad \text{with} \quad \dot{\mathbf{x}}_i := \begin{bmatrix} \dot{\mathbf{r}}_i \\ \dot{\boldsymbol{\omega}}_i \end{bmatrix} = \mathbf{J}_i(\mathbf{q}) \dot{\mathbf{q}} \in \mathfrak{se}. \end{aligned} \quad (2.2)$$

where $\mathbf{J}_i : \mathcal{Q} \rightarrow \mathbb{R}^{6 \times n_q}$ denotes the Jacobian matrix of the i -th body, determining the linear $\dot{\mathbf{r}}_i \in \mathbb{R}^3$ and angular $\dot{\boldsymbol{\omega}}_i \in \mathbb{R}^3$ velocities of that body, referenced to a frame attached to its **Center of Mass (CoM)**. Hence, $\mathbf{J}(\mathbf{q})$ is composed of n_b blocks of size $6 \times n_q$, one for each body.

Note that although the special Euclidean group is six-dimensional, $|\mathbb{SE}_3| = 6$, we adopt the common convention of representing a body configuration as a homogeneous transformation matrix, $\mathbf{X}_i \in \mathbb{R}^{4 \times 4}$; the standard matrix representation of elements of \mathbb{SE}_3 . Furthermore, the linear and angular velocities of a body—defining an element of the special Euclidean group Lie algebra \mathfrak{se} , also with $|\mathfrak{se}| = 6$ (see Selig (2004); Sola et al. (2018))—are represented as a 6D vector, $\dot{\mathbf{x}}_i \in \mathbb{R}^6$, following the standard convention of spatial algebra (Featherstone, 2008), used widely in the robot dynamics simulation and control literature Carpentier et al. (2019); Mastalli et al. (2020a).

For the Mini-Cheetah example (section 2.1), the forward kinematics function maps $\mathbf{q} \in \mathcal{Q}$ —composed of the base configuration together with twelve joint angles—to the position and orientation of the base of the robot and the shoulder, forearm, and lower-arm links of each leg in three-dimensional space.

¹Why \mathbb{SO}_2 ? For unbounded revolute joints q and $q + 2\pi$ represent the same angle, hence for these joints the configuration manifold is the unit circle, isomorphic to the \mathbb{SO}_2 manifold.

Correspondingly, the Jacobian matrix function maps the eighteen generalized velocities to the linear and angular velocities of each of the thirteen rigid bodies composing the robot.

Kinetic energy For a robotic system, the kinetic energy is defined as the sum of the linear and angular kinetic energies of each of the robot's n_b rigid bodies:

$$\text{Kinetic Energy} := \frac{1}{2} \sum_{i=1}^{n_b} \left(m_i \|\dot{\mathbf{r}}_i\|^2 + \boldsymbol{\omega}_i^\top \mathbf{I}_i \boldsymbol{\omega}_i \right) = \frac{1}{2} \sum_{i=1}^{n_b} \dot{\mathbf{x}}_i^\top \underbrace{\begin{bmatrix} m_i \text{Id}_3 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_i \end{bmatrix}}_{\mathbf{V}_i \in \mathbb{S}_{++}^6} \dot{\mathbf{x}}_i, \quad (2.3a)$$

where $m_i \in \mathbb{R}_{>0}$ and $\mathbf{I}_i \in \mathbb{S}_{++}^3 \subset \mathbb{R}^{3 \times 3}$ denote the i^{th} body's mass and rotational inertia matrix, referenced with respect to the frame attached to the body's CoM. Here, $\mathbf{V}_i \in \mathbb{S}_{++}^6 \subset \mathbb{R}^{6 \times 6}$ is known as the body's spatial inertia matrix—defined as an element of the set of 6-dimensional symmetric positive definite matrices; \mathbb{S}_{++}^6 (Featherstone, 2008).

Using the Jacobian matrix function defined in (2.2), the kinetic energy can be equivalently expressed in terms of the generalized coordinates as

$$T(\mathbf{q}, \dot{\mathbf{q}}) := \frac{1}{2} \dot{\mathbf{q}}^\top \mathbf{M}(\mathbf{q}) \dot{\mathbf{q}}, \quad \text{with} \quad \mathbf{M}(\mathbf{q}) := \sum_{i=1}^{n_b} \mathbf{J}_i(\mathbf{q})^\top \mathbf{V}_i \mathbf{J}_i(\mathbf{q}). \quad (2.3b)$$

Here, $\mathbf{M} : \mathcal{Q} \rightarrow \mathbb{S}_{++}^{n_q} \subset \mathbb{R}^{n_q \times n_q}$ denotes the generalized mass matrix function, mapping each configuration \mathbf{q} to a symmetric positive definite generalized mass matrix $\mathbf{M}(\mathbf{q})$. Just as the state of the real-world robotic system is numerically represented by a point $(\mathbf{q}, \dot{\mathbf{q}})$, the generalized mass matrix $\mathbf{M}(\mathbf{q})$ represents algebraically the robot's inertial properties at configuration \mathbf{q} .

Generalized force To obtain a generalized form of Newton's second law in the space of generalized coordinates, we must express all impressed Cartesian forces acting on the robot as forces acting along the admissible directions of motion of the robot's rigid bodies, that is, in the space of generalized coordinates. Let $\mathbf{h}_i = \begin{bmatrix} \mathbf{f}_i \\ \mathbf{p}_i \end{bmatrix} \in \mathbb{R}^6$ denote the impressed wrench, composed of the force $\mathbf{f}_i \in \mathbb{R}^3$ and torque $\mathbf{p}_i \in \mathbb{R}^3$ acting on the robot's i -th body CoM², and let $d\mathbf{X}_i \in \mathfrak{se}_3$ denote an admissible displacement of the same body consistent with the kinematic constraints. The instantaneous work performed by these impressed wrenches over admissible displacements is then defined as

$$\overline{dw} = \sum_{i=1}^{n_b} \langle \mathbf{h}_i, d\mathbf{X}_i \rangle = \sum_{i=1}^{n_b} \langle \mathbf{J}_i(\mathbf{q})^\top \mathbf{h}_i, d\mathbf{q} \rangle = \langle \boldsymbol{\tau}(\mathbf{q}, \dot{\mathbf{q}}), d\mathbf{q} \rangle, \quad \text{s.t.} \quad \boldsymbol{\tau}(\mathbf{q}, \dot{\mathbf{q}}) := \sum_{i=1}^{n_b} \mathbf{J}_i(\mathbf{q})^\top \mathbf{h}_i. \quad (2.4)$$

Here, $\boldsymbol{\tau} : \mathcal{TQ} \mapsto \mathcal{T}_q \mathcal{Q}^* \subseteq \mathbb{R}^{n_q}$ denotes the *generalized force vector field*, which assigns to each state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$ an instantaneous vector of forces acting along the directions of motion of each DoF. This instantaneous *generalized force* is the representation, in generalized coordinates, of the impressed Cartesian wrenches $\{\mathbf{h}_i\}_{i=1}^{n_b}$.

In practice, we will consider the generalized force vector field to be decomposed into components representing different sources of wrenches acting on the robot, including gravitational effects $\boldsymbol{\tau}_{\text{grav}}$,

²Given that generalized coordinates parameterize motions compatible with the system's kinematic constraints, force components enforcing these constraints perform no work and are therefore ignored.

joint damping and friction forces τ_{damp} , external contact forces τ_{cont} , and control or actuation forces τ_{ctrl} , namely:

$$\boldsymbol{\tau}(\mathbf{q}, \dot{\mathbf{q}}) = \boldsymbol{\tau}_{\text{grav}}(\mathbf{q}) + \boldsymbol{\tau}_{\text{damp}}(\mathbf{q}, \dot{\mathbf{q}}) + \boldsymbol{\tau}_{\text{cont}}(\mathbf{q}, \dot{\mathbf{q}}) + \boldsymbol{\tau}_{\text{ctrl}}(\mathbf{q}, \dot{\mathbf{q}}). \quad (2.5)$$

Having defined the generalized coordinates, mass, and forces, we proceed to express Newton's second law in the space of generalized coordinates via the principle of least action.

Principle of least action In analytical mechanics, the equations of motion of a system are derived through the minimization of a scalar functional known as the *action*, denoted by A (Lanczos, 2020). The action is defined as the time integral of the excess of kinetic energy over the work performed by moving forces along an arbitrary state trajectory $\vec{\mathbf{q}}_{t_0:t_1} = (\mathbf{q}_t, \dot{\mathbf{q}}_t)_{t \in [t_0, t_1] \subseteq \mathbb{T}}$ with fixed endpoints $\mathbf{q}_{t_0}, \mathbf{q}_{t_1} \in \mathcal{Q}$, namely

$$\text{Action: } A(\vec{\mathbf{q}}_{t_0:t_1}) = \int_{t_0}^{t_1} T(\mathbf{q}_t, \dot{\mathbf{q}}_t) - U(\mathbf{q}_t, \dot{\mathbf{q}}_t) dt. \quad (2.6)$$

Here, $U : \mathcal{TQ} \rightarrow \mathbb{R}$ denotes the work function, measuring the instantaneous work performed by moving forces, and $T : \mathcal{TQ} \rightarrow \mathbb{R}$ is the system's kinetic energy, defined in eq. 2.3b.

Analytical solutions to this variational problem—namely, the minimization of eq. 2.6—require the work function to be integrable. This condition is generally violated in the presence of damping or friction forces, contact forces, or arbitrary control inputs (i.e., τ_{damp} , τ_{cont} , and τ_{ctrl} in eq. 2.5) (Lanczos, 2020). For this reason, we first consider the idealized setting in which the system is subject only to gravitational forces $\boldsymbol{\tau}_{\text{grav}}$. In this case, the work function reduces to the gravitational potential energy function³ $V : \mathcal{Q} \rightarrow \mathbb{R}$. Consequently, the action functional simplifies to

$$A(\vec{\mathbf{q}}_{t_0:t_1}) := \int_{t_0}^{t_1} L(\mathbf{q}_t, \dot{\mathbf{q}}_t) dt = \int_{t_0}^{t_1} T(\mathbf{q}_t, \dot{\mathbf{q}}_t) - V(\mathbf{q}_t) dt, \quad (2.7)$$

where $L(\mathbf{q}, \dot{\mathbf{q}})$ denotes the system's Lagrangian function—defined as the excess of kinetic energy over gravitational potential energy—which will play a central role in the definition of symmetries in chapter 3. Minimizing eq. 2.7 amounts to finding a trajectory $\vec{\mathbf{q}}_{t_0:t_1}^*$ that renders the action stationary, i.e., such that the first variation of A with respect to infinitesimal perturbations of the trajectory vanishes. Trajectories satisfying this condition obey, at every time instant, the following necessary condition (Lanczos, 2020):

$$\frac{d}{dt} \frac{\partial L(\mathbf{q}_t, \dot{\mathbf{q}}_t)}{\partial \dot{\mathbf{q}}_t} - \frac{\partial L(\mathbf{q}_t, \dot{\mathbf{q}}_t)}{\partial \mathbf{q}_t} = 0. \quad (2.8)$$

Euler–Lagrange equations of motion Eq. 2.8 is known as the *Euler–Lagrange equations of motion* in analytical mechanics. It defines a second-order ODE that translates Newton's second law to the space of generalized coordinates and fully characterizes the system's dynamics for any mechanical

³Defined for robotic systems as $V(\mathbf{q}) := \sum_{i=1}^{n_b} m_i \langle \mathbf{g}, \mathbf{r}_i(\mathbf{q}) \rangle + \text{const}$, where \mathbf{g} denotes the gravitational acceleration vector, and $\mathbf{r}_i(\mathbf{q}) \in \mathbb{R}^3$ is the position of the i -th body CoM, obtained from the forward kinematics function (eq. 2.1).

system:

$$\frac{d}{dt} \frac{\partial T(\mathbf{q}_t, \dot{\mathbf{q}}_t)}{\partial \dot{\mathbf{q}}_t} - \frac{\partial T(\mathbf{q}_t, \dot{\mathbf{q}}_t)}{\partial \mathbf{q}_t} = \frac{\partial V(\mathbf{q}_t)}{\partial \mathbf{q}_t} \quad (2.9a)$$

$$\mathbf{M}(\mathbf{q}_t) \ddot{\mathbf{q}}_t + \mathbf{N}(\mathbf{q}_t, \dot{\mathbf{q}}_t) \dot{\mathbf{q}}_t = \boldsymbol{\tau}_{\text{grav}}(\mathbf{q}_t) \quad (2.9b)$$

$$\underbrace{\mathbf{M}(\mathbf{q}_t) \ddot{\mathbf{q}}_t}_{\text{Inertial forces}} = \underbrace{\boldsymbol{\tau}_{\text{grav}}(\mathbf{q}_t) + \boldsymbol{\tau}_{\text{nonlin}}(\mathbf{q}_t, \dot{\mathbf{q}}_t)}_{\text{Moving forces}}. \quad (2.9c)$$

The left-hand side of eq. 2.9, derived from the kinetic energy, represents the inertial forces appearing in Newton's second law. The right-hand side, derived from the work function, represents the moving forces acting on the robot. In this setting, these forces are composed of gravitational effects and nonlinear velocity-dependent forces $\boldsymbol{\tau}_{\text{nonlin}}(\mathbf{q}, \dot{\mathbf{q}}) := -\mathbf{N}(\mathbf{q}, \dot{\mathbf{q}}) \dot{\mathbf{q}}$, which arise from gyroscopic effects associated with the dynamics of moving rigid bodies (Bloch et al., 1996; Wieber, 2006) and, from a geometric perspective, define the curvature of the configuration manifold \mathcal{Q} (Lanczos, 2020)⁴.

Having defined the equations of motion for robotic systems, we now shift our focus to the structural properties they inherit from the symmetries of classical mechanics. This provides a first step toward understanding the different types of symmetry that arise in robot dynamics, discussed in chapter 3.

Ignorable variables and Noether's theorem In analytical mechanics, one often encounters *ignorable* or *kinosthenic* coordinates, defined as generalized position coordinates q_i that do not explicitly appear in the Lagrangian function, i.e., such that $\partial L(\mathbf{q}, \dot{\mathbf{q}}) / \partial q_i = 0$ (Lanczos, 2020). These coordinates play a central role in Noether's invariance theorem (Noether, 1983), which states that for every ignorable coordinate q_i there exists a corresponding conserved quantity, given by the generalized momentum conjugate to q_i , namely

$$\text{if } \frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial q_i} = 0, \text{ and eq. 2.8 holds, then } \frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{q}_i} = \text{const.} \quad (2.10)$$

In classical mechanics, the existence of ignorable coordinates is closely related to the fact that Newtonian physics is independent of the chosen inertial reference frame. This means that the dynamics do not depend on the absolute value of the ignorable coordinates q_i , but only on their time derivatives \dot{q}_i .

To illustrate this, consider the dynamics of the Atlas robot in fig. 1.1. Let $\mathbf{r}_B^o = [r_x, r_y, r_z] \in \mathbb{R}^3$ and $\mathbf{R}_B^o = [\theta_x, \theta_y, \theta_z] \in \mathbb{R}^3$ denote the position and orientation of the robot's base in a chosen inertial frame o . These six coordinates are part of the generalized coordinates \mathbf{q} and are ignorable, since any translation, rotation, or reflection of the reference frame—e.g., $\mathbf{r}_B^{g\Delta x \triangleright o} = [r_x + \Delta x, r_y, r_z]$ or $\mathbf{r}_B^{g-z \triangleright o} = [r_x, r_y, -r_z]$ —leaves both the robot's Lagrangian and its dynamics invariant. Consequently, deriving the equations of motion using eq. 2.9 in any such transformed reference frame yields equations that are linearly equivalent to the original ones, as they describe the same underlying physical behavior (Wheeler, 2014).

In this work, we adopt an interpretation of ignorable coordinates analogous to that introduced by Emmy Noether: each ignorable variable is associated with a symmetry group composed of translations and

⁴Formally, $\boldsymbol{\tau}_{\text{nonlin}}$ arise from the kinetic energy and are therefore classified as inertial forces in the framework of Lanczos (2020). In robotics, however, these terms typically reduce to centrifugal and Coriolis forces and are commonly grouped with other moving forces (Wieber, 2006).

reflections acting on that variable. Collectively, the translations and rotations of absolute position and orientation variables define the fundamental symmetry group of classical mechanics—the Euclidean group \mathbb{E}_3 (see fig. 1.1). This group plays a central role in the definition and interpretation of symmetries in robot dynamics, covered in detail in chapter 3.

2.2 Data-driven stochastic dynamics models

In section 2.1, we introduced an analytical robot-dynamics model that describes Newtonian mechanics in generalized coordinates through a system of second-order ordinary differential equations (eq. 2.9). This ODE yields deterministic dynamics for the system state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$, which, for most real-world robotic systems, can be written as the following underactuated control-affine system:

$$\mathbf{M}(\mathbf{q}_t)\ddot{\mathbf{q}}_t = \tau_{\text{grav}}(\mathbf{q}_t) + \tau_{\text{nonlin}}(\mathbf{q}_t, \dot{\mathbf{q}}_t) + \underbrace{\bar{\mathbf{B}}\mathbf{a}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{z}_t)}_{\tau_{\text{ctrl}}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{z}_t)} + \tau_{\text{unobs}}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{z}_t), \quad (2.11)$$

Here, $\bar{\mathbf{B}} \in \mathbb{R}^{n_q \times n_{\text{ctrl}}}$ is the robot’s actuation matrix, mapping the $n_{\text{ctrl}} \leq n_q$ controlled DoF to the full set of n_q generalized forces. Moreover, $\mathbf{a}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{z}_t) \in \mathcal{A}$ denotes the control torques applied at time t as a function of the current state and the robot–environment interaction state $\mathbf{z}_t \in \mathcal{Z}$. The variable \mathbf{z} is assumed to encode the proprioceptive and exteroceptive information required for policy decision-making, including task information and the state of the robot in the environment, in particular the robot’s contact mode—the set of active contact surfaces and contact wrenches. Finally, $\tau_{\text{unobs}}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{z}_t)$ denotes generalized forces arising from friction, damping, contact forces, and approximation errors, which are only partially observable.

Uncertainty of forcing terms Although the underlying physical phenomenon is deterministic, epistemic and measurement uncertainty in the partially observable generalized forces $\tau_{\text{unobs}}(\cdot) \in \mathcal{T}_q\mathcal{Q}^*$ motivate modelling robot dynamics in contact-rich environments via a stochastic differential equation (SDE) (Thrun et al., 2005) that defines a controlled nonlinear diffusion process (Theodorou et al., 2010) of the form:

$$\dot{\mathbf{s}}_t = \underbrace{f(\mathbf{s}_t) + \mathbf{B}\mathbf{a}(\mathbf{s}_t)}_{\text{Controlled deterministic drift}} + \underbrace{\mathbf{D}(\mathbf{s}_t)\boldsymbol{\xi}_t}_{\text{State-dependent diffusion}}, \quad \text{with } \boldsymbol{\xi}_t \sim \mathcal{N}(\mathbf{0}, \text{Id}_{|\mathcal{S}|}), \quad \mathbf{s}_t := \begin{bmatrix} \mathbf{q}_t \\ \dot{\mathbf{q}}_t \\ \mathbf{z}_t \end{bmatrix} \in \mathcal{S} := \mathcal{TQ} \times \mathcal{Z}. \quad (2.12a)$$

In discrete time, this can be approximated as:

$$\mathbf{s}_{t+\Delta t} = \underbrace{\mathbf{s}_t + f(\mathbf{s}_t)\Delta t + \mathbf{B}\mathbf{a}(\mathbf{s}_t)\Delta t}_{F_{\Delta t}(\mathbf{s}_t, \mathbf{a}_t)} + \mathbf{D}(\mathbf{s}_t)\sqrt{\Delta t}\boldsymbol{\xi}_t, \quad \text{with } \boldsymbol{\xi}_t \sim \mathcal{N}(\mathbf{0}, \text{Id}_{|\mathcal{S}|}). \quad (2.12b)$$

Here, $\mathbf{s} \in \mathcal{S}$ denotes the world state, composed of the robot state $(\mathbf{q}, \dot{\mathbf{q}})$ and the robot–environment interaction state $\mathbf{z} \in \mathcal{Z}$. In particular, because \mathbf{z} includes the robot’s contact mode and task information, the world state distinguishes between states with the same $(\mathbf{q}, \dot{\mathbf{q}})$ but different contact modes and control tasks. The map $f : \mathcal{S} \rightarrow \mathbb{R}^{|\mathcal{S}|}$ is a deterministic drift vector field defined by the known components of the dynamics in eq. 2.11; $\mathbf{B} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ is the actuation matrix in \mathcal{S} ; and $\mathbf{D} : \mathcal{S} \rightarrow \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ is a state-dependent diffusion matrix that modulates stochasticity as a function

of the current state. Through \mathbf{D} , the Gaussian perturbation $\boldsymbol{\xi}_t$ captures epistemic and measurement uncertainty in the unobservable generalized forces $\boldsymbol{\tau}_{\text{unobs}}(\cdot)$ and in the robot state $(\mathbf{q}, \dot{\mathbf{q}})$.

State-conditional transition distribution As is common in the literature, we model the system dynamics by approximating the controlled discrete-time Markovian dynamics in eq. 2.12b. This induces a state-conditional Gaussian probability distribution $\mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, often referred to as the Δt -transition distribution (Thrun et al., 2005), defined by the conditional law

$$\mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}(\cdot | \mathbf{a}_t, \mathbf{s}_t) \approx \mathcal{N}\left(F_{\Delta t}(\mathbf{s}_t, \mathbf{a}_t), \mathbf{D}(\mathbf{s}_t)\mathbf{D}(\mathbf{s}_t)^\top \Delta t\right), \quad \forall \mathbf{s}_t \in \mathcal{S}, \mathbf{a}_t \in \mathcal{A}, t \in \mathbb{R}. \quad (2.13)$$

This distribution quantifies the probability of transitioning into the subset \mathcal{S} after a Δt time step when taking action $\mathbf{a} \in \mathcal{A}$ in state $\mathbf{s} \in \mathcal{S}$.⁵ Formally, for each pair (\mathbf{s}, \mathbf{a}) , $\mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}$ defines a Gaussian conditional probability measure over next states, with mean and covariance given by the deterministic drift and diffusion matrix in eq. 2.12b (Theodorou et al., 2010). The Gaussian form of the conditional distribution follows from modelling the discrete-time world dynamics with eq. 2.12b, which implicitly assumes that $\mathbf{z} \in \mathcal{Z}$ contains the information describing the true robot–environment contact mode and the corresponding contact-mode transitions over the next Δt seconds. Under this assumption, stochasticity arises from uncertainty in the contact forces associated with the true contact mode at t and $t + \Delta t$, which are assumed to be Gaussian distributed for sufficiently small Δt .

Consequently, the multimodality of the transition distribution expected from uncertainty in the contact mode—for instance, uncertainty about whether the robot at time $t + \Delta t$ will gain or lose a given contact, or whether a given contact will be in a sticking or sliding friction regime—does not appear in this state-conditional transition distribution, as it is assumed that \mathbf{z} fully determines the contact mode at t and $t + \Delta t$. Instead, such multimodality arises from partial observability of the state.

Partially observable dynamics Given that the world state $\mathbf{s} \in \mathcal{S}$ is only partially observable, we assume the existence of a stochastic *observation model*, $\mathbb{P}_{\mathcal{O}|\mathcal{S}} : \Sigma_{\mathcal{O}} \times \mathcal{S} \rightarrow [0, 1]$, which defines the conditional probability of an observation $\mathbf{o}_t \in \mathcal{O}$ given a state $\mathbf{s}_t \in \mathcal{S}$. Let \mathbf{o}_t denote the component of the state that is directly measurable from sensors, so that the stochasticity in $\mathbb{P}_{\mathcal{O}|\mathcal{S}}$ directly captures measurement uncertainty. In practice, these observations commonly include joint positions, velocities, accelerations, and forces; RGBD images; point clouds; IMU readings; battery level; and related sensor-noise characteristics.

From a modelling perspective, these observations $\mathbf{o}_t \in \mathcal{O}$ are the variables we aim to predict in time. Their stochastic dynamics are defined by composing $\mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}$ and $\mathbb{P}_{\mathcal{O}|\mathcal{S}}$, yielding the following observation transition distribution:

$$\mathbb{P}_{\mathcal{O}|\mathcal{A}\mathcal{S}}^{\Delta t}(\mathcal{O} | \mathbf{a}_t, \mathbf{s}_t) = \int_{\mathcal{S}} \mathbb{P}_{\mathcal{O}|\mathcal{S}}(\mathcal{O} | \mathbf{s}') \mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}(d\mathbf{s}' | \mathbf{a}_t, \mathbf{s}_t), \quad \forall \mathbf{s}_t \in \mathcal{S}, \mathcal{O} \subseteq \mathcal{O}, \mathbf{a} \in \mathcal{A}, t \in \mathbb{R}, \quad (2.14)$$

Note that $\mathbb{P}_{\mathcal{O}|\mathcal{A}\mathcal{S}}^{\Delta t}$ represents the ideal model of robot–environment dynamics that we aim to approximate with data-driven models. However, because $\mathbf{s}_t \in \mathcal{S}$ is only partially observable, the next-best attainable ideal target is a model that leverages optimal state estimation from observation history. Let $\vec{\mathcal{O}} := \bigcup_{t \geq 0} \mathcal{O}^{t+1}$ denote the space of all finite observation histories, so that for a given state \mathbf{s}_t

⁵We implicitly assume that \mathbf{s} is the realization of a random variable s on the measurable space $(\mathcal{S}, \Sigma_{\mathcal{S}})$, with law $\mathbb{P}_{\mathbf{s}}$.

there exists an observation history $\vec{o}_t := (\mathbf{o}_0, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}$ (Thrun et al., 2005). Then, our target observation dynamics model is defined by the conditional probability distribution

$$\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \vec{o}_t) = \int_{\mathcal{S}} \mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \mathbf{s}) \mathbb{P}_{\mathbf{s}|\vec{o}}^{\Delta t}(d\mathbf{s}|\vec{o}_t), \quad \forall \vec{o}_t \in \vec{\mathcal{O}}, \mathbb{O} \subseteq \mathcal{O}, \mathbf{a}_t \in \mathcal{A}, t \in \mathbb{R}, \quad (2.15)$$

quantifying the probability of observing any $\mathbf{o}_{t+\Delta t} \in \mathbb{O}$ after a Δt time step when taking action $\mathbf{a} \in \mathcal{A}$ at time t , given the history of past observations \vec{o}_t . Here, $\mathbb{P}_{\mathbf{s}|\vec{o}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \rightarrow [0, 1]$ denotes the true conditional estimation of states given past observation histories.

Note that although $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ is assumed to be Gaussian, partial observability of the state makes $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}$ potentially non-Gaussian and multimodal. Intuitively, this multimodality captures uncertainty in the unobservable robot–environment contact mode: for a given observation history \vec{o}_t , multiple contact modes may be plausible, rendering $\mathbb{P}_{\mathbf{s}|\vec{o}}^{\Delta t}(\cdot|\vec{o}_t)$ multimodal, which in turn gives rise to multimodal observation transition distributions $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}$.

Data-driven stochastic dynamics modelling In this context, the learning problem of data-driven stochastic dynamics modelling can be formalized in its most general form as the approximation of the conditional distribution $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}$ (eq. 2.15) from data, using a stochastic model

$\Phi_{\theta} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \vec{\mathcal{O}} \rightarrow [0, 1]$, parameterized by $\theta \in \Theta$. Namely, by:

$$\theta^* = \arg \min_{\theta \in \Theta} \mathcal{L}_{\text{err}}(\Phi_{\theta}) = \arg \min_{\theta \in \Theta} \mathbb{E}_{\vec{o}_t, \mathbf{a}_t \sim \mathbb{P}_{\vec{o}, \mathbf{a}}} \left[\mathcal{D} \left(\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}(\cdot|\mathbf{a}_t, \vec{o}_t) \parallel \Phi_{\theta}(\cdot|\mathbf{a}_t, \vec{o}_t) \right) \right], \quad (2.16)$$

Where $\mathcal{D} : (\mathcal{O}, \Sigma_{\mathcal{O}}) \times (\mathcal{O}, \Sigma_{\mathcal{O}}) \rightarrow [0, +\infty)$ denotes a divergence measure between probability distributions, such as the Kullback–Leibler divergence, and $\mathbb{P}_{\vec{o}, \mathbf{a}}$ is a joint distribution over observation histories and actions, that weights the approximation error of Φ_{θ} across different regions of the observation-action space, assigning higher weights to conditioning points $(\vec{o}_t, \mathbf{a}_t)$ that are more likely under $\mathbb{P}_{\vec{o}, \mathbf{a}}$. In practice, $\mathbb{P}_{\vec{o}, \mathbf{a}}$ is typically induced by a chosen control policy, or a family of control policies, defining the distribution of trajectories under nominal operation of the robot (Thrun et al., 2005).

In practice, the expectation over $\mathbb{P}_{\vec{o}, \mathbf{a}}$ is intractable and only empirically approximable from data via Monte Carlo estimation, using a dataset of N observation histories and actions collected from the robot $\mathbb{D} = \{(\vec{o}_i, \mathbf{a}_i) \sim \mathbb{P}_{\vec{o}, \mathbf{a}}\}_{i=1}^N$. Likewise, the divergence \mathcal{D} is typically intractable and only empirically approximable from data. Note that, since $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}$ is expected to be multimodal—as a given observation history might not be able to provide sufficient information to disambiguate the true contact mode—the stochastic model Φ_{θ} should not be parameterized with a unimodal parametric distribution family, e.g., a Gaussian.

Note that, by minimizing eq. 2.16, the dynamics model Φ_{θ} is implicitly tackling the problem of optimal state estimation from observation history, since $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{o}}^{\Delta t}$ is defined by the composition of $\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}$ and $\mathbb{P}_{\mathbf{s}|\vec{o}}^{\Delta t}$ (see eq. 2.15).

Data-driven stochastic state estimation The learning problem of data-driven stochastic state estimation can be formalized, in its most general form, as the approximation of the conditional distribution $\mathbb{P}_{\mathbf{s}|\vec{o}}^{\Delta t}$ (eq. 2.15) from data, using a stochastic model $\Upsilon_{\theta} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \rightarrow [0, 1]$, parameterized

by $\theta \in \Theta$. Namely, by:

$$\theta^* = \arg \min_{\theta \in \Theta} \mathcal{L}_{\text{sc}}(\Upsilon_{\theta}) = \arg \min_{\theta \in \Theta} \mathbb{E}_{\vec{o}_t \sim \mathbb{P}_{\vec{o}}} \left[\mathcal{D} \left(\mathbb{P}_{\vec{s}|\vec{o}}^{\Delta t}(\cdot | \vec{o}_t) \parallel \Upsilon_{\theta}(\cdot | \vec{o}_t) \right) \right], \quad (2.17)$$

Where similarly as in eq. 2.16, \mathcal{D} denotes an (intractable) divergence measure between probability distributions on $(\mathcal{S}, \Sigma_{\mathcal{S}})$, and $\mathbb{P}_{\vec{o}}$ is a distribution over observation histories, derived from $\mathbb{P}_{\vec{o}, \mathbf{a}}$ by marginalization (Theodorou et al., 2010; Thrun et al., 2005).

2.3 Group and representation theory for mortals

This section provides a concise overview of the fundamental concepts in group and representation theory, which underpin the definitions of symmetries for dynamical systems and random variables throughout this work. If you are uninitiated in group theory, welcome, don't panic, readers are encouraged to skim or skip this section on a first reading, returning to it as needed when specific concepts are referenced in the main text. For comprehensive background on these topics in finite-dimensional vector spaces, see Steinberg (2012); Weiler et al. (2023); for infinite-dimensional spaces, as used in chapter 9, consult Knapp (1986).

Group theory provides the mathematical framework for describing symmetries in a wide range of contexts, including physics, chemistry, and machine learning. The central object of interest is a group: a set of elements closed under composition and inversion.

Definition 2.1 (Group). *A group is a set \mathbb{G} , endowed with a binary composition operator defined as:*

$$\begin{aligned} (\circ) : \quad \mathbb{G} \times \mathbb{G} &\longrightarrow \mathbb{G} \\ (g_1, g_2) &\longrightarrow g_1 \circ g_2, \end{aligned} \quad (2.18a)$$

such that the following axioms hold:

$$\text{Associativity: } (g_1 \circ g_2) \circ g_3 = g_1 \circ (g_2 \circ g_3), \quad \forall g_1, g_2, g_3 \in \mathbb{G}, \quad (2.18b)$$

$$\text{Identity: } \exists e \in \mathbb{G} \text{ such that } e \circ g = g = g \circ e, \quad \forall g \in \mathbb{G}, \quad (2.18c)$$

$$\text{Inverses: } \forall g \in \mathbb{G}, \exists g^{-1} \in \mathbb{G} \text{ such that } g \circ g^{-1} = e = g^{-1} \circ g. \quad (2.18d)$$

Remark: We use e to denote the identity (or trivial) element of a group \mathbb{G} , and $g \in \mathbb{G}$ to denote generic group elements. Furthermore, the cardinality of a group \mathbb{G} is denoted by $|\mathbb{G}|$ and referred to as the group's order.

In this work, we are primarily interested in *symmetry groups*, i.e., groups of symmetry transformations of an object, typically a set \mathcal{X} . Each transformation is a bijection on the set, termed a *symmetry* because it preserves a fundamental property of the set. To describe how a group acts on a given set, we introduce the concept of a group action: a map that associates each group element with a transformation of the set. We define two types of group actions, left and right, which differ only in the order of composition between group elements and transformations.

Definition 2.2 (Left group action). *Let \mathcal{X} be a set endowed with symmetry group \mathbb{G} . The (left) group action of the group \mathbb{G} on the set \mathcal{X} is a map:*

$$\begin{aligned} (\triangleright_{\mathcal{X}}) : \mathbb{G} \times \mathcal{X} &\longrightarrow \mathcal{X} \\ (g, \mathbf{x}) &\longrightarrow g \triangleright_{\mathcal{X}} \mathbf{x} \end{aligned} \quad (2.19a)$$

that is compatible with the group composition and identity element $e \in \mathbb{G}$, in the sense that:

$$\text{Identity: } e \triangleright_{\mathcal{X}} \mathbf{x} = \mathbf{x}, \quad \forall \mathbf{x} \in \mathcal{X} \quad (2.19b)$$

$$\text{Associativity: } (g_1 \circ g_2) \triangleright_{\mathcal{X}} \mathbf{x} = g_1 \triangleright_{\mathcal{X}} (g_2 \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall g_1, g_2 \in \mathbb{G}, \forall \mathbf{x} \in \mathcal{X}. \quad (2.19c)$$

Definition 2.3 (Right group action). *Let \mathcal{X} be a set endowed with symmetry group \mathbb{G} . The right group action of the group \mathbb{G} on the set \mathcal{X} is a map:*

$$\begin{aligned} (\triangleleft_{\mathcal{X}}) : \mathcal{X} \times \mathbb{G} &\longrightarrow \mathcal{X} \\ (\mathbf{x}, g) &\longrightarrow \mathbf{x} \triangleleft_{\mathcal{X}} g \end{aligned} \quad (2.20a)$$

that is compatible with the group composition and identity element $e \in \mathbb{G}$, in the sense that:

$$\text{Identity: } \mathbf{x} \triangleleft_{\mathcal{X}} e = \mathbf{x}, \quad \forall \mathbf{x} \in \mathcal{X} \quad (2.20b)$$

$$\text{Associativity: } \mathbf{x} \triangleleft_{\mathcal{X}} (g_1 \circ g_2) = (\mathbf{x} \triangleleft_{\mathcal{X}} g_1) \triangleleft_{\mathcal{X}} g_2, \quad \forall g_1, g_2 \in \mathbb{G}, \forall \mathbf{x} \in \mathcal{X}. \quad (2.20c)$$

These sets of bijections describe structural properties of the set \mathcal{X} . To study these properties, we will frequently refer to the group of symmetry related elements of to a given element $\mathbf{x} \in \mathcal{X}$ as its *group orbit* of \mathbf{x} :

Definition 2.4 (Group orbit). *Let \mathbf{x} be an element of the set \mathcal{X} endowed with symmetry group \mathbb{G} . The group orbit of \mathbf{x} is the the set of all symmetry related set elements, denoted by:*

$$\mathbb{G}\mathbf{x} := \{g \triangleright_{\mathcal{X}} \mathbf{x} \mid g \in \mathbb{G}\} \equiv \{\mathbf{x} \triangleleft_{\mathcal{X}} g \mid g \in \mathbb{G}\}.$$

Familiar groups in robotics Because robotics studies the dynamics of mechanical systems in three-dimensional Euclidean space, roboticists are familiar with Lie groups arising from Euclidean geometry. These include the rotation groups in two and three dimensions, \mathbb{SO}_2 and \mathbb{SO}_3 , respectively, and the groups of rigid translations and rotations in two and three dimensions, \mathbb{SE}_2 and \mathbb{SE}_3 , respectively. These Lie groups are also termed matrix Lie groups: their elements are represented in matrix form—i.e., as rotation and homogeneous transformation matrices—and the left and right group actions (definitions 2.2 and 2.3) reduce to left and right matrix multiplication of points in 2- and 3-dimensional space, respectively.

In robotics, these Lie groups are often studied from a Euclidean-geometric or differential-geometric perspective rather than from a group-theoretic one, largely because most of classical robotics does not use core results from group and representation theory. In this thesis, the underlying symmetry group of interest is the Euclidean group \mathbb{E}_3 —composed of rotations, reflections, and translations in 3-dimensional space—and its subgroups. However, we will rely heavily on an abstract group-theoretic analysis of this group, because we use several group- and representation-theoretic results and abstractions to characterize the implications of symmetries for robot dynamics and control. For instance, we will define several group actions of \mathbb{E}_3 on diverse sets and vector spaces whose

dimensionality differs from 3, including high-dimensional and even infinite-dimensional vector spaces. This will require us to define generalized matrix representations of \mathbb{E}_3 and its subgroups on arbitrary vector spaces. Furthermore, we will leverage the structure of these high-dimensional matrix representations to derive theoretical and practical results on the implications of symmetries for robot dynamics and control that would be difficult to obtain without a group-theoretic perspective.

For instance, in this work we frequently leverage symmetry transformations of three types of sets:

- ✱ **The state of a mechanical/robotic system:** The state space of a mechanical system, defined by the space of position and velocity generalized coordinates $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$ (see [section 2.1](#)), is said to have a symmetry group if each $g \in \mathbb{G}$ preserves the energy of the system (see [definition 3.1](#)).
- ✱ **A manifold:** In this work, a manifold \mathcal{M} is said to have a symmetry group \mathbb{G} if every $g \in \mathbb{G}$ preserves the scalar curvature of the manifold; that is, the trace of the manifold's curvature metric is the same at points related by the group action (see [fig. 3.4](#)).
- ✱ **A random variable:** A random variable \mathbf{y} is said to have a symmetry group \mathbb{G} if every $g \in \mathbb{G}$ preserves the marginal probability of an event; that is, the probability of an event is the same for events related by the group action (see [definition 2.16](#)).

Finite symmetry groups A second divergence from the classical study of Euclidean isometries is that we focus on discrete, often finite, symmetry groups arising from discretizations and finite subgroups of the Euclidean group \mathbb{E}_3 . As discussed in [chapter 4](#), morphological symmetry groups are often finite. Below, we list and depict graphically in [fig. 2.1](#) some of the finite symmetry groups that will recur in this work:

- ✱ **The reflection group** $\mathbb{C}_2 = \{e, g_r \mid g_r^2 = e\}$: This is the simplest non-trivial symmetry group, composed of a single non-trivial reflection symmetry g_r . It plays an important role in this thesis because it characterizes the dynamical symmetries of many robotic and biological systems with bilateral reflection symmetry in their morphology (see [fig. 1.1](#)).
- ✱ **Cyclic groups** $\mathbb{C}_n = \{e, g, g^2, \dots, g^{n-1} \mid g^n = e\}$: Cyclic groups of order n arise from discretizing the Lie group of 2-dimensional rotations, \mathbb{SO}_2 , into n equally spaced rotations. For instance, \mathbb{C}_4 is the symmetry group of a square under rotations, and \mathbb{C}_6 is the symmetry group of a regular hexagon under rotations.
- ✱ **Dihedral groups** $\mathbb{D}_n \simeq \mathbb{C}_n \times \mathbb{C}_2$: These groups arise from discretizing the group of rotations and reflections in 2 dimensions, \mathbb{O}_2 , into n equally spaced rotations and a single reflection. For instance, \mathbb{D}_4 is the full symmetry group of a square. A related finite group used to describe the dynamical symmetries of some quadrupedal robots is the Klein-four group \mathbb{K}_4 , generated by two commuting reflections and their composition (see [fig. 4.1](#)).
- ✱ **Discretizations of the sphere:** A wide variety of finite symmetry groups arise from discretizations and finite subgroups of the 3-dimensional rotation and orthogonal groups, \mathbb{SO}_3 and \mathbb{O}_3 . Familiar examples used in this thesis include the icosahedral group \mathbb{I}_h and the octahedral group, which are the symmetry groups of the regular icosahedron and octahedron, respectively.

Representation theory

In practice, to numerically represent the action of a symmetry group on a robot's configuration manifold, or on the value space of a random variable of interest, we choose a vector space and a basis




Figure 2.1 Cayley diagrams for representative finite symmetry groups used throughout this thesis. Cayley diagrams represent groups graphically, where nodes correspond to the individual elements of the group, and colored arrows denote right composition by the group generators, whose combinations result in the entire group. Algebraically, these groups are defined by their elements and constraints as $\mathbb{C}_n = \{e, g, \dots, g^{n-1} \mid g^n = e\}$, $\mathbb{D}_n = \{e, g, \dots, g^{n-1}, h, gh, \dots, g^{n-1}h \mid g^n = e, h^2 = e, h \circ g \circ h = g^{-1}\}$, $\mathbb{K}_4 \equiv \mathbb{D}_2$, and $\mathbb{K}_4 \times \mathbb{C}_2 = \{e, g_s, g_t, g_s g_t, g_f, g_s \circ g_f, g_t \circ g_f, g_s \circ g_t \circ g_f \mid g_s^2 = e, g_t^2 = e, g_f^2 = e, g_s g_t = g_t g_s, g_s g_f = g_f g_s, g_t g_f = g_f g_t\}$.

in which to express the relevant points. In this work, most group actions of interest are represented numerically using a linear group representation, which maps the symmetry group of an object to a group of invertible linear transformations of the vector space used to represent that object. Denoting by $\mathbb{GL}(\mathcal{X})$ the group of *all* invertible linear transformations on a vector space \mathcal{X} (i.e., the set of all invertible matrices), we can express a group action on \mathcal{X} via a linear group representation:

Definition 2.5 (Linear group representation). *Let $\mathcal{X} \subset \mathbb{R}^n$ be a vector space endowed with symmetry group \mathbb{G} . A linear representation of \mathbb{G} on \mathcal{X} is a map, denoted by $\rho_{\mathcal{X}}$, between symmetry transformation and invertible linear maps on \mathcal{X} (i.e., elements of the general linear group $\mathbb{GL}(\mathcal{X})$):*

$$\begin{aligned} \rho_{\mathcal{X}} : \mathbb{G} &\longrightarrow \mathbb{GL}(\mathcal{X}) \\ g &\longrightarrow \rho_{\mathcal{X}}(g) \in \mathbb{R}^{n \times n}, \end{aligned} \quad (2.21a)$$

such that the following properties hold:

$$\text{composition} : \rho_{\mathcal{X}}(g_1 \circ g_2) = \rho_{\mathcal{X}}(g_1) \rho_{\mathcal{X}}(g_2), \quad \forall g_1, g_2 \in \mathbb{G}, \quad (2.21b)$$

$$\text{inversion} : \rho_{\mathcal{X}}(g^{-1}) = \rho_{\mathcal{X}}(g)^{-1}, \quad \forall g \in \mathbb{G}. \quad (2.21c)$$

$$\text{identity} : \rho_{\mathcal{X}}(g \circ g^{-1}) = \rho_{\mathcal{X}}(e) = \mathbf{I}, \quad (2.21d)$$

Eqs. 2.21b to 2.21d show how the composition and inversion of symmetry transformations translate to matrix multiplication and inversion, respectively.

This enables us to express any linear group action $\triangleright_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$ via $\rho_{\mathcal{X}}$ as follows:

$$\begin{aligned} (\triangleright_{\mathcal{X}}) : \mathbb{G} \times \mathcal{X} &\longrightarrow \mathcal{X} \\ (g, \mathbf{x}) &\longrightarrow g \triangleright_{\mathcal{X}} \mathbf{x} := \rho_{\mathcal{X}}(g) \mathbf{x}. \end{aligned} \quad (2.21e)$$

Remark: We will default to using the group action notation $g \triangleright_{\mathcal{X}} \mathbf{x}$ rather than the representation notation $\rho_{\mathcal{X}}(g) \mathbf{x}$, unless the latter is needed for clarity or to avoid ambiguity about the basis. Furthermore, throughout this work, we assume group representations to be real, meaning that the representation maps to the real general linear group $\mathbb{GL}(\mathcal{X}, \mathbb{R})$ rather than to the complex general linear group $\mathbb{GL}(\mathcal{X}, \mathbb{C})$. This is a non-standard choice in representation theory, where most standard

results are stated for complex representations (Steinberg, 2012). However, it is standard in machine learning, where the real linear-algebra operations underlying group actions are more memory- and compute-efficient to implement.

Note that choosing a different basis set in which to represent points in \mathcal{X} will lead to $\triangleright_{\mathcal{X}}$ having different matrix form representations, that is a different linear group representation $\rho_{\mathcal{X}}$. However, the group action $\triangleright_{\mathcal{X}}$ itself remains invariant to the choice of basis, thus creating an equivalence relationship between group representations:

Definition 2.6 (Equivalent group representations). *Let \mathcal{X} be a vector space endowed with symmetry group \mathbb{G} , and let $\rho'_{\mathcal{X}}$ and $\rho_{\mathcal{X}}$ be two group representations of \mathbb{G} on \mathcal{X} . They are said to be equivalent, denoted by $\rho'_{\mathcal{X}} \sim \rho_{\mathcal{X}}$, if there exists an invertible change of basis $\mathbf{Q} \in \text{End}_{\mathbb{G}}(\mathcal{X})$ such that*

$$\rho'_{\mathcal{X}}(g) = \mathbf{Q}\rho_{\mathcal{X}}(g)\mathbf{Q}^{-1}, \quad \forall g \in \mathbb{G}. \quad (2.22)$$

Equivalent representations arise when the same group action $(\triangleright_{\mathcal{X}}) : \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$ is expressed in different coordinate frames or bases. For instance, let $\mathbb{A}_{\mathcal{X}}$ and $\mathbb{B}_{\mathcal{X}}$ be two bases for $\mathcal{X} = \text{span}(\mathbb{A}_{\mathcal{X}}) = \text{span}(\mathbb{B}_{\mathcal{X}})$, and let $\mathbf{Q}_{\mathbb{A}}^{\mathbb{B}} : \mathcal{X} \rightarrow \mathcal{X}$ denote the change of basis from $\mathbb{A}_{\mathcal{X}}$ to $\mathbb{B}_{\mathcal{X}}$, so that $\mathbf{x}^{\mathbb{B}} = \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}}\mathbf{x}^{\mathbb{A}}$ for all $\mathbf{x}^{\mathbb{A}} \in \mathcal{X}$. Then the group action admits equivalent representations, $\rho_{\mathcal{X}}^{\mathbb{A}} \sim \rho_{\mathcal{X}}^{\mathbb{B}}$, since

$$\begin{aligned} g \triangleright_{\mathcal{X}} \mathbf{x}^{\mathbb{B}} &:= \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}}(g \triangleright_{\mathcal{X}} \mathbf{x}^{\mathbb{A}}), \quad \forall g \in \mathbb{G}, \\ \rho_{\mathcal{X}}^{\mathbb{B}}(g)\mathbf{x}^{\mathbb{B}} &= \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}}(\rho_{\mathcal{X}}^{\mathbb{A}}(g)\mathbf{x}^{\mathbb{A}}) = \left(\mathbf{Q}_{\mathbb{A}}^{\mathbb{B}}\rho_{\mathcal{X}}^{\mathbb{A}}(g)\mathbf{Q}_{\mathbb{A}}^{\mathbb{B}-1}\right)\mathbf{x}^{\mathbb{B}}, \\ \rho_{\mathcal{X}}^{\mathbb{B}}(g) &= \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}}\rho_{\mathcal{X}}^{\mathbb{A}}(g)\mathbf{Q}_{\mathbb{A}}^{\mathbb{B}-1}. \end{aligned} \quad (2.23)$$

We will be frequently studying linear maps between symmetric vector spaces that preserve some or all of the group structure of the spaces. These maps are known as homomorphisms and isomorphisms, and they are defined as follows:

Definition 2.7 (Homomorphism and Isomorphism). *Let \mathcal{X} and \mathcal{Y} be two vector spaces endowed with the same symmetry group \mathbb{G} , with the respective group actions $\triangleright_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \mapsto \mathcal{X}$ and $\triangleright_{\mathcal{Y}} : \mathbb{G} \times \mathcal{Y} \mapsto \mathcal{Y}$. The spaces are said to be \mathbb{G} -homomorphic if there exists a linear map $\mathbf{A} : \mathcal{X} \mapsto \mathcal{Y}$ that commutes with the group action, such that $g \triangleright_{\mathcal{Y}}(\mathbf{A}\mathbf{x}) = \mathbf{A}(g \triangleright_{\mathcal{X}}\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$. They are said to be \mathbb{G} -isomorphic if the linear map is invertible. Graphically, \mathcal{X} and \mathcal{Y} are \mathbb{G} -homomorphic or \mathbb{G} -isomorphic if the following diagrams commute:*

$$\begin{array}{ccc} \mathcal{X} \xrightarrow{\triangleright_{\mathcal{X}}} \mathcal{X} & \mathbf{A} \in \text{Hom}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y}) & \text{or} & \mathcal{X} \xrightarrow{\triangleright_{\mathcal{X}}} \mathcal{X} & \mathbf{A} \in \text{Iso}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y}). & (2.24) \\ \downarrow \mathbf{A} & & & \downarrow \mathbf{A}^{-1} & & \\ \mathcal{Y} \xrightarrow{\triangleright_{\mathcal{Y}}} \mathcal{Y} & & & \mathcal{Y} \xrightarrow{\triangleright_{\mathcal{Y}}} \mathcal{Y} & & \\ \underbrace{\hspace{10em}} & & & \underbrace{\hspace{10em}} & & \\ \text{Homomorphism} & & & \text{Isomorphism} & & \end{array}$$

Here, $\text{Hom}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y})$ denotes the space of \mathbb{G} -equivariant linear maps between \mathcal{X} and \mathcal{Y} , and $\text{Iso}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y})$ denotes the space of \mathbb{G} -equivariant invertible linear maps between \mathcal{X} and \mathcal{Y} .

Whenever we refer to a linear map from a symmetric vector space to itself that preserves the group structure, we call it an *endomorphism*. This is a special case of homomorphisms and isomorphisms where the domain and codomain are the same space.

Definition 2.8 (Endomorphism). *Let \mathcal{X} be a vector space endowed with symmetry group \mathbb{G} , with the group action $\triangleright_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \mapsto \mathcal{X}$. A linear map $\mathbf{A} : \mathcal{X} \mapsto \mathcal{X}$ is said to be an endomorphism if it commutes with the group action, such that:*

$$\rho_{\mathcal{X}}(g)\mathbf{A} = \mathbf{A}\rho_{\mathcal{X}}(g), \quad \forall g \in \mathbb{G} \quad \iff \quad \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \downarrow \mathbf{A} & & \downarrow \mathbf{A} \\ \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \end{array}$$

We will denote the space of all endomorphisms of \mathcal{X} as $\text{End}_{\mathbb{G}}(\mathcal{X})$, such that any $\mathbf{A} \in \text{End}_{\mathbb{G}}(\mathcal{X})$ satisfies the above commutation property.

Structure of \mathbb{G} -symmetric vector spaces Symmetric vector spaces can be decomposed into smaller subspaces with fewer symmetries, analogous to how a symmetry group can be broken down into products and direct products of smaller groups. We will leverage this structure and decomposition in several results throughout the text, so it is important to introduce the relevant concepts and definitions here.

Definition 2.9 (\mathbb{G} -stable and irreducible subspaces). *Let \mathcal{X} be a vector space endowed with a group action $(\triangleright_{\mathcal{X}})$ of the symmetry group \mathbb{G} . A subspace $\mathcal{X}' \subseteq \mathcal{X}$ is said to be \mathbb{G} -stable if the action of any group element on any vector in the subspace remains within the subspace, that is,*

$$g \triangleright_{\mathcal{X}} \mathbf{x} \in \mathcal{X}', \quad \forall \mathbf{x} \in \mathcal{X}' \subseteq \mathcal{X}, \forall g \in \mathbb{G}.$$

If the only \mathbb{G} -stable subspaces of \mathcal{X} are $\{\mathbf{0}\}$ and \mathcal{X} itself, then \mathcal{X} is a irreducible \mathbb{G} -stable space. We will denote irreducible \mathbb{G} -stable spaces with an over bar, e.g., $\bar{\mathcal{V}}$.

Algebraically, the presence of a \mathbb{G} -stable subspace implies that the group representation on the entire space can be decomposed block-diagonally into group representations acting on each \mathbb{G} -stable subspace.

Definition 2.10 (Decomposable representation). *Let \mathcal{X} be a vector space with a group action $(\triangleright_{\mathcal{X}})$ defined by the representation $\rho_{\mathcal{X}}$ in a chosen basis $\mathbb{A}_{\mathcal{X}}$. The representation is decomposable if it is equivalent to a direct sum of two lower-dimensional representations, $\rho_{\mathcal{X}} \sim \rho_{\mathcal{X}_1} \oplus \rho_{\mathcal{X}_2}$, where \mathcal{X}_1 and \mathcal{X}_2 are \mathbb{G} -stable subspaces of \mathcal{X} . Equivalently, there exists a change of basis $\mathbf{Q}_{\mathbb{A}}^{\mathbb{B}} : \mathcal{X} \rightarrow \mathcal{X}$ such that*

$$\rho_{\mathcal{X}}^{\mathbb{B}} = \begin{bmatrix} \rho_{\mathcal{X}_1} & \mathbf{0} \\ \mathbf{0} & \rho_{\mathcal{X}_2} \end{bmatrix} = \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}} \rho_{\mathcal{X}} \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}-1}, \quad \text{and} \quad g \triangleright_{\mathcal{X}} \mathbf{x}^{\mathbb{B}} := \rho_{\mathcal{X}}^{\mathbb{B}}(g) \mathbf{x}^{\mathbb{B}} = \begin{bmatrix} \rho_{\mathcal{X}_1}(g) \mathbf{x}_1^{\mathbb{B}} \\ \rho_{\mathcal{X}_2}(g) \mathbf{x}_2^{\mathbb{B}} \end{bmatrix}, \quad \text{where} \quad \mathbf{Q}_{\mathbb{A}}^{\mathbb{B}} \mathbf{x} = \begin{bmatrix} \mathbf{x}_1^{\mathbb{B}} \in \mathcal{X}_1 \\ \mathbf{x}_2^{\mathbb{B}} \in \mathcal{X}_2 \end{bmatrix}$$

Hence, the representation's decomposition $\rho_{\mathcal{X}} \sim \rho_{\mathcal{X}_1} \oplus \rho_{\mathcal{X}_2}$ corresponds to **decomposing the vector space** into \mathbb{G} -stable subspaces, $\mathcal{X} = \mathcal{X}_1 \oplus \mathcal{X}_2$.

When iteratively applying the decomposition process, we eventually reach representations that cannot be further decomposed. These are known as irreducible representations, or *irreps*, and they serve as the fundamental building blocks for all representations of a compact symmetry group \mathbb{G} . From a vector space perspective, the irreducible \mathbb{G} -stable subspaces (definition 2.9) associated with these irreps are the elementary subspaces that comprise any symmetric vector space, analogous to how one-dimensional subspaces are the fundamental components of standard vector spaces.

Definition 2.11 (Irreducible representation). *Let \mathcal{X} be a vector space endowed with a group action $(\triangleright_{\mathcal{X}})$ of a symmetry group \mathbb{G} . A representation $\rho_{\mathcal{X}}$ of \mathbb{G} on \mathcal{X} is said to be irreducible if it cannot be*

decomposed into smaller representations acting on proper \mathbb{G} -stable subspaces (definition 2.9). That is, the only \mathbb{G} -stable subspaces $\mathcal{X}' \subseteq \mathcal{X}$ are $\mathcal{X}' = \{\mathbf{0}\}$ and $\mathcal{X}' = \mathcal{X}$ itself.

To differentiate irreps from decomposable representations we will denote the formers and their associated irreducible \mathbb{G} -stable spaces with an over bar: $\bar{\rho}_{\bar{\mathcal{V}}} : \mathbb{G} \rightarrow \mathbb{GL}(\bar{\mathcal{V}})$.

Crucially, any compact symmetry group \mathbb{G} has a unique set of countably many irreps, denoted by $\{\bar{\rho}_k\}_{k \in [1, n_{\text{iso}}]}$, where k denotes the irrep type and $n_{\text{iso}} \leq |\mathbb{G}|$ denotes the number of unique irreps of \mathbb{G} . A fundamental property of these irreps is that any two non-equivalent irreps act on vector spaces that are mutually *orthogonal*. This implies that whenever we decompose symmetric vector spaces into their irreducible subspaces we are inherently decomposing the space into orthogonal \mathbb{G} -stable subspaces, which will greatly simplify numerical and theoretical analyses. Formally, these orthogonality relations are a consequence of Schur's lemma, which we state below in its original form for the case of complex irreps and discuss its adaptation to real irreps.

Lemma 2.1 (Schur's Lemma for unitary (complex) representations (Knapp, 1986, Prop 1.5)). *Consider two complex Hilbert spaces, \mathcal{H} and \mathcal{H}' , endowed with the (complex) irreducible unitary representations $\bar{\phi} : \mathbb{G} \mapsto \mathbb{U}(\mathcal{H})$ and $\bar{\phi}' : \mathbb{G} \mapsto \mathbb{U}(\mathcal{H}')$, respectively. Let $\mathbf{T} : \mathcal{H} \mapsto \mathcal{H}'$ be a linear map commuting with the group actions, such that $\mathbf{T} \in \text{Hom}_{\mathbb{G}}(\mathcal{H}, \mathcal{H}')$. Then, if the irreducible representations are not equivalent, i.e., $\bar{\phi} \not\sim \bar{\phi}'$, then \mathbf{T} is the trivial (or zero) map. Conversely, if $\bar{\phi} \sim \bar{\phi}'$, then \mathbf{T} is a constant multiple of an isomorphism. Denoting \mathbf{I} as the identity operator, this can be expressed as:*

$$\bar{\phi} \not\sim \bar{\phi}' \iff \mathbf{0}_{\mathcal{H}'} = \mathbf{T}h \mid \forall h \in \mathcal{H} \quad (2.25a)$$

$$\bar{\phi} \sim \bar{\phi}' \iff \mathbf{T} = \alpha \mathbf{U}, \alpha \in \mathbb{C}, \mathbf{U} \cdot \mathbf{U}^H = \mathbf{I} \quad (2.25b)$$

$$\bar{\phi} = \bar{\phi}' \iff \mathbf{T} = \alpha \mathbf{I}, \alpha \in \mathbb{C} \quad (2.25c)$$

The most common interpretation of Schur's lemma is that whenever the irreps are equivalent, $\bar{\phi} \sim \bar{\phi}'$, their associated spaces are isomorphic, $\mathcal{H} \sim \mathcal{H}'$ and \mathbf{T} is an element of the endomorphism space $\text{End}_{\mathbb{G}}^{\mathbb{C}}(\bar{\mathcal{H}})$, with $\bar{\mathcal{H}} \sim \mathcal{H} \sim \mathcal{H}'$ (see definition 2.8). Consequently, (2.25b) implies that the endomorphism space is one-dimensional, i.e., $|\text{End}_{\mathbb{G}}^{\mathbb{C}}(\bar{\mathcal{H}})| = 1$, with $\alpha \in \mathbb{C}$ denoting the only degree of freedom, and (2.25c) denotes the scenario in which the basis sets for the two spaces are identical.

However, this result holds only for *complex* irreducible representations and requires adaptation for the case of real irreducible representations—given that computational linear algebra is mostly based on the real numbers. The main difference stems from the fact that if $\bar{\rho} : \mathbb{G} \rightarrow \mathbb{GL}(\bar{\mathcal{V}})$ is a real irrep, then the space of (real) endomorphisms, $\text{End}_{\mathbb{G}}(\bar{\mathcal{V}})$, is no longer one-dimensional, but rather it can be 1, 2, or 4 dimensional. Here, an *algebra* denotes a vector space equipped with a bilinear multiplication operation. The relevant endomorphism algebras are isomorphic to the real algebra ($\mathbb{R} = \text{span}\{1\}$), complex algebra ($\mathbb{C} = \text{span}\{1, i \mid i^2 = -1\}$), or quaternionic algebra ($\mathbb{H} = \text{span}\{1, i, j, k \mid i^2 = j^2 = k^2 = -1\}$), respectively. Denoting by $\Psi : \mathbb{K} \rightarrow \text{End}_{\mathbb{G}}(\bar{\mathcal{V}})$ the isomorphism of basis elements of $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$ with the basis elements of $\text{End}_{\mathbb{G}}(\bar{\mathcal{V}})$, we can summarize the basis




Figure 2.2 Example of an endomorphism acting on a \mathbb{C}_3 -stable irreducible 2D space. The irreducible representation is of complex type, with endomorphism space $\text{End}_{\mathbb{C}_3}(\mathbb{R}^2) \sim \mathbb{C} = \text{span}\{1, i\}$, comprising all transformations that uniformly scale and rotate/reflect the plane.

sets of the three cases as follows (see Cesa et al. (2021, Appendix C) for details):

$$\begin{aligned}
 \mathbb{R} &\sim \text{End}_{\mathbb{G}}(\mathcal{V}) = \text{span}\{\Psi(1) = \text{Id}_{|\bar{\rho}|}\} \\
 \mathbb{C} &\sim \text{End}_{\mathbb{G}}(\mathcal{V}) = \text{span}\left\{\Psi(1) = \begin{bmatrix} \text{Id}_n & \mathbf{0} \\ \mathbf{0} & \text{Id}_n \end{bmatrix}, \Psi(i) = \begin{bmatrix} \mathbf{0} & -\text{Id}_n \\ \text{Id}_n & \mathbf{0} \end{bmatrix}\right\} \\
 \mathbb{H} &\sim \text{End}_{\mathbb{G}}(\mathcal{V}) = \text{span}\left\{ \begin{array}{l} \Psi(1) = \begin{bmatrix} \text{Id}_n & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \text{Id}_n & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \text{Id}_n & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \text{Id}_n \end{bmatrix}, \Psi(i) = \begin{bmatrix} \mathbf{0} & \mathbf{0} & -\text{Id}_n & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & -\text{Id}_n \\ \text{Id}_n & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \text{Id}_n & \mathbf{0} & \mathbf{0} \end{bmatrix} \\ \Psi(j) = \begin{bmatrix} \mathbf{0} & -\text{Id}_n & \mathbf{0} & \mathbf{0} \\ \text{Id}_n & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \text{Id}_n \\ \mathbf{0} & \mathbf{0} & -\text{Id}_n & \mathbf{0} \end{bmatrix}, \Psi(k) = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & -\text{Id}_n \\ \mathbf{0} & \mathbf{0} & \text{Id}_n & \mathbf{0} \\ \mathbf{0} & -\text{Id}_n & \mathbf{0} & \mathbf{0} \\ \text{Id}_n & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \end{array} \right\} \quad (2.26)
 \end{aligned}$$

While this result might appear complex, its interpretation is straightforward: given a \mathbb{G} -stable irreducible space $\bar{\mathcal{V}}$, the space of linear maps from the space to itself that preserve the symmetry structure consists of *linear transformations that scale all dimensions of $\bar{\mathcal{V}}$ uniformly, and possibly rotate or reflect the space*. Algebraically, this means that any element of the algebra has a unique singular space, with a single singular value determined by the element's coefficients in the basis of (2.26). We summarize this result in the following proposition.

Proposition 2.1 (A real endomorphism has a single singular space). *Let \mathbb{G} be a compact symmetry group, $(\bar{\rho}, \bar{\mathcal{V}})$ be an irreducible representation and its associated \mathbb{G} -stable space, and let $\text{End}_{\mathbb{G}}(\bar{\mathcal{V}})$ denote the space endomorphism algebra. Then every $\mathbf{A} \in \text{End}_{\mathbb{G}}(\bar{\mathcal{V}})$ admits an *singular value decomposition (SVD)**

$$\mathbf{A} = \mathbf{U} \gamma \text{Id}_d \mathbf{V}^\top,$$

where $\gamma \in \mathbb{R}_{\geq 0}$ is the single singular value, repeated with multiplicity $d = |\bar{\rho}| = |\mathcal{V}_k|$. The right singular basis \mathbf{V} can, without loss of generality, be taken as the canonical orthonormal basis of $\bar{\mathcal{V}}$, while the left singular basis is then $\mathbf{U} = \gamma^{-1} \mathbf{A} \mathbf{V}$, which is an orthogonal rotation/reflection of \mathbf{V} .

(\mathbb{R}) **Real case.** Any $\mathbf{A} \in \text{End}_{\mathbb{G}}(\bar{\mathcal{V}})$ is of the form

$$\mathbf{A} = a \Psi(1) = a \text{Id}_d, \quad a \in \mathbb{R}.$$

Hence

$$\mathbf{A}^\top \mathbf{A} = a^2 \text{Id}_d, \quad \sigma(\mathbf{A}) = \{|a|\}^{\times d}.$$

(\mathbb{C}) **Complex case.** Every element can be written as

$$\mathbf{A} = a \Psi(1) + b \Psi(i) = \begin{bmatrix} a \text{Id}_n & -b \text{Id}_n \\ b \text{Id}_n & a \text{Id}_n \end{bmatrix}, \quad a, b \in \mathbb{R}.$$

Using $\Psi(i)^\top = -\Psi(i)$ and $\Psi(i)^\top \Psi(i) = \text{Id}$,

$$\mathbf{A}^\top \mathbf{A} = (a^2 + b^2) \text{Id}_d, \quad \sigma(\mathbf{A}) = \{\sqrt{a^2 + b^2}\}^{\times d}.$$

(\mathbb{H}) **Quaternionic case.** Each element admits the expansion

$$\mathbf{A} = a \Psi(1) + b \Psi(i) + c \Psi(j) + d \Psi(k), \quad a, b, c, d \in \mathbb{R},$$

where $\Psi(i), \Psi(j), \Psi(k)$ are the quaternionic structure matrices from (2.26), satisfying $\Psi(\alpha)^\top = -\Psi(\alpha)$, $\Psi(\alpha)^2 = -\text{Id}$, and $\Psi(\alpha)^\top \Psi(\alpha) = \text{Id}$, with the usual anti-commutation rules. Consequently,

$$\mathbf{A}^\top \mathbf{A} = (a^2 + b^2 + c^2 + d^2) \text{Id}_d, \quad \sigma(\mathbf{A}) = \{\sqrt{a^2 + b^2 + c^2 + d^2}\}^{\times d}.$$

As an intuitive low-dimensional example, consider the case of a 2D rotational irrep of the cyclic group \mathbb{C}_3 . The dimension of the irreducible \mathbb{G} -stable subspace is $|\bar{\mathcal{V}}| = 2$, and the irreducible representation is of complex type, $\bar{\rho} : \mathbb{C}_3 \rightarrow \mathbb{GL}(\bar{\mathcal{V}})$, with $\text{End}_{\mathbb{G}}(\bar{\mathcal{V}}) \sim \mathbb{C} = \text{span}\{1, i\}$ denoting the space of all rotations, reflections, and uniform scalings of the plane (see fig. 2.2).

2.3.1 Decomposition of symmetric vector spaces

We now have all the necessary tools to decompose symmetric vector spaces into their fundamental building blocks: irreducible \mathbb{G} -stable subspaces (definition 2.9). By Maschke's theorem (Knapp, 1986), we have that irreducible representations are the fundamental building blocks of the representations of a compact symmetry group \mathbb{G} , given that any group representation $\rho_{\mathcal{X}} : \mathbb{G} \rightarrow \mathbb{GL}(\mathcal{X})$ can be decomposed into a direct sum of irreducible representations, $\rho_{\mathcal{X}} \sim \bigoplus_{i=1}^n \rho_{\mathcal{X}_i}$, where each $\rho_{\mathcal{X}_i}$ is isomorphic to one of the group's $n_{\text{iso}} \leq |\mathbb{G}|$ irreducible representations (definitions 2.10 and 2.11).

This decomposition will play a crucial role in facilitating numerical and theoretical analysis of operations on symmetric vector spaces. Therefore, we will frequently choose a convenient basis of the symmetric vector space which readily exposes this decomposition, termed isotypic basis. For the sake of generality we consider below the more general case of (finite and infinite dimensional) separable Hilbert spaces, which will enable us to extend these results to function spaces.

Definition 2.12 (Isotypic Basis). Let $\rho_{\mathcal{H}} : \mathbb{G} \rightarrow \mathbb{U}(\mathcal{H})$ be a unitary group representation of a compact group \mathbb{G} on a separable Hilbert Space \mathcal{H} . The representation is said to be defined in an isotypic basis if it is defined by a direct sum of irreducible representations grouped by their type, that is, if:

$$\rho_{\mathcal{H}} = \bigoplus_{k=1}^{n_{\text{iso}}} \bigoplus_{p=1}^{m_k} \bar{\rho}_k, \quad (2.27)$$

where $\{\bar{\rho}_k : \mathbb{G} \rightarrow \mathbb{U}(\bar{\mathcal{H}}_k)\}_{k=1}^{n_{\text{iso}}}$ are the $n_{\text{iso}} \leq |\mathbb{G}|$ irreducible representations of \mathbb{G} , and $m_k \leq \infty$ is the multiplicity (i.e., number of copies) of the irrep type k in the representation $\rho_{\mathcal{H}}$.

Remark 2.1. Note that multiple isotypic bases exist for a given representation $\rho_{\mathcal{H}}$, as both the irrep ordering and each irrep's multiplicity ordering can be arbitrarily permuted.

The utility of an isotypic basis stems from Schur’s orthogonality relations (lemma 2.1), which ensure that any symmetric vector space decomposes into at most n_{iso} orthogonal subspaces.

Theorem 2.1 (Isotypic decomposition of symmetric Hilbert spaces (Knapp, 1986)). *Let \mathbb{G} be a compact group and \mathcal{H} a separable Hilbert space with a unitary group representation $\rho_{\mathcal{H}} : \mathbb{G} \rightarrow \mathbb{U}(\mathcal{H})$. Then we can identify $n_{\text{iso}} \leq |\mathbb{G}|$ irreducible representations $\bar{\rho}_k : \mathbb{G} \rightarrow \mathbb{U}(\bar{\mathcal{H}}_k)$ that allow us to decompose \mathcal{H} into a sum of orthogonal subspaces, denoted isotypic subspaces: $\mathcal{H} = \bigoplus_{1 \leq k \leq n_{\text{iso}}} \mathcal{H}^{(k)}$ where each $\mathcal{H}^{(k)} = \bigoplus_{p=1}^{m_k} \mathcal{H}_p^{(k)}$ is the sum of at most $m_k \leq \infty$ countably many subspaces isometrically isomorphic to $\bar{\mathcal{H}}_k$.*

Remark: In practice, the isotypic decomposition of any symmetric finite dimensional symmetric Hilbert space, is found via a unitary/orthogonal change of basis, which can be computed numerically using Dixon’s reduction method for finite-dimensional unitary (complex) representations (Dixon, 1970), with minor additional logic for the orthogonal (real) case. This method is implemented in `symm_learning` in the `cplx_isotypic_decomposition` function.

Disentangled representations The concept of isotypic decomposition is intricately linked to the idea of disentangled representations, introduced by Higgins et al. (2018) in the representation learning literature, restated below for completeness.

Definition 2.13 (Disentangled representation (Higgins et al. (2018))). *A vector representation is called a disentangled representation with respect to a particular decomposition of a symmetry group into subgroups, if it decomposes into independent subspaces, where each subspace is affected by the action of a single subgroup, and the actions of all other subgroups leave the subspace unaffected.*

Note that the independent subspaces of definition 2.13 refer to the orthogonal isotypic subspaces $\{\mathcal{H}^{(k)}\}_{k=1}^{n_{\text{iso}}}$, each of which is acted upon by a unique quotient group⁶ defined by the kernel of the irreducible action on that isotypic subspace:

$$\mathbb{G}^{(k)} = \mathbb{G} / \mathbb{N}_k, \quad \text{where} \quad \mathbb{N}_k := \ker(\bar{\rho}_k) = \{g \in \mathbb{G} \mid \bar{\rho}_k(g) = \text{Id}_{d_k}\}. \quad (2.28)$$

Where each $\mathbb{G}^{(k)}$ is a well-defined group of cosets generated by the normal subgroup \mathbb{N}_k . In practice, each $\mathbb{G}^{(k)}$ is isomorphic to the effective (matrix) group encoded by each irreducible representation $\bar{\rho}_k : \mathbb{G} \mapsto \mathbb{U}(\bar{\mathcal{H}}_k)$.

In practice we will leverage this decomposition in section 4.3 to decompose the configuration manifold of a robotic system into orthogonal submanifolds, and on section 9.2.1 to approximate \mathbb{G} -invariant conditional probability distributions using G -equivariant linear operators.

2.3.2 Maps between symmetric vector spaces

We will frequently study and use linear and non-linear maps between symmetric vector spaces. Our focus is on maps that preserve entirely or partially the group structure of the vector spaces. These types of maps can be classified as \mathbb{G} -equivariant, \mathbb{G} -invariant maps:

Definition 2.14 (\mathbb{G} -equivariant and \mathbb{G} -invariant maps). *Let \mathcal{X} and \mathcal{Y} be two vector spaces endowed with the same symmetry group \mathbb{G} , with the respective group actions $\triangleright_{\mathcal{X}}$ and $\triangleright_{\mathcal{Y}}$. A map $f : \mathcal{X} \mapsto \mathcal{Y}$*

⁶The original definition of disentangled representations refers to subgroups, but in general the quotient group $\mathbb{G}^{(k)}$ need not be a **subgroup** of \mathbb{G} .

is said to be \mathbb{G} -equivariant if it commutes with the group action, such that:

$$\begin{aligned} g \triangleright_{\mathcal{Y}} \mathbf{y} = g \triangleright_{\mathcal{Y}} f(\mathbf{x}) = f(g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, g \in \mathbb{G}. \\ \rho_{\mathcal{Y}}(g)f(\mathbf{x}) = f(\rho_{\mathcal{X}}(g)\mathbf{x}) \end{aligned} \iff \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \downarrow f & & \downarrow f \\ \mathcal{Y} & \xrightarrow{\triangleright_{\mathcal{Y}}} & \mathcal{Y} \end{array} \quad (2.29a)$$

A specific case of \mathbb{G} -equivariant maps are the \mathbb{G} -invariant ones, which are maps that commute with the group action and have trivial output group actions $\triangleright_{\mathcal{Y}}$ such that $\rho_{\mathcal{Y}}(g) = \mathbf{I}$ for all $g \in \mathbb{G}$. That is:

$$\begin{aligned} \mathbf{y} = g \triangleright_{\mathcal{Y}} f(\mathbf{x}) = f(g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, g \in \mathbb{G}. \\ \mathbf{y} = \rho_{\mathcal{Y}}(g)f(\mathbf{x}) = f(\rho_{\mathcal{X}}(g)\mathbf{x}) \end{aligned} \iff \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \searrow f & & \downarrow f \\ & & \mathcal{Y} \end{array} \quad (2.29b)$$

$\curvearrowright_{\triangleright_{\mathcal{Y}}}$

We will also frequently study matrix-valued functions that map a symmetric vector space \mathcal{X} to the space of linear maps between two other symmetric vector spaces \mathcal{Y} and \mathcal{Z} . These functions are said to be \mathbb{G} -equivariant if they satisfy a similar equivariance constraint as in [definition 2.14](#), but with the group action on the output space being the conjugation action of the group on the space of linear maps between \mathcal{Y} and \mathcal{Z} , which is defined as follows:

Definition 2.15 (\mathbb{G} -equivariant matrix-valued function). *Let \mathcal{X} , \mathcal{Y} and \mathcal{Z} be three \mathbb{G} -symmetric vector spaces, with respective group actions $\triangleright_{\mathcal{X}}$, $\triangleright_{\mathcal{Y}}$ and $\triangleright_{\mathcal{Z}}$. A matrix-valued function $\mathbf{F} : \mathcal{X} \mapsto \mathbb{R}^{|\mathcal{Z}| \times |\mathcal{Y}|}$ is said to be \mathbb{G} -equivariant if its image is the space of \mathbb{G} -homomorphisms between \mathcal{Y} and \mathcal{Z} ; $\text{Hom}_{\mathbb{G}}(\mathcal{Y}, \mathcal{Z})$, if it satisfies:*

$$g \triangleright_{\mathcal{Z}} \mathbf{F}(\mathbf{x}) \triangleright_{\mathcal{Y}}^{-1} = \mathbf{F}(g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, g \in \mathbb{G}. \iff \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \mathbf{F} \downarrow & & \downarrow \mathbf{F} \\ \text{Hom}_{\mathbb{G}}(\mathcal{Y}, \mathcal{Z}) & \xrightarrow{\triangleright_{\mathcal{Z}} \cdot \triangleright_{\mathcal{Y}}^{-1}} & \text{Hom}_{\mathbb{G}}(\mathcal{Y}, \mathcal{Z}) \end{array}$$

Such a constraint is also referred to as \mathbb{G} -steerability constraint in [Weiler et al. \(2021\)](#), and has well known algebraic structural constraints fully characterized in [Lang and Weiler \(2020\)](#).

Structure of \mathbb{G} -equivariant linear maps When restricting our focus to linear maps between symmetric vector spaces, the \mathbb{G} -equivariance constraint imposes strong structural constraints on the degrees of freedom of the map. We will frequently leverage this known structure, which is summarized in the following proposition:

Proposition 2.2 (Structure of \mathbb{G} -homomorphisms / intertwiners / \mathbb{G} -equivariant linear maps). *Let \mathbb{G} be a compact group and $\mathbf{A} \in \text{Hom}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y})$ be a \mathbb{G} -equivariant linear map between two (real) \mathbb{G} -symmetric vector spaces \mathcal{X} and \mathcal{Y} , with isotypic decompositions ([theorem 2.1](#)) given by:*

$$\mathcal{X} = \bigoplus_{k=1}^{n_{\text{iso}}} \mathcal{X}^{(k)} = \bigoplus_{k=1}^{n_{\text{iso}}} \bigoplus_{i=1}^{m_k^{\mathcal{X}}} \mathcal{X}_i^{(k)} \quad \text{and} \quad \mathcal{Y} = \bigoplus_{k=1}^{n_{\text{iso}}} \mathcal{Y}^{(k)} = \bigoplus_{k=1}^{n_{\text{iso}}} \bigoplus_{j=1}^{m_k^{\mathcal{Y}}} \mathcal{Y}_j^{(k)},$$

where n_{iso} denotes the number of isotypic subspaces, and $m_k^{\mathcal{X}}$ and $m_k^{\mathcal{Y}}$ denote the multiplicities of the irreducible representation $\bar{\rho}_k : \mathbb{G} \rightarrow \mathbb{GL}(\bar{\mathcal{V}}_k)$ in \mathcal{X} and \mathcal{Y} , respectively. Each $\mathcal{X}_i^{(k)}$ and $\mathcal{Y}_j^{(k)}$ is

isometrically isomorphic to $\bar{\mathcal{V}}_k$ (see [theorem 2.1](#)). Hence, in the isotypic bases, the map \mathbf{A} decomposes block-diagonally into n_{iso} blocks corresponding to homomorphisms between isotypic subspaces of the same type, that is:

$$\mathbf{A} = \bigoplus_{k=1}^{n_{\text{iso}}} \mathbf{A}^{(k)} \quad \text{where} \quad \mathbf{A}^{(k)} \in \text{Hom}_{\mathbb{G}}(\mathcal{X}^{(k)}, \mathcal{Y}^{(k)}).$$

Furthermore, the map $\mathbf{A}^{(k)}$ decomposes into $m_k^x \times m_k^y$ blocks of endomorphisms of the irreducible subspace $\bar{\mathcal{V}}_k$. That is:

$$\mathbf{A}^{(k)} = \begin{bmatrix} \mathbf{A}_{1,1}^{(k)} & \cdots & \mathbf{A}_{1,m_k^x}^{(k)} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{m_k^y,1}^{(k)} & \cdots & \mathbf{A}_{m_k^y,m_k^x}^{(k)} \end{bmatrix} \quad \text{where} \quad \mathbf{A}_{i,j}^{(k)} \in \text{End}_{\mathbb{G}}(\bar{\mathcal{V}}_k), \forall i \in [1, m_k^y], j \in [1, m_k^x].$$

Consequently, depending on the type of irreducible representation $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$, each sub-block is constrained to be in the span of the corresponding basis elements in [eq. 2.26](#). Consequently, if we denote by \mathbb{B} the basis set of \mathbb{K} , we have that the map $\mathbf{A}^{(k)}$ can be expressed in tensor product form as:

$$\mathbf{A}^{(k)} = \sum_{b \in \mathbb{B}} \Theta_b^{(k)} \otimes \Psi_k(b), \quad \text{where} \quad \Theta_b^{(k)} \in \mathbb{R}^{m_k^y \times m_k^x}, \Psi_k : \mathbb{B} \rightarrow \text{End}_{\mathbb{G}}(\bar{\mathcal{V}}_k) \quad (2.30)$$

With $[\Theta_b^{(k)}]_{i,j} = \langle \mathbf{A}_{i,j}^{(k)}, \Psi_k(b) \rangle$ denoting the basis expansion coefficient of the i -th, j -th endomorphism sub-block with the basis element $\Psi_k(b)$.

2.4 \mathbb{G} -invariant random variables

This section presents the background on probability theory in the presence of symmetry priors necessary to derive the core results of the thesis.

To begin, we recall the standard definition of a random variable taking values in a probability space. Let $(\mathcal{X}, \Sigma_{\mathcal{X}}, \mathbb{P}_{\mathbf{x}})$ be a probability space, with \mathcal{X} a set, $\Sigma_{\mathcal{X}}$ a σ -algebra of subsets of \mathcal{X} , and $\mathbb{P}_{\mathbf{x}} : \Sigma_{\mathcal{X}} \mapsto [0, 1]$ the marginal probability measure of \mathbf{x} . A σ -algebra $\Sigma_{\mathcal{X}}$ is the collection of subsets of \mathcal{X} that are treated as measurable events; it contains \mathcal{X} , is closed under complements, and is closed under countable unions. We denote a random variable \mathbf{x} taking values in \mathcal{X} by $\mathbf{x} : \Omega \rightarrow \mathcal{X}$, where Ω is the underlying sample space. Furthermore, we denote the marginal probability of an event $\mathbb{X} \subseteq \mathcal{X}$ by $\mathbb{P}(\mathbf{x} \in \mathbb{X}) := \mathbb{P}_{\mathbf{x}}(\mathbb{X})$ and the conditional probability of an event $\mathbb{Y} \subseteq \mathcal{Y}$ given $\mathbf{x} = \mathbf{x}$ by $\mathbb{P}(\mathbf{y} \in \mathbb{Y} | \mathbf{x} = \mathbf{x}) := \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y} | \mathbf{x})$, where $\mathbb{P}_{\mathbf{y}|\mathbf{x}} : \Sigma_{\mathcal{Y}} \times \mathcal{X} \mapsto [0, 1]$ is the conditional probability distribution of \mathbf{y} given \mathbf{x} . Finally, we denote realizations of \mathbf{x} by $\mathbf{x} \in \mathcal{X}$. Expectations of an observable $f : \mathcal{X} \rightarrow \mathbb{R}$ under \mathbf{x} are denoted by $\mathbb{E}_{\mathbf{x}}[f] = \int_{\mathcal{X}} f(\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x})$, and conditional expectation of an observable $h : \mathcal{Y} \rightarrow \mathbb{R}$ given $\mathbf{x} = \mathbf{x}$ by $\mathbb{E}_{\mathbf{y}|\mathbf{x}}[h(\mathbf{y})] = \int_{\mathcal{Y}} h(\mathbf{y}) \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y} | \mathbf{x})$.

Probability symmetry priors The first fundamental symmetry prior we will encounter is the invariance of the marginal distribution of a random variable under the action of a compact symmetry group \mathbb{G} :

Definition 2.16 (\mathbb{G} -invariant random variable). *Let $(\mathcal{X}, \Sigma_{\mathcal{X}}, \mathbb{P}_{\mathbf{x}})$ be a probability space, and let \mathbb{G} be a compact symmetry group acting on \mathcal{X} via the group action $\triangleright_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$ ([definition 2.2](#)).*

A random variable \mathbf{x} taking values in \mathcal{X} is said to be symmetric with respect to \mathbb{G} if the probability measure $\mathbb{P}_{\mathbf{x}}$ is invariant under the group action:

$$\mathbb{P}_{\mathbf{x}}(\mathbb{X}) = \mathbb{P}_{\mathbf{x}}(g \triangleright_{\mathcal{X}} \mathbb{X}), \quad \forall g \in \mathbb{G}, \mathbb{X} \in \Sigma_{\mathcal{X}}. \quad \Longleftrightarrow \quad \begin{array}{ccc} \Sigma_{\mathcal{X}} & \xrightarrow{\triangleright_{\mathcal{X}}} & \Sigma_{\mathcal{X}} \\ & \searrow \mathbb{P}_{\mathbf{x}} & \downarrow \mathbb{P}_{\mathbf{x}} \\ & & [0, 1] \end{array} \quad (2.31)$$

As we will see in chapters 6 to 8, the distributions of world states $\mathbb{P}_{\mathbf{s}}$, robot actions $\mathbb{P}_{\mathbf{a}}$, and state observations $\mathbb{P}_{\mathbf{o}}$ are common examples of such \mathbb{G} -invariant random variables.

The second fundamental symmetry prior we will encounter is the invariance of the conditional distribution of a random variable given another random variable under the action of a compact symmetry group \mathbb{G} :

Definition 2.17 (\mathbb{G} -invariant conditional distribution). *Let $(\mathcal{X}, \Sigma_{\mathcal{X}}, \mathbb{P}_{\mathbf{x}})$ and $(\mathcal{Y}, \Sigma_{\mathcal{Y}}, \mathbb{P}_{\mathbf{y}})$ be probability spaces, and let \mathbb{G} be a compact symmetry group acting on both \mathcal{X} and \mathcal{Y} via the group actions $\triangleright_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$ and $\triangleright_{\mathcal{Y}} : \mathbb{G} \times \mathcal{Y} \rightarrow \mathcal{Y}$, respectively. The conditional distribution $\mathbb{P}_{\mathbf{y}|\mathbf{x}} : \Sigma_{\mathcal{Y}} \times \mathcal{X} \mapsto [0, 1]$ of \mathbf{y} given \mathbf{x} is said to be \mathbb{G} -invariant if the following conditions hold:*

$$\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) = \mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_{\mathcal{Y}} \mathbb{Y} | g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall g \in \mathbb{G}, \mathbf{x} \in \mathcal{X}, \mathbb{Y} \in \Sigma_{\mathcal{Y}}, \quad \Longleftrightarrow \quad \begin{array}{ccc} \Sigma_{\mathcal{Y}} \times \mathcal{X} & \xrightarrow{(\triangleright_{\mathcal{Y}}, \triangleright_{\mathcal{X}})} & \Sigma_{\mathcal{Y}} \times \mathcal{X} \\ & \searrow \mathbb{P}_{\mathbf{y}|\mathbf{x}} & \downarrow \mathbb{P}_{\mathbf{y}|\mathbf{x}} \\ & & [0, 1] \end{array} \quad (2.32)$$

As we will see in chapters 6 to 8, the world transition dynamics $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ (eq. 2.13), observation transition dynamics $\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}$ (eq. 2.15), optimal control policies π_* , sensor models $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ (eq. 2.14), and optimal state estimation models $\mathbb{P}_{\mathbf{s}|\mathbf{o}}^{\Delta t}$ (eq. 2.15) are common examples of such \mathbb{G} -invariant conditional distributions.

Intuitively, these conditional distributions can be viewed as probabilistic generalizations of equivariant maps between the spaces \mathcal{X} and \mathcal{Y} . To make this precise, we focus on a frequently encountered scenario in this thesis: for a given pair of random variables (\mathbf{x}, \mathbf{y}) , \mathbf{x} is a \mathbb{G} -invariant random variable and the conditional distribution of \mathbf{y} given \mathbf{x} is \mathbb{G} -invariant. In this setting, the marginal, joint, and conditional expectations inherit the symmetry constraints stated in the following proposition.

Proposition 2.3. *Let (\mathbf{x}, \mathbf{y}) be two random variables such that \mathbf{x} is \mathbb{G} -symmetric (definition 2.16) and the conditional distribution of \mathbf{y} given \mathbf{x} is \mathbb{G} -invariant (eq. 2.32). Then the marginal distribution of \mathbf{y} is also \mathbb{G} -invariant; therefore, \mathbf{y} is \mathbb{G} -symmetric. More precisely, if*

$$\mathbb{P}_{\mathbf{x}}(\mathbb{X}) = \mathbb{P}_{\mathbf{x}}(g \triangleright_{\mathcal{X}} \mathbb{X}), \quad \text{and} \quad \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) = \mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_{\mathcal{Y}} \mathbb{Y} | g \triangleright_{\mathcal{X}} \mathbf{x}), \\ \forall g \in \mathbb{G}, \mathbf{x} \in \mathcal{X}, \mathbb{X} \in \Sigma_{\mathcal{X}}, \mathbb{Y} \in \Sigma_{\mathcal{Y}},$$

then the marginal distribution of \mathbf{y} is \mathbb{G} -invariant:

$$\mathbb{P}_{\mathbf{y}}(\mathbb{Y}) = \mathbb{P}_{\mathbf{y}}(g \triangleright_{\mathcal{Y}} \mathbb{Y}), \quad \forall g \in \mathbb{G}, \mathbb{Y} \in \Sigma_{\mathcal{Y}}, \quad (2.34)$$

the joint distribution of (\mathbf{x}, \mathbf{y}) is \mathbb{G} -invariant:

$$\mathbb{P}_{\mathbf{xy}}(\mathbb{X}, \mathbb{Y}) = \mathbb{P}_{\mathbf{xy}}(g \triangleright_{\mathcal{X}} \mathbb{X}, g \triangleright_{\mathcal{Y}} \mathbb{Y}), \quad \forall g \in \mathbb{G}, \mathbb{X} \in \Sigma_{\mathcal{X}}, \mathbb{Y} \in \Sigma_{\mathcal{Y}}. \quad (2.35)$$

Moreover, if the conditional first moment is uniformly bounded: $\sup_{\mathbf{x} \in \mathcal{X}} \int_{\mathcal{Y}} \|\mathbf{y}\| \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) < \infty$, the conditional expectation of \mathbf{y} given \mathbf{x} is a \mathbb{G} -equivariant map (definition 2.14):

$$\mathbb{E}[\mathbf{y}|\mathbf{x} = g \triangleright_{\mathcal{X}} \mathbf{x}] = g \triangleright_{\mathcal{Y}} \mathbb{E}[\mathbf{y}|\mathbf{x} = \mathbf{x}] \quad \forall g \in \mathbb{G}, \mathbf{x} \in \mathcal{X}. \quad \iff \quad \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \downarrow \mathbb{E}[\mathbf{y}|\mathbf{x}=\cdot] & & \downarrow \mathbb{E}[\mathbf{y}|\mathbf{x}=\triangleright_{\mathcal{X}}\cdot] \\ \mathcal{Y} & \xrightarrow{\triangleright_{\mathcal{Y}}} & \mathcal{Y} \end{array} \quad (2.36)$$

Proof. To prove eq. 2.34, let $g \in \mathbb{G}$ and $\mathbb{Y} \in \Sigma_{\mathcal{Y}}$. Then

$$\begin{aligned} \mathbb{P}_{\mathbf{y}}(g \triangleright_{\mathcal{Y}} \mathbb{Y}) &= \int_{\mathcal{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_{\mathcal{Y}} \mathbb{Y}|\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \\ &= \int_{\mathcal{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|g^{-1} \triangleright_{\mathcal{X}} \mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \quad (\text{by eq. 2.32}) \\ &= \int_{\mathcal{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \quad (\text{change of variables and } \mathbb{P}_{\mathbf{x}} \text{ } \mathbb{G}\text{-invariance}) \\ &= \mathbb{P}_{\mathbf{y}}(\mathbb{Y}). \end{aligned}$$

To prove eq. 2.35, let $g \in \mathbb{G}$, $\mathbb{X} \in \Sigma_{\mathcal{X}}$, and $\mathbb{Y} \in \Sigma_{\mathcal{Y}}$. By construction, we have,

$$\mathbb{P}_{\mathbf{xy}}(\mathbb{X}, \mathbb{Y}) = \int_{\mathbb{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}), \quad \forall \mathbb{X} \in \Sigma_{\mathcal{X}}, \mathbb{Y} \in \Sigma_{\mathcal{Y}}.$$

Therefore,

$$\begin{aligned} \mathbb{P}_{\mathbf{xy}}(g \triangleright_{\mathcal{X}} \mathbb{X}, g \triangleright_{\mathcal{Y}} \mathbb{Y}) &= \int_{g \triangleright_{\mathcal{X}} \mathbb{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_{\mathcal{Y}} \mathbb{Y}|\tilde{\mathbf{x}}) \mathbb{P}_{\mathbf{x}}(d\tilde{\mathbf{x}}) \\ &= \int_{\mathbb{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_{\mathcal{Y}} \mathbb{Y}|g \triangleright_{\mathcal{X}} \mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \quad (\tilde{\mathbf{x}} = g \triangleright_{\mathcal{X}} \mathbf{x} \text{ and } \mathbb{P}_{\mathbf{x}} \text{ } \mathbb{G}\text{-invariance}) \\ &= \int_{\mathbb{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \quad (\text{by eq. 2.32}) \\ &= \mathbb{P}_{\mathbf{xy}}(\mathbb{X}, \mathbb{Y}). \end{aligned}$$

To prove eq. 2.36, we write

$$\begin{aligned} \mathbb{E}[\mathbf{y}|\mathbf{x} := g \triangleright_{\mathcal{X}} \mathbf{x}] &= \int_{\mathcal{Y}} \mathbf{y} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|g \triangleright_{\mathcal{X}} \mathbf{x}) \\ &= \int_{\mathcal{Y}} g \triangleright_{\mathcal{Y}} \mathbf{y} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) \quad (\text{change of variables and eq. 2.32}) \\ &= g \triangleright_{\mathcal{Y}} \int_{\mathcal{Y}} \mathbf{y} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) \quad (\text{linearity of the action on } \mathcal{Y}) \\ &= g \triangleright_{\mathcal{Y}} \mathbb{E}[\mathbf{y}|\mathbf{x} = \mathbf{x}]. \end{aligned}$$

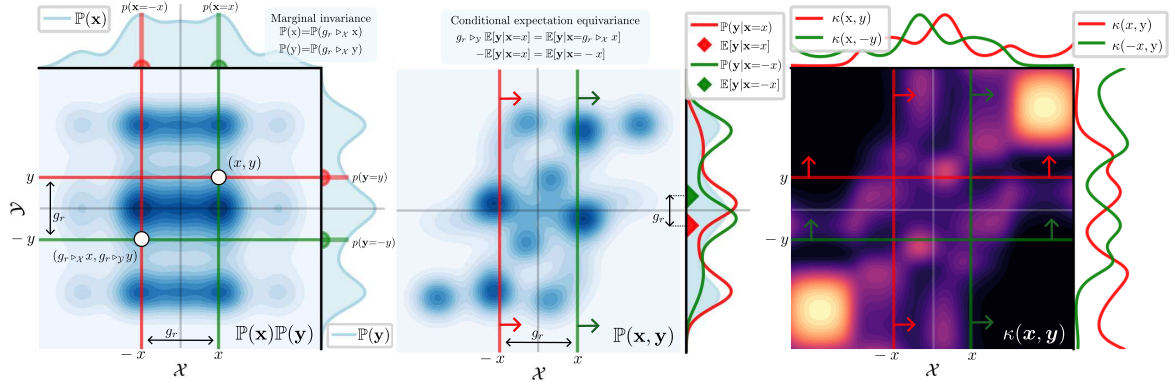


Figure 2.3 Example of \mathbb{G} -invariant random variables $(x, y) \sim \mathcal{X} \times \mathcal{Y} \subset \mathbb{R} \times \mathbb{R}$, whose marginals $\mathbb{P}_x(x)$ and $\mathbb{P}_y(y)$; joint $\mathbb{P}_{x,y}(x, y)$; and conditional $\mathbb{P}_{y|x}(y|x)$ distributions are invariant to reflections of the data: $g_r \triangleright_{\mathcal{X}} x = -x$ and $g_r \triangleright_{\mathcal{Y}} y = -y$, where g_r denotes the reflection element of the reflection symmetry group $\mathbb{C}_2 := \{e, g_r | g_r^2 = e\}$ (see prop. 2.3). Consequently, the PMD $\kappa(x, y)$ is \mathbb{C}_2 -invariant. □

A self-contained illustration of prop. 2.3 is shown in fig. 2.3 for two scalar random variables with the simplest non-trivial symmetry group, the reflection symmetry group $\mathbb{C}_2 := \{e, g_r | g_r^2 = e\}$ (see fig. 2.1), where e denotes the identity element and g_r denotes the reflection element (see section 2.3).

A common algebraic trick we will use throughout the thesis is to leverage the relation between the conditional distributions defined by symmetry related conditioning events, such as $\mathbf{x} = x$ and $\mathbf{x} = g_r \triangleright_{\mathcal{X}} x$, to change variable in the integral of the conditional expectation of \mathbf{y} given \mathbf{x} . This is formalized in the following corollary.

Corollary 2.1. *Let (x, y) be two random variables such that \mathbf{x} is \mathbb{G} -symmetric (definition 2.16) and the conditional distribution of \mathbf{y} given \mathbf{x} is \mathbb{G} -invariant (eq. 2.32). Then, we have that*

$$\int_{\mathcal{Y}} \mathbb{P}_{y|x}(dy|x) = \int_{g_r \triangleright_{\mathcal{Y}} y} \mathbb{P}_{y|x}(g_r \triangleright_{\mathcal{Y}} dy|x) = \int_{\mathcal{Y}} \mathbb{P}_{y|x}(dy|g^{-1} \triangleright_{\mathcal{X}} x) \quad \forall g \in \mathbb{G}.$$

Optimal transport paths between \mathbb{G} -invariant distributions

In chapter 8, we study the implications of symmetry priors in generative modelling paradigms used to parameterize stochastic optimal control policies. This leads to the problem of optimal transport between two probability distributions that are \mathbb{G} -invariant with respect to a given compact symmetry group \mathbb{G} . This problem has been studied in the context of \mathbb{G} -equivariant flow matching by Klein et al. (2023), specifically for optimal transport between two *marginal* probability distributions $\mathbb{P}_0 : \Sigma_{\mathcal{Y}} \mapsto [0, 1]$ and $\mathbb{P}_1 : \Sigma_{\mathcal{Y}} \mapsto [0, 1]$ that are both \mathbb{G} -invariant (definition 2.16).

Given that in optimal control problems, the control policy is parameterized by a *conditional* probability distributions of control actions given conditioning states, our focus is on optimal transport maps between two \mathbb{G} -invariant *conditional* probability distributions $\mathbb{P}_0 : \Sigma_{\mathcal{Y}} \times \mathcal{X} \rightarrow [0, 1]$ and $\mathbb{P}_1 : \Sigma_{\mathcal{Y}} \times \mathcal{X} \rightarrow [0, 1]$, where the solution is a family of optimal transport paths indexed by the realizations of the conditioning random variable \mathbf{x} . To tackle this scenario, we first restate the main results of Klein et al. (2023) in our measure theoretic notation, and then extend these results to the conditional setting.

Optimal transport between marginal distributions In optimal transport, given two marginal probability distributions \mathbb{P}_0 and \mathbb{P}_1 on a shared domain \mathcal{Y} , the goal is to find a transport plan $T : \Sigma_{\mathcal{Y}} \times \Sigma_{\mathcal{Y}} \mapsto [0, 1]$ that minimizes the expected transport cost between the two distributions. Where the transport cost $c : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a strictly convex distance metric $c(\mathbf{y}_0, \mathbf{y}_1) = d(\mathbf{y}_0, \mathbf{y}_1)^p$ for some $p \geq 1$. The optimal transport problem can be formalized as the following optimization problem over the space of transport plans $\Pi(\mathbb{P}_0, \mathbb{P}_1)$, which consists of all probability measures on $\mathcal{Y} \times \mathcal{Y}$ with marginals \mathbb{P}_0 and \mathbb{P}_1 :

$$\inf_T \int_{\mathcal{Y} \times \mathcal{Y}} c(\mathbf{y}, \mathbf{y}') T(d\mathbf{y}, d\mathbf{y}'), \quad \text{s.t.} \quad \int_{\mathcal{Y}} T(d\mathbb{Y}, d\mathbf{y}) = \mathbb{P}_0(\mathbb{Y}), \quad \int_{\mathcal{Y}} T(d\mathbf{y}, d\mathbb{Y}) = \mathbb{P}_1(\mathbb{Y}), \quad \forall \mathbb{Y} \in \Sigma_{\mathcal{Y}}. \quad (2.37)$$

In the context where the source \mathbb{P}_0 and target \mathbb{P}_1 distributions are \mathbb{G} -invariant (eq. 2.31), we have that the domain \mathcal{Y} admits a group action $\triangleright_{\mathcal{Y}} : \mathbb{G} \times \mathcal{Y} \rightarrow \mathcal{Y}$ that preserve distances in \mathcal{Y} . Consequently, the transport cost becomes a \mathbb{G} -invariant map, that is $c(g \triangleright_{\mathcal{Y}} \mathbf{y}_0, g \triangleright_{\mathcal{Y}} \mathbf{y}_1) = c(\mathbf{y}_0, \mathbf{y}_1)$ for all $g \in \mathbb{G}$ and $\mathbf{y}_0, \mathbf{y}_1 \in \mathcal{Y}$. In this setting, Klein et al. (2023) introduces the symmetrized cost $c^{\mathbb{G}} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ defined as the minimum transport cost between points in the same \mathbb{G} -orbit:

$$c^{\mathbb{G}}(\mathbf{y}_0, \mathbf{y}_1) := \min_{g \in \mathbb{G}} c(\mathbf{y}_0, g \triangleright_{\mathcal{Y}} \mathbf{y}_1) = \min_{g \in \mathbb{G}} c(g \triangleright_{\mathcal{Y}} \mathbf{y}_0, \mathbf{y}_1), \quad \forall \mathbf{y}_0, \mathbf{y}_1 \in \mathcal{Y}. \quad (2.38)$$

The following theorem formalizes the optimal transport properties of marginal invariant distributions, translating the core findings of Klein et al. (2023) into the measure-theoretic notation of this thesis.

Theorem 2.2 (Equivariant OT for \mathbb{G} -invariant marginal distributions (Klein et al., 2023)). *Let \mathbb{P}_0 and \mathbb{P}_1 be two \mathbb{G} -invariant marginal probability distributions (definition 2.16) on the measure space $(\mathcal{Y}, \Sigma_{\mathcal{Y}})$. Let $c : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a strictly convex, \mathbb{G} -invariant cost function. Then:*

1. *The optimal transport plan $T(\cdot, \cdot) \in \Pi(\mathbb{P}_0, \mathbb{P}_1)$ is \mathbb{G} -invariant:*

$$T(\mathbb{Y}_0, \mathbb{Y}_1) = T(g \triangleright_{\mathcal{Y}} \mathbb{Y}_0, g \triangleright_{\mathcal{Y}} \mathbb{Y}_1), \quad \forall g \in \mathbb{G}, \mathbb{Y}_0, \mathbb{Y}_1 \in \Sigma_{\mathcal{Y}}. \quad (2.39)$$

2. *If the optimal transport map $w^* : \mathcal{Y} \rightarrow \mathcal{Y}$ exists and is unique, it is \mathbb{G} -equivariant:*

$$w^*(g \triangleright_{\mathcal{Y}} \mathbf{y}) = g \triangleright_{\mathcal{Y}} w^*(\mathbf{y}), \quad \forall g \in \mathbb{G}, \mathbf{y} \in \mathcal{Y}. \quad (2.40)$$

3. *The optimal transport map w^* is also an optimal transport map for the symmetrized cost $c^{\mathbb{G}}$.*
4. *Assuming \mathcal{Y} is a vector space, the optimal velocity field $\mathbf{v}_*(\mathbf{y}, k)$ generating the optimal transport-map interpolation $\mathbf{y}^{(k)} = (1 - k)\mathbf{y}^{(0)} + k w(\mathbf{y}^{(0)})$ is \mathbb{G} -equivariant:*

$$\mathbf{v}_*(g \triangleright_{\mathcal{Y}} \mathbf{y}, k) = g \triangleright_{\mathcal{Y}} \mathbf{v}_*(\mathbf{y}, k), \quad \forall g \in \mathbb{G}, k \in [0, 1]. \quad (2.41)$$

Proof. A detailed proof is provided by (Klein et al., 2023, Appendix B.1). □

Conditional optimal transport We now switch our focus to conditional optimal transport, where the goal is to find a family of optimal transport plans $\{\Pi_{\mathbf{x}}(\mathbb{P}_0, \mathbb{P}_1)\}_{\mathbf{x} \in \mathcal{X}}$ between *conditional* probability distributions $\mathbb{P}_0 : \Sigma_{\mathcal{Y}} \times \mathcal{X} \rightarrow [0, 1]$ and $\mathbb{P}_1 : \Sigma_{\mathcal{Y}} \times \mathcal{X} \rightarrow [0, 1]$, conditioned on the realizations of the a random variable \mathbf{x} , taking values in the measure space $(\mathcal{X}, \Sigma_{\mathcal{X}})$. That is, to solve the following

family of optimal transport problems:

$$\inf_T \int_{\mathcal{Y} \times \mathcal{Y}} c(\mathbf{y}, \mathbf{y}' | \mathbf{x}) T(d\mathbf{y}, d\mathbf{y}' | \mathbf{x}), \quad \text{s.t.} \quad \begin{cases} \int_{\mathcal{Y}} T(\mathbb{Y}, d\mathbf{y} | \mathbf{x}) = \mathbb{P}_0(\mathbb{Y} | \mathbf{x}), \\ \int_{\mathcal{Y}} T(d\mathbf{y}, \mathbb{Y} | \mathbf{x}) = \mathbb{P}_1(\mathbb{Y} | \mathbf{x}), \end{cases} \quad \forall \mathbb{Y} \in \Sigma_{\mathcal{Y}}, \mathbf{x} \in \mathcal{X}. \quad (2.42)$$

This section extends the previous analysis to a family of optimal transport paths—indexed by a secondary \mathbb{G} -invariant random variable \mathbf{x} —between two *conditional* probability distributions $\mathbb{P}_0(\cdot | \mathbf{x})$ and $\mathbb{P}_1(\cdot | \mathbf{x})$ that are both \mathbb{G} -invariant (definition 2.17).

Because the conditional distributions $\mathbb{P}_i(\cdot | \mathbf{x})$ behave identically to marginal distributions for any fixed realization of the conditioning variable \mathbf{x} , the properties from theorem 2.2 trivially extend to the conditional optimal transport plan $T(\cdot, \cdot | \mathbf{x})$ and map $T(\cdot | \mathbf{x})$.

Corollary 2.2 (Conditional Equivariant Optimal Transport). *Let $\mathbb{P}_0(\cdot | \mathbf{x})$ and $\mathbb{P}_1(\cdot | \mathbf{x})$ be two \mathbb{G} -invariant conditional probability distributions. Under the identical assumptions of theorem 2.2, for any realization $\mathbf{x} \in \mathcal{X}$:*

1. *The optimal transport plan is conditionally invariant:*

$$T(\mathbb{Y}_0, \mathbb{Y}_1 | \mathbf{x}) = T(g \triangleright_{\mathcal{Y}} \mathbb{Y}_0, g \triangleright_{\mathcal{Y}} \mathbb{Y}_1 | g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall g \in \mathbb{G}, \mathbb{Y}_0, \mathbb{Y}_1 \in \Sigma_{\mathcal{Y}}.$$

2. *If the conditional optimal transport map $w^* : \mathcal{Y} \times \mathcal{X} \rightarrow \mathcal{Y}$ exists and is unique for any given \mathbf{x} , it is conditionally \mathbb{G} -equivariant:*

$$g \triangleright_{\mathcal{Y}} w^*(\mathbf{y}, \mathbf{x}) = w^*(g \triangleright_{\mathcal{Y}} \mathbf{y}, g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall g \in \mathbb{G}, \mathbf{y} \in \mathcal{Y}.$$

3. *The optimal transport map w^* is also an optimal transport map for the symmetrized cost $c^{\mathbb{G}}$, for any fixed realization of the conditioning variable \mathbf{x} .*
4. *Assuming \mathcal{Y} is a vector space, the optimal velocity field $\mathbf{v}_*(\mathbf{y}, \mathbf{x}, k)$ generating the optimal transport-map interpolation $\mathbf{y}^{(k)} = (1 - k)\mathbf{y}^{(0)} + k w^*(\mathbf{y}^{(0)}, \mathbf{x})$ is \mathbb{G} -equivariant:*

$$g \triangleright_{\mathcal{Y}} \mathbf{v}_*(\mathbf{y}, \mathbf{x}, k) = \mathbf{v}_*(g \triangleright_{\mathcal{Y}} \mathbf{y}, g \triangleright_{\mathcal{X}} \mathbf{x}, k), \quad \forall g \in \mathbb{G}, k \in [0, 1].$$

Proof. For any fixed realization $\mathbf{x} \in \mathcal{X}$, the \mathbb{G} -invariance of the conditional distributions implies $\mathbb{P}_i(\mathbb{Y} | \mathbf{x}) = \mathbb{P}_i(g \triangleright_{\mathcal{Y}} \mathbb{Y} | g \triangleright_{\mathcal{X}} \mathbf{x})$ for $i \in \{0, 1\}$. Evaluating the optimal transport problem between these distributions at the transformed conditioning state $g \triangleright_{\mathcal{X}} \mathbf{x}$ is strictly equivalent to solving the marginal optimal transport problem between the transformed measures $\mathbb{P}_0(g \triangleright_{\mathcal{Y}} \cdot | \mathbf{x})$ and $\mathbb{P}_1(g \triangleright_{\mathcal{Y}} \cdot | \mathbf{x})$. Applying theorem 2.2 directly to these fixed slices yields the conditional invariance of the plan and the conditional equivariance of both the map and its corresponding velocity field. \square

Chapter 3

Symmetries of robot dynamics

This chapter introduces the different types of symmetry present in robot dynamics and their implications. We begin in [section 3.1](#) by formalizing symmetries as energy-preserving transformations of the robot’s state associated with a Euclidean isometry—that is, with a rigid translation, rotation, or reflection of the Euclidean space in which the robot evolves. Then, in [chapter 4](#), we introduce morphological symmetries—a specific class of symmetries associated with structural regularities in a robot’s morphology, featuring a rich set of geometric and algebraic constraints that we will leverage in later chapters as inductive biases for data-driven robot modelling, estimation, and control.

Note: In [section 2.1](#), we characterized the symmetries of Newtonian (classical) mechanics as the rigid transformations of three-dimensional Euclidean space, i.e., the *Euclidean group* \mathbb{E}_3 , which comprises all rotations, translations, and reflections of \mathbb{R}^3 . In the classical interpretation, these symmetries are considered *passive* transformations affecting only the chosen inertial reference frame ([Bloch et al., 1996](#); [Murray et al., 2017](#)). Their primary implication is the invariance of Newtonian physics to the choice of inertial reference frame, given that translating, rotating, or reflecting that reference frame does not alter the underlying dynamics ([Bloch et al., 1996](#); [Noether, 1983](#)). In this chapter, we reinterpret these symmetries as *active* transformations acting on the robot and its operational environment, leaving the chosen inertial reference frame unchanged.

3.1 Symmetries as energy-preserving transformations

In the context of dynamical systems, a symmetry is a transformation that relates two distinct system states with identical—or equivalent—dynamics. Concretely, this means that when subjected to equivalent forcing, symmetry-related states undergo equivalent temporal evolution. Identifying such symmetries is therefore highly valuable, as they act as powerful structural priors for modelling and control: the ability to model or control a given state $(\mathbf{q}, \dot{\mathbf{q}})$ immediately extends to all of its symmetry-related counterparts (see [fig. 3.1](#)).

The goal of this section is to formalize the notion of dynamics equivalence and describe how symmetries in dynamics translate to geometric and algebraic constraints on the system’s state space, equations of motion, and generalized mass matrix.

In robotics, symmetries are energy-preserving transformations—associated with reflections, translations, or rotations in Euclidean space—that relate dynamically equivalent states ([Bloch et al., 1996](#);




Figure 3.1 Example symmetry transformations relating two dynamically equivalent states. **Left:** Double pendulum with reflection symmetry $\mathbb{G} = \mathbb{C}_2 = \{e, g_s \mid g_s^2 = e\}$. The state $(\mathbf{q} = [\theta_1, \theta_2], \dot{\mathbf{q}} = [\dot{\theta}_1, \dot{\theta}_2])$ is transformed by reflection symmetry g_s defined via $(g_s \triangleright_{\mathcal{Q}} \mathbf{q} = [-\theta_1, -\theta_2], g_s \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} = [-\dot{\theta}_1, -\dot{\theta}_2])$. **Right:** Cart-pole with translation symmetry $\mathbb{G} = \mathbb{T}_x = \{g_{\Delta x} \mid \Delta x \in \mathbb{R}\}$. The state $(\mathbf{q} = [x, \theta], \dot{\mathbf{q}} = [\dot{x}, \dot{\theta}])$ is transformed by translation symmetry $g_{\Delta x}$ defined via $(g_{\Delta x} \triangleright_{\mathcal{Q}} \mathbf{q} = [x + \Delta x, \theta], g_{\Delta x} \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} = [\dot{x}, \dot{\theta}])$.

Ordoñez-Apaez et al., 2025; Wieber, 2006). The action of a symmetry on the robot’s state is formalized through a *group action* on the configuration space, $(\triangleright_{\mathcal{Q}}) : \mathbb{G} \times \mathcal{Q} \mapsto \mathcal{Q}$ (see definition 2.2). This action defines an (affine or linear) bijection on the configuration space for each symmetry transformation, where \mathbb{G} denotes the set of transformations forming the system’s symmetry group (see definition 2.1). For intuition, fig. 3.1 illustrates simple examples of group actions: a double pendulum with a reflection symmetry (linear action) and a cart-pole with a horizontal translation symmetry (affine action).

The action of a symmetry transformation on the robot’s generalized velocity coordinates and on accelerations or forces is derived via the push-forward and pull-back of the configuration-space action on the manifold \mathcal{Q} . These operations lift the symmetry action to the corresponding tangent and cotangent spaces, $\mathcal{T}_q \mathcal{Q}$ and $\mathcal{T}_q^* \mathcal{Q}$, respectively, yielding the following group actions:

$$\begin{aligned}
 (\triangleright_{\mathcal{T}_q \mathcal{Q}}) : \mathbb{G} \times \mathcal{T}_q \mathcal{Q} &\mapsto \mathcal{T}_q \mathcal{Q}, & (\triangleright_{\mathcal{T}_q^* \mathcal{Q}}) : \mathbb{G} \times \mathcal{T}_q^* \mathcal{Q} &\mapsto \mathcal{T}_q^* \mathcal{Q}, \\
 (g, \dot{\mathbf{q}}) &\mapsto g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} := \frac{\partial(g \triangleright_{\mathcal{Q}} \mathbf{q})}{\partial \mathbf{q}} \dot{\mathbf{q}} & (g, \boldsymbol{\tau}) &\mapsto g \triangleright_{\mathcal{T}_q^* \mathcal{Q}} \boldsymbol{\tau} := \left(\frac{\partial(g \triangleright_{\mathcal{Q}} \mathbf{q})}{\partial \mathbf{q}} \right)^{-\top} \boldsymbol{\tau}
 \end{aligned} \tag{3.1}$$

Once again, fig. 3.1 grounds this formalism in intuitive transformations of velocities and forces for the double pendulum and cart-pole systems. The purpose of introducing this unified notation—although it may initially seem unnecessarily complex—is to enable the treatment of both simple and complex robotic systems, potentially exhibiting simple or intricate symmetry groups, within a single consistent mathematical framework based on linear algebra; see figs. 3.1 and 3.2 and definition 2.2.

Dynamics equivalence Having defined how symmetry transformations act on the system’s state space, we can formalize the concept of dynamics equivalence between two symmetry-related states, which intuitively implies that when the two states are subjected to equivalent forcing, they undergo equivalent temporal evolution. In classical mechanics, dynamics equivalence is commonly expressed via the invariance of the system’s Lagrangian function to symmetry transformations (Bloch et al., 1996; Murray et al., 2017; Wieber, 2006):

Definition 3.1 (Symmetric dynamical system). *Let $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}$ be the state of a dynamical system and let $L : \mathcal{T}\mathcal{Q} \mapsto \mathbb{R}$ be its Lagrangian function. The system is deemed to possess a symmetry group \mathbb{G} if its Lagrangian is \mathbb{G} -invariant. That is if:*

$$L(\mathbf{q}, \dot{\mathbf{q}}) = L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}), \quad \forall g \in \mathbb{G}, \forall (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}. \tag{3.2}$$

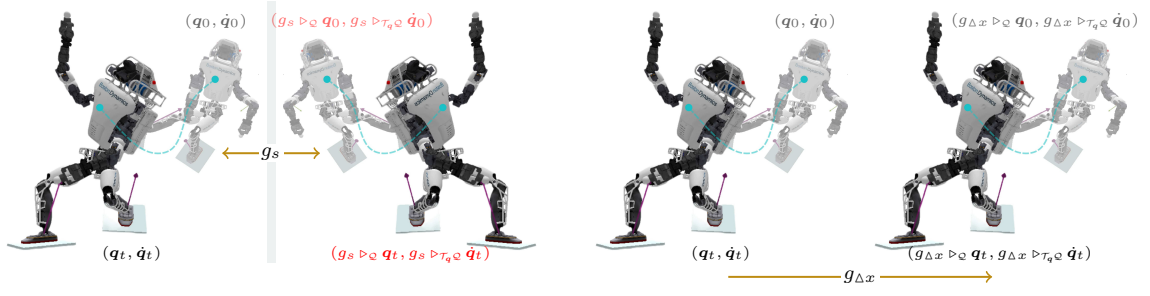


Figure 3.2 Example symmetry transformations of the dynamics of the Atlas humanoid robot. **Left:** action of the reflection of space g_s , relating the dynamically equivalent states (q_t, \dot{q}_t) and $(g_s \triangleright_Q q_t, g_s \triangleright_{T_q Q} \dot{q}_t)$. Crucially, the defined reflection action results in a state that is not reachable, i.e. $g_s \triangleright_Q q_t \notin \mathcal{Q}$, given that real-world rigid bodies cannot undergo a *true* reflection. **Right:** Example action of the translation symmetry group $\mathbb{G} = \mathbb{T}_x = \{g_{\Delta x} | \Delta x \in \mathbb{R}\}$, which for any horizontal translation $g_{\Delta x}$, relates the dynamically equivalent states (q_t, \dot{q}_t) and $(g_{\Delta x} \triangleright_Q q_t, g_{\Delta x} \triangleright_{T_q Q} \dot{q}_t)$.

Nomenclature: A dynamical system satisfying eq. 3.2 is referred to as a symmetric dynamical system, with \mathbb{G} referred to as its symmetry group. Furthermore, any pair of states related by a symmetry transformation are said to be symmetric states with equivalent dynamics.

The reason eq. 3.2 defines a *dynamical equivalence relation* is that, as recalled in section 2.1, the robot's equations of motion are derived directly from the Lagrangian via the principle of least action (see eq. 2.9). Consequently, deriving the equations of motion for two symmetry-related trajectories with fixed endpoints—namely a trajectory \bar{q}_{t_0, t_1} and its symmetry-transformed counterpart $g \triangleright \bar{q}_{t_0, t_1} := \{g \triangleright_Q q_t | q_t \in \bar{q}_{t_0, t_1}, t \in [t_0, t_1]\}$, for any $g \in \mathbb{G}$ (e.g., see fig. 3.1)—yields equivalent variational problems:

$$\begin{aligned} \arg \min_{\bar{q}_{t_0, t_1}} \int_{t_0}^{t_1} L(q_t, \dot{q}_t) dt \quad \text{and} \quad \arg \min_{g \triangleright \bar{q}_{t_0, t_1}} \int_{t_0}^{t_1} L(g \triangleright_Q q_t, g \triangleright_{T_q Q} \dot{q}_t) dt \\ \frac{d}{dt} \frac{\partial L(q, \dot{q})}{\partial(\dot{q})} - \frac{\partial L(q, \dot{q})}{\partial(q)} = \mathbf{0} \in \mathcal{T}_q \mathcal{Q}^*, \quad \frac{d}{dt} \frac{\partial L(g \triangleright_Q q, g \triangleright_{T_q Q} \dot{q})}{\partial(g \triangleright_{T_q Q} \dot{q})} - \frac{\partial L(g \triangleright_Q q, g \triangleright_{T_q Q} \dot{q})}{\partial(g \triangleright_Q q)} = \mathbf{0} \in \mathcal{T}_q \mathcal{Q}^*. \end{aligned} \quad (3.3a)$$

As the solutions to both variational problems are the robot's Euler–Lagrange equations of motion, eq. 3.3a induces a symmetry relation between inertial and moving forces at symmetry-related states (e.g., see fig. 3.1). To see this, note that by the linearity and invertibility of group actions (definition 2.2) and the chain rule, the partial derivatives of the Lagrangian with respect to variations in q satisfy

$$\begin{aligned} L(q, \dot{q}) &= L(g \triangleright_Q q, g \triangleright_{T_q Q} \dot{q}), \\ \frac{\partial L(q, \dot{q})}{\partial q} &= \frac{\partial L(g \triangleright_Q q, g \triangleright_{T_q Q} \dot{q})}{\partial q} \\ \frac{\partial L(q, \dot{q})}{\partial q} &= \left(\frac{\partial(g \triangleright_Q q)}{\partial q} \right)^\top \frac{\partial L(g \triangleright_Q q, g \triangleright_{T_q Q} \dot{q})}{\partial g \triangleright_Q q} \\ g \triangleright_{T_q Q} \left(\frac{\partial L(q, \dot{q})}{\partial q} \right) &= \frac{\partial L(g \triangleright_Q q, g \triangleright_{T_q Q} \dot{q})}{\partial g \triangleright_Q q}, \quad \text{by eq. 3.1.} \end{aligned} \quad (3.3b)$$

Likewise, the partial derivatives of the Lagrangian with respect to variations in $\dot{\mathbf{q}}$ satisfy

$$\begin{aligned}
\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} &= \frac{\partial L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} \\
\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} &= \left(\frac{\partial (g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} \right)^\top \frac{\partial L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}} \\
\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} &= \left(\frac{\partial (g \triangleright_{\mathcal{Q}} \mathbf{q})}{\partial \mathbf{q}} \right)^\top \frac{\partial L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}} \quad \text{by eq. 3.1} \\
g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} \left(\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} \right) &= \frac{\partial L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}}, \quad \text{by eq. 3.1.}
\end{aligned} \tag{3.3c}$$

Then, substituting eqs. 3.3b and 3.3c into eq. 3.3a and noting that symmetry transformations are time-invariant, we derive the equivariance constraints on the Euler–Lagrange equations at symmetry-related states, which hold for the entire symmetry group:

$$g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} \left(\frac{d}{dt} \frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}} - \frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \mathbf{q}} \right) = \frac{d}{dt} \frac{\partial L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial (g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})} - \frac{\partial L(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}{\partial (g \triangleright_{\mathcal{Q}} \mathbf{q})} \quad \forall g \in \mathbb{G}, \quad (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}. \tag{3.3d}$$

Recalling from section 2.2 that the robot’s equations of motion are derived from the Euler–Lagrange equations, with generalized inertial forces arising from $\frac{d}{dt} \frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \dot{\mathbf{q}}}$ and generalized moving forces arising from $\frac{\partial L(\mathbf{q}, \dot{\mathbf{q}})}{\partial \mathbf{q}}$, we can rewrite eq. 3.3d as \mathbb{G} -equivariance constraints on the robot’s equations of motion at symmetry-related states:

$$\begin{aligned}
g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} (\mathbf{M}(\mathbf{q}) \ddot{\mathbf{q}} - \boldsymbol{\tau}(\mathbf{q}, \dot{\mathbf{q}})) &= \mathbf{M}(g \triangleright_{\mathcal{Q}} \mathbf{q}) g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} \ddot{\mathbf{q}} - \boldsymbol{\tau}(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}) \quad \forall g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}. \\
\underbrace{(g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} \mathbf{M}(\mathbf{q}) - \mathbf{M}(g \triangleright_{\mathcal{Q}} \mathbf{q}) \tau_{q \mathcal{Q}^* \triangleleft g})}_{\text{Equivariance constraint on mass matrix}} \ddot{\mathbf{q}} &= \underbrace{g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} \boldsymbol{\tau}(\mathbf{q}, \dot{\mathbf{q}}) - \boldsymbol{\tau}(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})}_{\text{Equivariance constraint on generalized forces}}
\end{aligned} \tag{3.3e}$$

Here, $\boldsymbol{\tau}(\mathbf{q}, \dot{\mathbf{q}}) := \boldsymbol{\tau}_{\text{grav}}(\mathbf{q}) + \boldsymbol{\tau}_{\text{nonlin}}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}_q \mathcal{Q}^*$ denotes the generalized force vector assumed by the principle of least action, composed of gravitational and nonlinear terms (see eq. 2.9).

As already hinted, the core implication of the dynamical equivalence between two states—formalized by eq. 3.3e—is that both states exhibit identical temporal evolution when subjected to equivalent forcing (see fig. 3.1). Consequently, knowing how to model and control the system along a given trajectory immediately provides knowledge of how to model and control the system along any symmetry-related trajectory. Since this form of knowledge transfer constitutes the key prior we aim to leverage, developing clear intuition around it is essential.

If eq. 3.3e is required to hold for arbitrary generalized accelerations $\ddot{\mathbf{q}} \in \mathcal{T}_q \mathcal{Q}^*$, then the robot’s mass matrix must satisfy the following matrix-valued equivariance constraint (see definition 2.15):

$$g \triangleright_{\mathcal{T}_q \mathcal{Q}^*} \mathbf{M}(\mathbf{q}) \tau_{q \mathcal{Q}^* \triangleleft g}^{-1} = \mathbf{M}(g \triangleright_{\mathcal{Q}} \mathbf{q}), \quad \forall g \in \mathbb{G}, \mathbf{q} \in \mathcal{Q}, \tag{3.4}$$

which endows the mass matrix with rich, symmetry-group-dependent algebraic structure, well characterized for compact groups originally by the Wigner–Eckart theorem in quantum mechanics and in Lang and Weiler (2020) in the context of GDL. Violating this constraint for any $g \in \mathbb{G}$ implies that g does not define a symmetry of the robot’s dynamics. For example, the symmetry in the temporal evolution of the double pendulum shown in fig. 3.1 is broken when the links exhibit an asymmetric mass distribution, as illustrated in fig. 3.3-(a-b).




Figure 3.3 **Left:** Symmetry breaking in the double pendulum due to disturbances in the mass distribution (black box on second link), resulting in the violation eq. 3.4, and causing states (a) and (b) to no longer exhibit equivalent instantaneous dynamics. **Right:** Introduction of an external disturbance force $f \in \mathbb{R}^2$ leading to the violation eq. 3.5, thereby breaking the equivariance relation between the symmetry-related states (q_0, \dot{q}_0) in (d) and $(g_r \triangleright_{\mathcal{Q}} q_0, g_r \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}} \dot{q}_0)$ in (e). Note, however, that under the assumption that inertial and moving disturbances and their reflected counterparts are equally likely, each disturbed state $(g_r \triangleright_{\mathcal{Q}} q_0, g_r \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}} \dot{q}_0)$ in (b) and (e) admits a symmetric equivalent disturbed state, shown in (c) and (f), respectively.

In contrast to eq. 3.4, which encodes a constraint on the system’s morphology, the constraint on generalized moving forces required for eq. 3.3e to hold,

$$g \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}^*} \tau(\mathbf{q}, \dot{\mathbf{q}}) = \tau(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}} \dot{\mathbf{q}}), \quad \forall g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}, \quad (3.5)$$

simply states that dynamically equivalent states will evolve equivalently only when subjected to equivalent—or symmetry-related—forcing. This condition applies even for general force terms, including control inputs, contact forces, and non-conservative effects (see eq. 2.5). For instance, in the double pendulum example of fig. 3.1, applying an external disturbance to only one of two dynamically equivalent states results in distinct temporal evolutions, as shown in fig. 3.3-(d-e).

Crucially, in real-world robotic systems, both inertial and moving-force disturbances are unavoidable. Our focus is therefore not on idealized symmetry-preserving scenarios (fig. 3.1), but on practical settings in which a broad distribution of disturbances acts on the system at every time step. The central objective is thus to model and control system dynamics under such disturbance distributions. In this regime, the relevant symmetry relations are those between disturbed states, as illustrated in fig. 3.3-(b,c) and (e,f), since they enable the transfer of modelling and control knowledge from a system under a given perturbation to its symmetry-related state under the corresponding symmetric perturbation. This prior is of high practical relevance when symmetric disturbances are likely to be encountered during operation, and it applies to both simple and complex robotic systems, as depicted in figs. 3.5 and 3.6. This observation is formalized in the following corollary, which generalizes eq. 3.3e to arbitrary generalized force fields.

Corollary 3.1 (Equivariant equations of motion). *Let $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}$ be the state of a robotic system with symmetry group \mathbb{G} (definition 3.1), and let $\ddot{\mathbf{q}}(\mathbf{q}, \dot{\mathbf{q}}) := \mathbf{M}(\mathbf{q})^{-1} \tau(\mathbf{q}, \dot{\mathbf{q}})$ denote the system’s acceleration at state $(\mathbf{q}, \dot{\mathbf{q}})$, induced by an arbitrary generalized moving force vector $\tau(\mathbf{q}, \dot{\mathbf{q}})$ (eq. 2.5). Then, the instantaneous acceleration at any symmetry-related state is related to the acceleration at $(\mathbf{q}, \dot{\mathbf{q}})$ through the following equivariance constraint:*

$$\begin{aligned} g \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}^*} \ddot{\mathbf{q}}(\mathbf{q}, \dot{\mathbf{q}}) &= \ddot{\mathbf{q}}(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}} \dot{\mathbf{q}}), \quad \forall g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}. \\ g \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}^*} (\mathbf{M}(\mathbf{q})^{-1} \tau(\mathbf{q}, \dot{\mathbf{q}})) &= \mathbf{M}(g \triangleright_{\mathcal{Q}} \mathbf{q})^{-1} g \triangleright_{\mathcal{T}_{q_0} \mathcal{Q}^*} \tau(\mathbf{q}, \dot{\mathbf{q}}) \end{aligned} \quad (3.6)$$

3.2 Taxonomy of symmetries in robotics

The structure and implications of the symmetry constraints imposed on the system's equations of motion (eq. 3.3e) depend critically on the type of symmetry under consideration. This subsection provides a brief taxonomy of symmetries in robotics, illustrated in fig. 1.2, and thereby sets the stage for the introduction of morphological symmetries (chapter 4), which constitute the main focus of this thesis.

Symmetries in classical mechanics and robotics can be broadly classified into two categories: *spatial symmetries* and *temporal symmetries*. Spatial symmetries are transformations of the system's state space that relate states with equivalent dynamics (eq. 3.2), whose implications were analyzed in the previous section. Temporal symmetries, by contrast, are transformations of time itself—namely time translations and time reflections—which also relate states with equivalent dynamics, but whose analysis and implications require a distinct, though conceptually analogous, development. Since practitioners often conflate temporal and spatial symmetries, we briefly discuss temporal symmetries to provide a complete context and clearly differentiate between these two classes of symmetry priors.

3.2.1 Temporal symmetries

Classical (Newtonian) mechanics admits two fundamental temporal symmetry priors: time translation and time reflection. These symmetries arise because the dynamics of mechanical systems are, by construction, invariant under transformations of time. Intuitively, this means that the dynamics governing a robotic system today are the same as those governing that same system tomorrow or yesterday. Moreover, since Newton's laws of motion are time-independent and describe deterministic dynamics, the same equations remain valid under a reversal of the time variable, at least at the level of idealized mechanical models.

Accordingly, in sections 2.1 and 3.1 we incorporated this prior by assuming that the Lagrangian of our dynamical systems does not depend on time. In full generality, however, one should consider a time-dependent Lagrangian $L : \mathcal{TQ} \times \mathbb{T} \mapsto \mathbb{R}$, where $\mathbb{T} \subseteq \mathbb{R}$ denotes the time domain. Then, analogously to the formal definition of spatial symmetries (definition 3.1), temporal symmetries can be defined through the invariance of the Lagrangian under transformations of time:

$$L(\mathbf{q}, \dot{\mathbf{q}}, t) = L(\mathbf{q}, h_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}}, h_{\triangleright_{\mathbb{T}}} t), \quad \forall h \in \mathbb{H}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}, t \in \mathbb{T}. \quad (3.7)$$

Here, \mathbb{H} denotes the temporal symmetry group, $(\triangleright_{\mathbb{T}}) : \mathbb{H} \times \mathbb{T} \mapsto \mathbb{T}$ the group action of a temporal symmetry on time, and $(\triangleright_{\mathcal{T}_q \mathcal{Q}}) : \mathbb{H} \times \mathcal{T}_q \mathcal{Q} \mapsto \mathcal{T}_q \mathcal{Q}$ the induced action on the system's tangent space at any configuration $\mathbf{q} \in \mathcal{Q}$. In classical mechanics, the temporal symmetry group is given by $\mathbb{H} := \mathbb{R} \times \mathbb{C}_2$. Here, $\mathbb{R} = \{h_{\Delta t} \mid \Delta t \in \mathbb{R}\}$ denotes the group of time translations, acting as $h_{\Delta t} \triangleright_{\mathbb{T}} t = t + \Delta t$ and leaving velocities invariant, i.e., $h_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}} = \dot{\mathbf{q}}$. The factor $\mathbb{C}_2 = \{e, h_r \mid h_r^2 = e\}$ denotes the time-reflection group, acting via $h_r \triangleright_{\mathbb{T}} t = -t$ and reversing velocities, $h_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}} = -\dot{\mathbf{q}}$.

The invariance of the dynamics under time translations is commonly referred to as *time homogeneity* and underlies the Markov assumption for the robot's discretized dynamics adopted in most optimal-control and reinforcement-learning formulations. Since these priors are standard in robotics modelling, we do not incorporate them explicitly in our analysis, for the sake of notational simplicity.

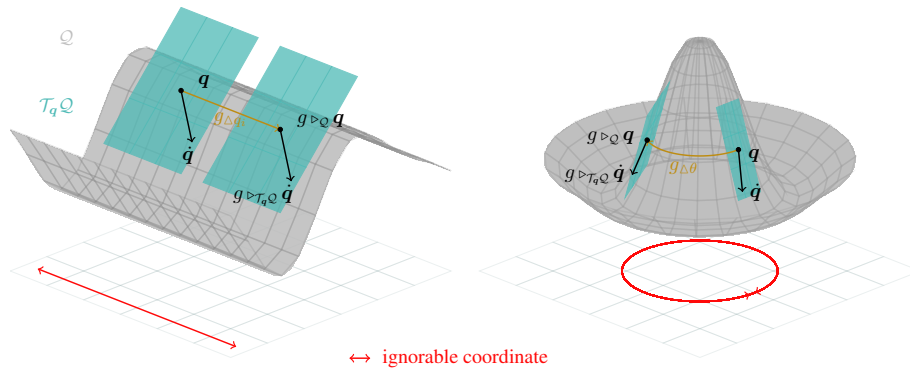


Figure 3.4 Illustration of the geometric constraints induced on the configuration manifold \mathcal{Q} by ignorable coordinates. **Left:** a translational ignorable coordinate (red axis) constrains the curvature of \mathcal{Q} to remain constant under translations and reflections along that coordinate. **Right:** a rotational ignorable coordinate constrains the curvature of \mathcal{Q} to remain constant under rotations and reflections about the corresponding axis (vertical rotation axis).

3.2.2 Spatial symmetries

In robotics, spatial symmetries are transformations acting on the robot’s configuration manifold \mathcal{Q} , and—via eq. 3.1—on the associated tangent $\mathcal{T}_q \mathcal{Q}$ and cotangent spaces $\mathcal{T}_q^* \mathcal{Q}$ at each configuration $q \in \mathcal{Q}$. We classify spatial symmetries into two categories: *environmental symmetries* and *morphological symmetries*. Both are induced by the spatial symmetry group underlying Newtonian mechanics, namely the Euclidean group \mathbb{E}_3 of isometries comprising rigid translations, rotations, and reflections of space. Crucially, whereas environmental symmetries are present in every mechanical system, morphological symmetries are robot-specific and arise from intrinsic regularities in the robot’s morphology.

Environmental/base symmetries Environmental symmetries, often referred to as *base symmetries* in the context of floating-base robots (Bloch et al., 1996), arise from transformations of the *ignorable variables* in classical mechanics. These correspond to transformations of the chosen inertial reference frame (see fig. 1.1), which induce offsets in the x , y , and z position coordinates and in the roll, pitch, and yaw orientation coordinates of the robot and the objects in its environment. As recalled in section 2.1, an ignorable variable is a generalized coordinate q_i that does not appear explicitly in the system’s Lagrangian (see eq. 2.10). Consequently, the system dynamics are independent of the absolute value of q_i and depend only on its velocity \dot{q}_i and on the remaining generalized coordinates (Bloch et al., 1996; Murray et al., 2017; Wieber, 2006). This is a universal property of mechanical systems governed by classical mechanics (see fig. 1.1).

Exploiting environmental symmetries in data-driven methods is both straightforward and common practice. In dynamics modelling, environmental symmetries are typically leveraged by avoiding the direct prediction of *absolute* positions and orientations. Instead, models are parameterized to predict relative positions and orientations, or directly velocities, while global coordinates are recovered by integration from an initial condition (Ghaffari et al., 2022). Similarly, in optimal control, data-driven policies should avoid dependence on absolute coordinates and instead rely on velocities and relative measurements, such as the relative pose between the robot and a target object rather than the object’s global position. Ignoring the environmental symmetries of classical mechanics in data-driven modelling and control is harmful and frequently found in practice. If parametric models of dynamics




Figure 3.5 **Left:** The Atlas humanoid robot exhibits a bilateral reflection symmetry group, $\mathbb{C}_2 = \{e, g_s \mid g_s^2 = e\}$, where g_s denotes sagittal reflection. **Right:** The robot’s configuration manifold inherits this symmetry, imposing geometric constraints on its curvature under high-dimensional reflections g_s .

or control policies are made directly dependable of global coordinates, the models will be biased toward the range of global coordinate values present in the training distribution, and will fail to generalize under trivial translations or rotations.

In geometric terms, the configuration manifold \mathcal{Q} is constrained to have constant curvature along transformations of the ignorable position and orientation coordinates (see fig. 3.4). Ignorable absolute positions—i.e., global x , y , or z —induce invariance under translation and reflection groups isomorphic to $\mathbb{Z} \times \mathbb{C}_2$, see fig. 3.4-left. Likewise, ignorable absolute orientations—i.e., roll, pitch, or yaw—induce invariance under rotation and reflection groups isomorphic to \mathbb{O}_2 , as illustrated in fig. 3.4-right.

Morphological symmetries This thesis focuses on the second class of spatial symmetries: *morphological symmetries* (see chapter 4). These symmetries induce richer geometric constraints on the configuration manifold, the system’s mass matrix, and optimal control policies. Morphological symmetries arise from a mechanical system’s ability to reconfigure itself to emulate a Euclidean isometry. For example, a double pendulum can be reconfigured (see fig. 3.1) to achieve a state with dynamics equivalent to a true spatial reflection. The next chapter provides a formal definition and detailed analysis of morphological symmetries, as leveraging them in robot learning requires a more intricate approach. Before proceeding, we offer a high-level intuition for the geometric differences between morphological and environmental symmetries.

Unlike environmental symmetries, morphological symmetry groups are finite and impose nontrivial geometric constraints on the configuration manifold of mechanical and robotic systems. For instance, in a robot with bilateral symmetry—where the left side is a reflection of the right—this reflection constrains the configuration manifold, introducing a high-dimensional reflection plane that relates the curvature on one side to that on the other (see fig. 3.5). More generally, robots with larger morphological symmetry groups, such as the quadruped Solo or Mini-Cheetah, have configuration manifolds whose curvature is similarly constrained by high-dimensional symmetry transformations (see fig. 3.6).

As discussed in chapters 4, 6 and 7, these symmetries also transfer to the observation and action spaces of the **Markov Decision Processes (MDPs)** that define the relevant control problems in robotics. Rendering relevant inductive priors for robot modelling, control, and estimation.




Figure 3.6 **Left:** The Solo quadruped robot possesses a morphological symmetry group of order 4, $\mathbb{K}_4 = \{e, g_s, g_t, g_r \mid g_s^2 = g_t^2 = g_r^2 = e, g_s \circ g_t = g_r\}$, where g_s and g_t denote sagittal and transverse reflections, and g_r denotes a 180° rotation. **Right:** The configuration manifold inherits this symmetry group, resulting in geometric constraints on its curvature under the high-dimensional actions of g_s , g_t , and g_r .

Summary 3.2.2

For any mechanical/robotic system, a symmetry is defined as an energy preserving transformation relating two system's states with equivalent dynamics (corollary 3.1). Meaning that both states have equivalent temporal evolution when subjected to equivalent forcing. This is a valuable prior for modelling and control as it allows to transfer knowledge from one state to its symmetry-related states.

Symmetries in robotics can be classified into two categories: temporal symmetries, which are transformations of time itself, and spatial symmetries, which are transformations of the robot's configuration manifold. Spatial symmetries can be further classified into environmental symmetries, which arise from ignorable coordinates (section 2.2) and morphological symmetries, which arise from regularities in the robot's morphology.

Chapter 4

Morphological Symmetries (MSs)

Morphological symmetries capture a robot’s ability to emulate Euclidean isometries—such as reflections, rotations, or translations—through *feasible* reconfigurations of its body. In particular, a robot can reach a dynamically equivalent state corresponding to the application of a Euclidean isometry to the robot and its environment, even when that isometry itself would produce an infeasible physical configuration.

To formalize this distinction, we separate the standard action of Euclidean isometries—here referred to as the environmental group action—from the morphological group action. The environmental action corresponds to the canonical action of \mathbb{E}_3 on the global world-state space Ω , acting jointly on the robot and its environment, with the robot’s configuration manifold \mathcal{Q} being a submanifold of Ω . In contrast, the morphological group action is defined only for a subgroup $\mathbb{G} \leq \mathbb{E}_3$ and acts directly on the robot’s configuration manifold through feasible state transformations:

$$\underbrace{\mathbb{P}_\Omega : \mathbb{E}_3 \times \Omega \mapsto \Omega}_{\text{Environmental action}} \quad \text{vs.} \quad \underbrace{\triangleright_{\mathcal{Q}} : \mathbb{G} \times \mathcal{Q} \mapsto \mathcal{Q}}_{\text{Morphological action}}. \quad (4.1)$$

For a given state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}$, we denote the resulting states under these actions by $(g^{\mathbb{P}_\Omega} \mathbf{q}, g^{\mathbb{P}_\Omega} \dot{\mathbf{q}})$ and $(g^{\triangleright_{\mathcal{Q}}} \mathbf{q}, g^{\triangleright_{\mathcal{Q}}} \dot{\mathbf{q}})$, respectively (see eq. 3.1 and fig. 4.1). While environmental actions may produce dynamically equivalent but physically unreachable robot configurations, morphological actions produce feasible states with equivalent dynamics through admissible reconfigurations of the robot’s body and joints.

Any humanoid robot provides the simplest example (see fig. 4.1-left): its bilateral body symmetry induces a morphological reflection group $\mathbb{G} = \mathbb{C}_2$ (see fig. 2.1). Although true spatial reflections are infeasible for any robot, humanoids can mimic reflections via the corresponding morphological action—reorienting the base, hips, and head; and exchanging the left and right limb configurations—thereby yielding a feasible state with equivalent dynamics. A similar phenomenon can arise for other Euclidean isometries, including multiple reflections, rotations, and translations, when robots exhibit larger symmetry groups. This is the case for some quadruped robots (see fig. 4.1-center-&-right), where sagittal and transverse body symmetries induce a Klein-four morphological symmetry group $\mathbb{G} = \mathbb{K}_4$. In this case, some Euclidean isometries—namely rotations—correspond to both feasible environmental and morphological actions, whereas reflections are feasible only morphologically.

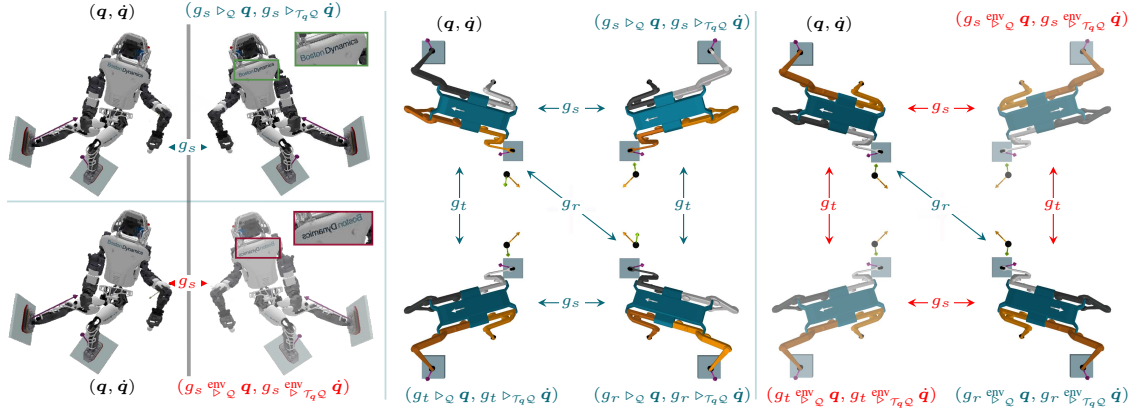


Figure 4.1 Distinction between environmental and morphological group actions (eq. 4.1), which define inequivalent realizations of the same Euclidean isometries. **Left:** Atlas humanoid robot with a bilateral sagittal reflection symmetry group $\mathbb{G} = \mathbb{C}_2 = \{e, g_s \mid g_s^2 = e\}$. The reflection g_s admits two realizations: (i) an environmental symmetry action, corresponding to a true spatial reflection of the robot and environment, which produces the unreachable state $(g_s^{\text{env}} \circledast \mathbf{q}, g_s^{\text{env}} \circledast \tau_{\mathbf{q}} \dot{\mathbf{q}})$, given that rigid bodies cannot undergo reflections; and (ii) a morphological symmetry action, implemented as a feasible robot reconfiguration that reorients the base and exchanges left and right limbs, yielding a reachable state with equivalent dynamics. The robot chest text orientation highlights the distinction. **Center-&Right:** Solo quadruped robot with a Klein-four morphological symmetry group $\mathbb{K}_4 = \{e, g_s, g_t, g_r \mid g_r^2 = g_s^2 = g_t^2 = e, g_s \circ g_t = g_r\}$, comprising sagittal and transverse reflections (g_s, g_t) and a 180° rotation (g_r). Environmental actions implement these isometries as true spatial reflections or rotations acting on both robot and environment (right), while morphological actions realize them as feasible robot state transformations that reorient the base and permute limb roles (center). Limb color labels and forward arrows indicate the induced permutations.

In both examples, for every Euclidean isometry $g \in \mathbb{G} \subset \mathbb{E}_3$, the environmental symmetry action transforms both the robot and its environment, producing a system state $(g^{\text{env}} \circledast \mathbf{q}, g^{\text{env}} \circledast \tau_{\mathbf{q}} \dot{\mathbf{q}})$ that may be infeasible. In contrast, a morphological symmetry acts directly on the robot state, yielding a *feasible* state $(g \circledast \mathbf{q}, g \circledast \tau_{\mathbf{q}} \dot{\mathbf{q}})$ that is dynamically equivalent to $(\mathbf{q}, \dot{\mathbf{q}})$. Consequently, knowledge about modelling and controlling the system at $(\mathbf{q}, \dot{\mathbf{q}})$ transfers directly to any morphologically symmetric state $(g \circledast \mathbf{q}, g \circledast \tau_{\mathbf{q}} \dot{\mathbf{q}})$ for all $g \in \mathbb{G}$. Crucially, whereas this transfer is straightforward for environmental symmetry transformations (see section 3.2), exploiting it for morphological symmetry transformations requires a more elaborate analysis and implementation, which we develop in the rest of this chapter.

Having built intuition through the examples in fig. 4.1, we now enumerate the characteristic properties of morphological symmetries and then provide a formal definition that captures them in full generality for any robotic system. Morphological symmetries have the following defining characteristics:

- MS1: Feasibility:** the action of a morphological symmetry transformation on a robot state $(\mathbf{q}, \dot{\mathbf{q}})$ yields a feasible state $(g \circledast \mathbf{q}, g \circledast \tau_{\mathbf{q}} \dot{\mathbf{q}}) \in \mathcal{TQ}$, which is dynamically equivalent to $(\mathbf{q}, \dot{\mathbf{q}})$.
- MS2: Non-triviality in joint-space:** a morphological symmetry transformation entails a non-trivial transformation of the robot's joint-space configuration, required to imitate the action of the corresponding Euclidean isometry on the robot's morphology (see fig. 4.1).

These properties can be formalized for any robotic system as follows:

Definition 4.1 (Morphological Symmetry). Let $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$ denote the state of a robotic system with Lagrangian function $L : \mathcal{TQ} \mapsto \mathbb{R}$. A Euclidean isometry $g \in \mathbb{E}_3$ is deemed a morphological symmetry of the system if it induces two inequivalent group actions yielding distinct states $(g \circledast \mathbf{q},$

$g_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}}$) and $(g_{\mathbb{F}_{\mathcal{Q}}} \mathbf{q}, g_{\mathbb{F}_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}})$, with $(g_{\triangleright_{\mathcal{Q}}} \mathbf{q}, g_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}})$ being a feasible state dynamically equivalent to $(\mathbf{q}, \dot{\mathbf{q}})$, i.e.:

$$\begin{aligned} L(\mathbf{q}, \dot{\mathbf{q}}) &= L(g_{\triangleright_{\mathcal{Q}}} \mathbf{q}, g_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}}) = L(g_{\mathbb{F}_{\mathcal{Q}}} \mathbf{q}, g_{\mathbb{F}_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}}), & \forall (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}. \\ \text{s.t. } & (g_{\mathbb{F}_{\mathcal{Q}}} \mathbf{q}, g_{\mathbb{F}_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}}) \neq (g_{\triangleright_{\mathcal{Q}}} \mathbf{q}, g_{\triangleright_{\mathcal{T}_q \mathcal{Q}}} \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}, \end{aligned}$$

Nomenclature: The set of all morphological symmetries of a robotic system forms its morphological symmetry group $\mathbb{G} < \mathbb{E}_3$.

Because morphological symmetries are spatial symmetries, they induce \mathbb{G} -equivariance of the system's equations of motion (see eq. 3.3e), impose non-trivial structural constraints on the mass matrix (see eq. 3.4), and provide relevant symmetry priors for data-driven dynamics modelling (chapter 6), optimal control (chapter 8), and estimation (chapter 7).

Next, we focus on floating-base rigid-body systems. For this class of robots, we develop a kinematic and dynamic analysis of how morphological-symmetry constraints arise in practice, with particular emphasis on the induced algebraic structure of the mass matrix and its implications for numerical physics simulation.

This analysis is primarily intended to support the identification of morphological symmetries in floating-base robots and to clarify the resulting dynamical constraints. Readers primarily interested in the implications for data-driven modelling, control, and estimation may jump directly to chapter 6, chapter 8, and chapter 7. Morphological symmetries can also arise in other robotic systems, including soft, continuum, and modular robots; extending this analysis to those settings is relatively straightforward, and their applications remain an interesting direction for future research.

The case of floating-base rigid-body systems

In this section, we characterize the constraints that rigid-body systems must satisfy to admit a morphological symmetry group $\mathbb{G} < \mathbb{E}_3$. Specifically, for robotic systems consisting of n_b interconnected rigid bodies arranged in a kinematic tree, we study how the presence of a morphological symmetry implies one or both of the following properties: (i) permutation symmetries of kinematic-tree chains, and (ii) symmetry constraints on the mass distributions of unique bodies and of bodies across symmetry-related kinematic chains. Both properties can be expressed as algebraic constraints on the inertial and kinematic parameters that define the robot's morphology, which are precisely the core parameters appearing in robot-description files such as URDFs and SDFs (Carpentier et al., 2019; Todorov et al., 2012). By understanding these parameter constraints, we aim to clarify how to

1. Define morphological symmetry group actions/transformations in practice through linear algebraic operations.
2. Identify morphological symmetries algorithmically from robot description files.
3. Establish the core constraints needed for computational design of robot morphologies with arbitrary symmetry groups, thus providing the path to generalize the Universal Animal design space (UniMal) (Gupta et al., 2021) to consider larger symmetry groups than \mathbb{C}_2 for the design of robot morphologies.

Configuration manifold of floating-base systems For generality, we consider floating-base robotic systems composed of n_b rigid bodies, including a base body (e.g., the torso body in the humanoid and quadruped examples of fig. 4.1) and $n_b - 1$ additional bodies (e.g., limbs), interconnected by n_j joints. For this class of systems, the configuration manifold decomposes into the floating-base configuration manifold \mathbb{SE}_3 —which is a Lie group—and the *joint-space* configuration manifold $\mathcal{M} \subseteq \mathbb{R}^{n_{js}}$, which encodes the positions of all joint degrees of freedom. That is, $\mathcal{Q} = \mathbb{SE}_3 \times \mathcal{M} \subset \mathbb{R}^{n_q}$ with $n_q = 6 + n_{js}$ and $n_{js} = \sum_{i=1}^{n_j} \text{number of } DoF \text{ of joint } i$ (refer to section 2.1 for details). Hence, we will often denote the decomposition of the generalized coordinates by

$$\mathbf{q} = \underbrace{(\mathbf{X}_B, \mathbf{q}_{js})}_{\begin{bmatrix} \mathbf{R}_B & \mathbf{r}_B \\ \mathbf{0} & \mathbf{1} \end{bmatrix}} \in \mathcal{Q} := \mathbb{SE}_3 \times \mathcal{M}, \quad \text{and} \quad \dot{\mathbf{q}} = \underbrace{(\dot{\mathbf{x}}_B, \dot{\mathbf{q}}_{js})}_{\begin{bmatrix} \dot{\mathbf{r}}_B \\ \boldsymbol{\omega}_B \end{bmatrix}} \in \mathcal{T}_q \mathcal{Q} := \mathfrak{se}_3 \times \mathcal{T}_q \mathcal{M}, \quad (4.2)$$

where $\mathbf{X}_B \in \mathbb{SE}_3$ denotes the base body configuration—represented as a 4×4 homogeneous matrix defined by the base position $\mathbf{r}_B \in \mathbb{R}^3$ and rotation matrix representation $\mathbf{R}_B \in \mathbb{SO}_3$ (see eq. 2.1)—and $\mathbf{q}_{js} \in \mathcal{M}$ denotes the joint-space configuration. Likewise, $\dot{\mathbf{x}}_B \in \mathfrak{se}_3$ denotes the base-body linear and angular velocity—represented as a 6×1 spatial velocity vector defined by the base linear velocity $\dot{\mathbf{r}}_B \in \mathbb{R}^3$ and angular velocity $\boldsymbol{\omega}_B \in \mathbb{R}^3$ (see eq. 2.2)—and $\dot{\mathbf{q}}_{js} \in \mathcal{T}_q \mathcal{M}$ denotes the joint-space generalized velocity coordinates.

Note that although $|\mathbb{SE}_3| = |\mathfrak{se}_3| = 6$, we rely on the 4×4 matrix-group representation of body configurations and 6×1 spatial velocity vectors to align with the notation used in the robot-dynamics simulation literature (Carpentier et al., 2019; Featherstone, 2008; Mastalli et al., 2020a). One may choose a different orientation parameterization, such as quaternions or Euler angles, but the underlying nature of the symmetry transformations on these objects remains unchanged.

Morphological symmetry group actions for floating-base systems In this context, the morphological and environmental group actions (see eq. 4.1), acting on the configuration manifold \mathcal{Q} of floating-base robotic systems, are defined as follows:

$$\begin{aligned} g \triangleright_{\mathcal{Q}} \mathbf{q} &:= (g \diamond_{\mathbb{E}_3} \mathbf{X}_B, g \triangleright_{\mathcal{M}} \mathbf{q}_{js}) & \text{and} & \quad g^{\text{env}}_{\mathbb{Q}} \mathbf{q} := (g \triangleright_{\mathbb{E}_3} \mathbf{X}_B, \mathbf{q}_{js}). \\ &= \left(\underbrace{\rho_{\mathbb{E}_3}(g) \mathbf{X}_B \rho_{\mathbb{E}_3}(g^{-1})}_{\begin{bmatrix} \mathbf{R}_g & \mathbf{r}_g \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{R}_B & \mathbf{r}_B \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{R}_g & \mathbf{r}_g \\ \mathbf{0} & \mathbf{1} \end{bmatrix}^{-1}}, \rho_{\mathcal{M}}(g) \mathbf{q}_{js} \right) & = \left(\underbrace{\rho_{\mathbb{E}_3}(g) \mathbf{X}_B}_{\begin{bmatrix} \mathbf{R}_g & \mathbf{r}_g \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \begin{bmatrix} \mathbf{R}_B & \mathbf{r}_B \\ \mathbf{0} & \mathbf{1} \end{bmatrix}}, \mathbf{q}_{js} \right). \end{aligned} \quad (4.3a)$$

The environmental action (${}^{\text{env}}_{\mathbb{Q}}$) is defined by the standard action of Euclidean isometries on rigid bodies, which simply rotates or reflects and translates the base-body configuration in space. Algebraically, this action is represented simply by composition of homogeneous transformation matrices:

$$\begin{aligned} (\triangleright_{\mathbb{E}_3}) : \mathbb{G} \times \mathbb{E}_3 &\mapsto \mathbb{E}_3, \\ (g, \mathbf{X}) &\mapsto \rho_{\mathbb{E}_3}(g) \mathbf{X}. \end{aligned} \quad (4.3b)$$

where $\rho_{\mathbb{E}_3}(g) \in \mathbb{R}^{4 \times 4}$ denotes the (*proper or improper*) homogeneous-transformation-matrix representation of the Euclidean isometry $g \in \mathbb{E}_3$, composed of a *roto-reflection* $\mathbf{R}_g \in \mathbb{O}_3$ and a translation $\mathbf{r}_g \in \mathbb{R}^3$ (see eq. 4.3a). In contrast, the morphological group action ($\triangleright_{\mathcal{Q}}$) : $\mathbb{G} \times \mathcal{Q} \mapsto \mathcal{Q}$ is defined by the standard *conjugation action* on \mathbb{E}_3 , ($\diamond_{\mathbb{E}_3}$), and by a group action on the joint-space configuration

manifold, $(\triangleright_{\mathcal{M}})$, defined as:

$$\begin{aligned} (\diamond_{\mathbb{E}_3}) : \mathbb{G} \times \mathbb{E}_3 &\mapsto \mathbb{E}_3, & \triangleright_{\mathcal{M}} : \mathbb{G} \times \mathcal{M} &\mapsto \mathcal{M}, \\ (g, \mathbf{X}) &\mapsto \boldsymbol{\rho}_{\mathbb{E}_3}(g) \mathbf{X} \boldsymbol{\rho}_{\mathbb{E}_3}(g^{-1}), & (g, \mathbf{q}_{\text{js}}) &\mapsto \boldsymbol{\rho}_{\mathcal{M}}(g) \mathbf{q}_{\text{js}}. \end{aligned} \quad (4.3c)$$

Here, $\boldsymbol{\rho}_{\mathcal{M}}(g) \in \mathbb{R}^{n_{\text{js}} \times n_{\text{js}}}$ defines the transformation of the joint-space configuration required to imitate the corresponding Euclidean isometry on the robot's morphology. In the context of the examples in fig. 4.1, $(\diamond_{\mathbb{E}_3})$ is responsible for the reorientation and relocation of the robot's base, and $(\triangleright_{\mathcal{M}})$ is responsible for the permutation of limb roles.

It is important to note that $(\mathbb{E}_{\mathcal{Q}})$ is not a well-defined group action (definition 2.2) on \mathcal{Q} because the resulting state may lie outside the configuration manifold, i.e., $(g \mathbb{E}_{\mathcal{Q}} \mathbf{q}, g \mathbb{E}_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}) \notin \mathcal{T}\mathcal{Q}$. This occurs when the Euclidean isometry g is a reflection, such as $g_s, g_t \in \mathbb{E}_3$ in fig. 4.1, since for any $\mathbf{R}_B \in \mathbb{S}\mathbb{O}_3$, the result of $g \triangleright_{\mathbb{E}_3} \mathbf{R}_B$ is a *roto-reflection*, implying that $g \triangleright_{\mathbb{E}_3} \mathbf{R}_B \notin \mathbb{S}\mathbb{O}_3$. This is the algebraic manifestation of the physical fact that rigid bodies cannot undergo reflections, depicted in fig. 4.1 by rendering infeasible states translucent. In contrast, note that the morphological group action always yields a feasible base configuration, since for any $\mathbf{R}_B \in \mathbb{S}\mathbb{O}_3$ and $g \in \mathbb{G}$, we have $g \diamond_{\mathbb{E}_3} \mathbf{R}_B \in \mathbb{S}\mathbb{O}_3$.

Then, following eq. 3.1, we can derive the morphological and environmental group actions on the tangent space $\mathcal{T}_q \mathcal{Q}$ of floating-base robotic systems as follows:

$$\begin{aligned} g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} &:= (g \triangleright_{\mathbb{E}_3} \dot{\mathbf{x}}_B, \boldsymbol{\rho}_{\mathcal{T}_q \mathcal{M}}(g) \dot{\mathbf{q}}_{\text{js}}) & \text{and} & & g \mathbb{E}_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} &:= (g \triangleright_{\mathbb{E}_3} \dot{\mathbf{x}}_B, \dot{\mathbf{q}}_{\text{js}}), \\ &= (\underbrace{\boldsymbol{\rho}_{\mathbb{E}_3}(g)}_{\mathbf{R}_g} \dot{\mathbf{x}}_B, \boldsymbol{\rho}_{\mathcal{T}_q \mathcal{M}}(g) \dot{\mathbf{q}}_{\text{js}}) & & & &= (\underbrace{\boldsymbol{\rho}_{\mathbb{E}_3}(g)}_{\mathbf{R}_g} \dot{\mathbf{x}}_B, \dot{\mathbf{q}}_{\text{js}}). \end{aligned} \quad (4.3d)$$

$$\begin{aligned} &\begin{bmatrix} \mathbf{R}_g & \mathbf{0} \\ \mathbf{0} & \det(\mathbf{R}_g) \mathbf{R}_g \end{bmatrix} \begin{bmatrix} \dot{\mathbf{r}}_B \\ \boldsymbol{\omega}_B \end{bmatrix} & & & \begin{bmatrix} \mathbf{R}_g & \mathbf{0} \\ \mathbf{0} & \det(\mathbf{R}_g) \mathbf{R}_g \end{bmatrix} \begin{bmatrix} \dot{\mathbf{r}}_B \\ \boldsymbol{\omega}_B \end{bmatrix} \end{aligned}$$

Here, the determinant factor appears in both angular-velocity blocks because $\boldsymbol{\omega}_B$ is represented as an axial vector under a fixed right-handed convention. Equivalently, for any $\mathbf{R}_g \in \mathbb{O}_3$ and $[\boldsymbol{\omega}_B]_{\times} \in \mathfrak{so}$, one has $\mathbf{R}_g [\boldsymbol{\omega}_B]_{\times} \mathbf{R}_g^{\top} = [\det(\mathbf{R}_g) \mathbf{R}_g \boldsymbol{\omega}_B]_{\times}$. Thus, reflections act on linear velocities as polar vectors but on angular velocities as pseudovectors. Consequently, both base-velocity actions share the same representation on $(\dot{\mathbf{r}}_B, \boldsymbol{\omega}_B)$; their distinction lies in the action on the base configuration and on the joint-space coordinates.

Next, we discuss the structure of the joint-space group action $(\triangleright_{\mathcal{M}})$, which captures the reconfiguration of the robot's body needed to mimic the corresponding Euclidean isometry on the robot's morphology.

4.1 Modularity and symmetry in the robot's morphology

Robotic systems often exhibit a regular distribution of replicated kinematic branches within their kinematic tree. Examples include the four identical legs of the Mini-Cheetah robot and the mirrored arms and legs of the Atlas robot (see fig. 4.1). As we characterize in this section, modularity in a robot's morphology is one of the necessary conditions for the existence of a morphological symmetry $g \in \mathbb{G}$ acting on the robot's state, thereby inducing the group action on the joint-space configuration manifold $\triangleright_{\mathcal{M}} : \mathbb{G} \times \mathcal{M} \rightarrow \mathcal{M}$ that handles the permutation of limb roles in the examples of figs. 4.1 and 9.8. In section 4.2, we then characterize the constraints on the mass distribution of the robot's bodies and on the kinematic-tree structure that are necessary for $\triangleright_{\mathcal{M}}$ to exist.




Figure 4.2 Example modular kinematic trees used in this section. **Left:** Atlas, with torso base body B , one head branch instance b_1^{head} , one hip branch instance b_1^{hip} , two arm branch instances b_1^{arm} and b_2^{arm} , and two leg branch instances b_1^{leg} and b_2^{leg} . **Right:** Mini-Cheetah, with torso base body B and four replicated leg branch instances $b_1^{\text{leg}}, \dots, b_4^{\text{leg}}$. Bodies are shown as nodes and joints as links.

Note: The definitions in this section are used to define the morphological group actions for the suite of robotic systems included in the `morpho_symm` library (see section 1.5.1), which identifies the group actions $\triangleright_{\mathcal{Q}}$ and $\triangleright_{\mathcal{T}_q\mathcal{Q}}$ of eqs. 4.3a and 4.3d for a large collection of robotic systems.

Modular kinematic tree structure Assume that the robot's kinematic tree is composed of a base body B and n_k unique kinematic branches, represented by the label set $\mathbb{B} = \{b^1, \dots, b^{n_k}\}$. Each branch b^i has $n_{\text{dof}}(b^i) \in \mathbb{N}$ DoF, and we denote by $n_{\text{rep}}(b^i) \in \mathbb{N}$ the number of times that branch is replicated within the tree. The instances of a kinematic branch $b^i \in \mathbb{B}$ are denoted by $\mathbb{B}_i = \{b_1^i, \dots, b_{n_{\text{rep}}(b^i)}^i\}$. As a working example, we consider the humanoid robot Atlas, whose unique kinematic branches are the arm, leg, head, and hip, that is, $\mathbb{B} = \{b^{\text{arm}}, b^{\text{leg}}, b^{\text{head}}, b^{\text{hip}}\}$. These substructures occur with multiplicities $n_{\text{rep}}(b^{\text{arm}}) = n_{\text{rep}}(b^{\text{leg}}) = 2$ and $n_{\text{rep}}(b^{\text{head}}) = n_{\text{rep}}(b^{\text{hip}}) = 1$. We also consider the quadruped robot Mini-Cheetah, which features a single unique kinematic branch, $\mathbb{B} = \{b^{\text{leg}}\}$, with multiplicity $n_{\text{rep}}(b^{\text{leg}}) = 4$. See fig. 4.2.

The action of a morphological symmetry on the kinematic tree can be defined as a permutation that preserves kinematic-branch types. Equivalently, it can be viewed as a permutation of instances of the same branch type, specified by a group action on each branch type. That is, each \mathbb{B}_i is endowed with a group action:

$$\begin{aligned} \triangleright_{\mathbb{B}_i} : \mathbb{G} \times \mathbb{B}_i &\rightarrow \mathbb{B}_i \\ (g, b_j^i) &\mapsto g \triangleright b_j^i := b_{\gamma_g(j)}^i, \quad \text{with } \gamma_g : \{1, \dots, n_{\text{rep}}(b^i)\} \rightarrow \{1, \dots, n_{\text{rep}}(b^i)\}, \end{aligned} \quad (4.4)$$

which permutes the instances of the same branch type. In particular, $g \triangleright_{\mathbb{B}_i} b_j^i := b_{\gamma_g(j)}^i \in \mathbb{B}_i$ denotes that, under the symmetry transformation g , instance j of branch type i is mapped to the instance indexed by $\gamma_g(j)$.

For example, for the Atlas robot (see fig. 4.2), the group action on the arm branch type is defined as a permutation of the right and left arms:

$$\begin{aligned} \triangleright_{\mathbb{B}_{\text{arm}}} : \mathbb{C}_2 \times \mathbb{B}_{\text{arm}} &\rightarrow \mathbb{B}_{\text{arm}}, \\ (g, b_j^{\text{arm}}) &\mapsto g \triangleright b_j^{\text{arm}} := b_{\gamma_g(j)}^{\text{arm}} \end{aligned} \quad \text{with} \quad \begin{aligned} \gamma_e(\text{right}) &= \text{right}, & \gamma_e(\text{left}) &= \text{left}, \\ \gamma_{g_s}(\text{right}) &= \text{left}, & \gamma_{g_s}(\text{left}) &= \text{right}. \end{aligned}$$




Figure 4.3 Permutations of the Mini-Cheetah leg-branch instances induced by the elements of $\mathbb{G} = \{e, g_s\} \times \{e, g_t\} \times \{e, g_f\}$ (see fig. 1.3). Each panel shows the kinematic tree of the robot, with bodies as nodes and joints as links. Dashed bidirectional arrows connect the branch instances exchanged by the corresponding symmetry transformation. The reflection g_f acts through the within-branch joint-space transformation but does not permute the leg-branch instances.

Likewise, fig. 4.3 shows the permutations of the leg branches of the Mini-Cheetah robot induced by the 8 symmetry transformations composing its morphological symmetry group (see fig. 1.3).

When exchanging the roles of kinematic branches of the same type, it is often necessary not only to permute the branch joint configurations, but also to apply specific transformations to the joints within each branch so that the resulting state remains feasible and dynamically equivalent to the original one. For instance, when permuting the right and left arms of the Atlas robot, one must apply specific reflections to some joints because corresponding joints in the right and left arms may use different conventions for the positive direction of their DoF (e.g., the positive direction of the shoulder-roll joint in the right arm may be defined as clockwise, whereas in the left arm it may be defined as counterclockwise). Another example is the transformations induced by the dorsal reflection symmetry g_f of the Mini-Cheetah robot which flips the robot upside down (see fig. 1.3). While this symmetry has a trivial permutation of the leg branches, it requires specific reflections of the joints within each leg to yield a feasible and dynamically equivalent configuration (see fig. 4.3).

Let $\mathcal{M}_{b^i} \subseteq \mathbb{R}^{n_{dof}(b^i)}$ denote the configuration space of a single instance of the kinematic branch type b^i . Then each symmetry transformation $g \in \mathbb{G}$ defines a unique group action on \mathcal{M}_{b^i} that applies the required branch-joint-space reconfiguration:

$$\begin{aligned} \triangleright_{\mathcal{M}_{b^i}} : \mathbb{G} \times \mathcal{M}_{b^i} &\rightarrow \mathcal{M}_{b^i}, \\ (g, \mathbf{q}_{js, b^i}) &\mapsto g \triangleright \mathbf{q}_{js, b^i} := \boldsymbol{\rho}_{\mathcal{M}_{b^i}}(g) \mathbf{q}_{js, b^i}, \quad \text{with } \boldsymbol{\rho}_{\mathcal{M}_{b^i}}(g) \in \mathbb{R}^{n_{dof}(b^i) \times n_{dof}(b^i)}. \end{aligned} \quad (4.5)$$

Each $\boldsymbol{\rho}_{\mathcal{M}_{b^i}}(g)$ is a linear invertible transformation (definition 2.5) that specifies how the morphological symmetry transformation g acts on the joint configuration of a single instance of branch type b^i . This

enables us to define the group action on the joint-space configuration manifold collecting all instances of branch type i . Let $\mathcal{M}_{\mathbb{B}_i} \subseteq \mathbb{R}^{n_{\text{dof}}(\mathbb{B}_i)}$ denote the joint-space configuration manifold of all instances of the kinematic branch type b^i , with $\mathcal{M}_{\mathbb{B}_i} = \prod_{j=1}^{n_{\text{rep}}(b^i)} \mathcal{M}_{b_j^i}$. Then the action on this subset of the joint-space configuration manifold is defined as the composition of eqs. 4.4 and 4.5:

$$\begin{aligned} \triangleright_{\mathcal{M}_{\mathbb{B}_i}} : \mathbb{G} \times \mathcal{M}_{\mathbb{B}_i} &\rightarrow \mathcal{M}_{\mathbb{B}_i}, \\ (g, \mathbf{q}_{\text{js}, \mathbb{B}_i}) &\mapsto g \triangleright_{\mathcal{M}_{\mathbb{B}_i}} \mathbf{q}_{\text{js}, \mathbb{B}_i} := \boldsymbol{\rho}_{\mathcal{M}_{\mathbb{B}_i}}(g) \mathbf{q}_{\text{js}, \mathbb{B}_i} = (\boldsymbol{\rho}_{\mathbb{B}_i}(g) \otimes \boldsymbol{\rho}_{\mathcal{M}_{b^i}}(g)) \mathbf{q}_{\text{js}, \mathbb{B}_i} \end{aligned} \quad (4.6)$$

Here, $\boldsymbol{\rho}_{\mathbb{B}_i}(g) \in \mathbb{R}^{n_{\text{rep}}(b^i) \times n_{\text{rep}}(b^i)}$ is the permutation-matrix representation of the group action on the instances of the kinematic branch type b^i (eq. 4.4), and \otimes denotes the Kronecker product. Intuitively, this action applies eq. 4.5 to each instance of the kinematic branch type b^i and then permutes the resulting configurations according to the branch permutation defined by $\boldsymbol{\rho}_{\mathbb{B}_i}(g)$ (eq. 4.4).

Crucially, because the joint-space configuration manifold \mathcal{M} decomposes into the configuration manifolds associated with each unique kinematic branch type, that is, $\mathcal{M} = \prod_{i=1}^{n_k} \mathcal{M}_{\mathbb{B}_i}$, the group action on the full joint-space configuration manifold is defined as:

$$\begin{aligned} \triangleright_{\mathcal{M}} : \mathbb{G} \times \mathcal{M} &\rightarrow \mathcal{M}, \\ (g, \mathbf{q}_{\text{js}}) &\mapsto g \triangleright \mathbf{q}_{\text{js}} := \boldsymbol{\rho}_{\mathcal{M}}(g) \mathbf{q}_{\text{js}} := \left(\bigoplus_{i=1}^{n_k} \boldsymbol{\rho}_{\mathcal{M}_{\mathbb{B}_i}}(g) \right) \mathbf{q}_{\text{js}} = \begin{bmatrix} \boldsymbol{\rho}_{\mathcal{M}_{\mathbb{B}_1}}(g) & & \\ & \ddots & \\ & & \boldsymbol{\rho}_{\mathcal{M}_{\mathbb{B}_{n_k}}}(g) \end{bmatrix} \begin{bmatrix} \mathbf{q}_{\text{js}, \mathbb{B}_1} \\ \vdots \\ \mathbf{q}_{\text{js}, \mathbb{B}_{n_k}} \end{bmatrix} \end{aligned} \quad (4.7)$$

If a point $\mathbf{q}_{\text{js}} \in \mathcal{M}$ is represented in a basis that respects the decomposition of \mathcal{M} into the configuration manifolds of each unique kinematic branch type, i.e., $\mathbf{q}_{\text{js}} = (\mathbf{q}_{\text{js}, \mathbb{B}_1}, \dots, \mathbf{q}_{\text{js}, \mathbb{B}_{n_k}})$, then the action of a morphological symmetry $g \in \mathbb{G}$ on \mathbf{q}_{js} is given by the block-diagonal matrix $\boldsymbol{\rho}_{\mathcal{M}}(g)$, where each block corresponds to the action on the joint-space configuration manifold of a unique kinematic branch type. This structure allows us to systematically construct the group action on the full joint-space configuration manifold from the actions defined on each unique kinematic branch type (eq. 4.6).

Summary 4.1

A morphological symmetry $g \in \mathbb{G}$ induces a non-trivial transformation of the joint-space configuration $\mathbf{q}_{\text{js}} \in \mathcal{M}$, required to reconfigure the robot's body so as to mimic the corresponding Euclidean isometry. This transformation is defined by the group action $\triangleright_{\mathcal{M}} : \mathbb{G} \times \mathcal{M} \rightarrow \mathcal{M}$ (eq. 4.7), which is assembled from the actions on the joint-space configurations associated with each unique kinematic branch type, $\triangleright_{\mathcal{M}_{\mathbb{B}_i}} : \mathbb{G} \times \mathcal{M}_{\mathbb{B}_i} \rightarrow \mathcal{M}_{\mathbb{B}_i}$. Examples include the joint-space configuration of all legs of the Mini-Cheetah robot and that of both arms of the Atlas robot. In turn, each $\triangleright_{\mathcal{M}_{\mathbb{B}_i}}$ is defined by combining two actions: (i) the permutation of the instances of branch type b^i , given by $\triangleright_{\mathbb{B}_i} : \mathbb{G} \times \mathbb{B}_i \rightarrow \mathbb{B}_i$ (eq. 4.4), for example, the permutation of the right and left arms of the Atlas robot; and (ii) the transformation of the joint-space configuration of a single instance of branch type b^i , given by $\triangleright_{\mathcal{M}_{b^i}} : \mathbb{G} \times \mathcal{M}_{b^i} \rightarrow \mathcal{M}_{b^i}$ (eq. 4.5), for example, the reflection applied to some joints when permuting the right and left arms of the Atlas robot.

Next, we discuss the algebraic constraints on the robot's morphology that are necessary for the existence of the group actions defined above.

4.2 Morphology constraints induced by symmetries

The previous section described the group action on the joint-space configuration manifold, $\triangleright_{\mathcal{M}} : \mathbb{G} \times \mathcal{M} \rightarrow \mathcal{M}$, as the composition of a permutation of kinematic branches with transformations of their individual joint-configuration manifolds. These transformations implicitly define, for every $g \in \mathbb{G}$, a permutation of bodies within the kinematic tree. In this section, we derive the algebraic constraints on the robot's kinematic-tree structure and mass distribution required for such body permutations to exist, since they are necessary conditions for the existence of a morphological symmetry $g \in \mathbb{G}$.

To derive morphological constraints on the parameters of a robot's kinematic tree and mass distribution, we study the conditions under which the robot's kinetic energy is \mathbb{G} -invariant, which is necessary because morphological symmetries are energy-preserving transformations (definition 4.1).

Recall that the robot's kinetic energy is defined as the aggregate of energies of all the n_b bodies in the kinematic tree (see eq. 2.3a and section 2.1):

$$T(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2} \dot{\mathbf{q}}^\top \mathbf{M}(\mathbf{q}) \dot{\mathbf{q}} = \sum_{i=1}^{n_b} \underbrace{\frac{1}{2} \dot{\mathbf{x}}_i^\top \mathbf{V}_i \dot{\mathbf{x}}_i}_{T_i(\mathbf{q}, \dot{\mathbf{q}})}, \quad \text{with } \mathbf{V}_i := \begin{bmatrix} m_i \text{Id}_3 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_i \end{bmatrix} \in \mathbb{S}_{++}^6. \quad (4.8)$$

where $\dot{\mathbf{x}}_i = \begin{bmatrix} \dot{\mathbf{r}}_i \\ \dot{\boldsymbol{\omega}}_i \end{bmatrix} \in \mathbb{R}^6$ denotes the i^{th} body's spatial velocity vector, composed of the linear velocity $\dot{\mathbf{r}}_i \in \mathbb{R}^3$ and angular velocity $\dot{\boldsymbol{\omega}}_i \in \mathbb{R}^3$ (see eq. 2.2); $\mathbf{V}_i \in \mathbb{S}_{++}^6$ denotes the body's spatial inertia matrix, composed of the mass $m_i \in \mathbb{R}$ and inertia matrix $\mathbf{I}_i \in \mathbb{S}_{++}^3$ (defined w.r.t. a frame at the body's CoM aligned with its principal axes of inertia (Featherstone, 2008)); and $T_i(\mathbf{q}, \dot{\mathbf{q}})$ denotes the body's kinetic energy.

For eq. 4.8 to be \mathbb{G} -invariant, the sum of the kinetic energies of all bodies at any symmetry-transformed state $(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})$ must equal the kinetic energy at the original state $(\mathbf{q}, \dot{\mathbf{q}})$. This can only happen if, given the set of transformed body kinetic energies $\{T_i(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})\}_{i=1}^{n_b}$, there exists a permutation $\sigma_g : \{1, \dots, n_b\} \mapsto \{1, \dots, n_b\}$ that matches the energies of the g -transformed bodies to those of the original bodies. Note that this permutation is induced by the permutation of kinematic branches of the same type defined in eqs. 4.4 and 4.7 and depicted in fig. 4.3 for the Mini-Cheetah robot.

Specifically, if \mathbb{G} is the morphological symmetry group of a robot, then for every $g \in \mathbb{G}$ there must exist a permutation σ_g such that the following algebraic constraints hold for all $g \in \mathbb{G}$, $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}$, and $i \in \{1, \dots, n_b\}$:

$$T_i(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}) = T_{\sigma_g(i)}(\mathbf{q}, \dot{\mathbf{q}}) \quad (4.9a)$$

By substituting eq. 4.8, we can derive *kinematic constraints* on the robot's morphology that enforce that the spatial velocity of the i^{th} body at the symmetry-transformed state $(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})$ is the g -transformed spatial velocity of the $\sigma_g(i)^{\text{th}}$ body at state $(\mathbf{q}, \dot{\mathbf{q}})$. Recalling that the spatial velocity vector of the i^{th} body, $\dot{\mathbf{x}}_i(\mathbf{q}, \dot{\mathbf{q}}) := \mathbf{J}_i(\mathbf{q}) \dot{\mathbf{q}}$, is determined by the body's Jacobian matrix-valued function $\mathbf{J}_i : \mathcal{Q} \mapsto \mathbb{R}^{6 \times n_q}$ (see eq. 2.2), these kinematic constraints translate to the robot's kinematic tree as matrix-equivariance constraints on the robot's Jacobians:

$$\begin{aligned} \dot{\mathbf{x}}_i(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}}) &:= g \triangleright_{\epsilon} \dot{\mathbf{x}}_{\sigma_g(i)}(\mathbf{q}, \dot{\mathbf{q}}) \quad \forall g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{T}\mathcal{Q}, i \in [1, n_b] \\ \mathbf{J}_i(g \triangleright_{\mathcal{Q}} \mathbf{q}) g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} &:= g \triangleright_{\epsilon} \mathbf{J}_{\sigma_g(i)}(\mathbf{q}) \dot{\mathbf{q}} \\ \mathbf{J}_i(g \triangleright_{\mathcal{Q}} \mathbf{q}) \tau_{q \mathcal{Q} \triangleleft g} &:= g \triangleright_{\epsilon} \mathbf{J}_{\sigma_g(i)}(\mathbf{q}) \end{aligned} \quad (4.9b)$$

Likewise, we can derive *inertial constraints* on the robot's morphology that relate the mass and inertia of the i^{th} body at $(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})$ to those of the $\sigma_g(i)^{\text{th}}$ body at $(\mathbf{q}, \dot{\mathbf{q}})$.

$$\begin{aligned} \dot{\mathbf{x}}_i^\top \mathbf{V}_i \dot{\mathbf{x}}_i &= (g \triangleright_{\mathcal{C}} \dot{\mathbf{x}}_{\sigma_g(i)})^\top \mathbf{V}_{\sigma_g(i)} (g \triangleright_{\mathcal{C}} \dot{\mathbf{x}}_{\sigma_g(i)}), \quad \forall g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}, i \in [1, n_b] \\ [\dot{\mathbf{r}}_i]^\top \begin{bmatrix} m_i \text{Id}_3 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_i \end{bmatrix} [\dot{\mathbf{r}}_i] &= \left(g \triangleright_{\mathcal{C}} \begin{bmatrix} \dot{\mathbf{r}}_{\sigma_g(i)} \\ \boldsymbol{\omega}_{\sigma_g(i)} \end{bmatrix} \right)^\top \begin{bmatrix} m_{\sigma_g(i)} \text{Id}_3 & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{\sigma_g(i)} \end{bmatrix} \left(g \triangleright_{\mathcal{C}} \begin{bmatrix} \dot{\mathbf{r}}_{\sigma_g(i)} \\ \boldsymbol{\omega}_{\sigma_g(i)} \end{bmatrix} \right) \end{aligned} \quad (4.9c)$$

These yield mass-equality constraints for bodies related by σ_g , together with conjugation constraints on their inertia matrices:

$$\begin{aligned} m_i \text{Id}_3 &= \mathbf{R}_g^\top m_{\sigma_g(i)} \text{Id}_3 \mathbf{R}_g & \omega_i^\top \mathbf{I}_i \omega_i &= \omega_{\sigma_g(i)}^\top \mathbf{I}_{\sigma_g(i)} \omega_{\sigma_g(i)} \\ m_i &= m_{\sigma_g(i)} & \omega_i^\top \mathbf{I}_i \omega_i &= (\det(\mathbf{R}_g) \mathbf{R}_g^\top \omega_i)^\top \mathbf{I}_{\sigma_g(i)} (\det(\mathbf{R}_g) \mathbf{R}_g^\top \omega_i) \\ & & \mathbf{I}_i &= \mathbf{R}_g \mathbf{I}_{\sigma_g(i)} \mathbf{R}_g^\top \end{aligned} \quad (4.9d)$$

$\forall g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}, i \in [1, n_b]$

Here, the left-hand side of eq. 4.9d states that bodies related by σ_g must have the same mass, and the right-hand side states that their inertia matrices must be related by conjugation with \mathbf{R}_g . This ensures that the inertia of body i at $(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})$ is g -equivalent to the inertia of body $\sigma_g(i)$ at $(\mathbf{q}, \dot{\mathbf{q}})$. To develop intuition behind these kinematic and inertial constraints, we study two cases: (i) trivial permutation of bodies, where for a given $g \in \mathbb{G}$ and body i we have $\sigma_g(i) = i$, and (ii) non-trivial permutation of bodies, where $\sigma_g(i) \neq i$.

- * $\sigma_g(i) = i$: For a symmetry transformation g that induces a trivial permutation of body i —for instance, the torso, hip, or head of the Atlas robot under a sagittal reflection g_s (see fig. 4.1)—the kinematic and inertial constraints in eqs. 4.9b and 4.9d reduce to:

$$\mathbf{J}_i(g \triangleright_{\mathcal{Q}} \mathbf{q})_{\mathcal{T}_q \mathcal{Q} \triangleleft g} := g \triangleright_{\mathcal{C}} \mathbf{J}_i(\mathbf{q}) \quad \text{and} \quad \mathbf{I}_i \mathbf{R}_g = \mathbf{R}_g \mathbf{I}_i \quad \forall \sigma_g(i) = i, g \in \mathbb{G}, (\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ} \quad (4.10)$$

Both constraints in this case can be interpreted as morphological constraints on the mass distribution of body i . From an algebraic perspective, the commutativity constraint on the inertia matrix reduces its free DoF (see details in prop. 2.2). The equivariance constraint on the Jacobian ensures that the body frame—aligned with the body's principal axes of inertia—at $(g \triangleright_{\mathcal{Q}} \mathbf{q}, g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}})$ is the g -transformed body frame at $(\mathbf{q}, \dot{\mathbf{q}})$. In practice, these constraints ensure that the body's mass distribution is invariant under conjugation by the isometry \mathbf{R}_g , such that if $g \in \mathbb{G}$ is a reflection, the body has a reflectional symmetry in its mass distribution (see fig. 4.5), and if $g \in \mathbb{G}$ is a rotation, the body has a rotational symmetry in its mass distribution.

Intuitively, for the case of reflections in fig. 4.1, eq. 4.10 is responsible for the sagittal reflection symmetry of the torso, head, and hip mass distribution in the humanoid robot Atlas under $g_s \in \mathbb{C}_2$, and the sagittal and transverse reflection symmetries together with the 180° rotational symmetry of the torso mass distribution in the Mini-Cheetah robot under $g_s, g_t, g_r \in \mathbb{C}_2 \times \mathbb{C}_2 \times \mathbb{C}_2$.

- * $\sigma_g(i) \neq i$: For any $g \in \mathbb{G}$ and body i for which the symmetry transformation implies a permutation of the dynamical roles of two bodies—such as the left and right limbs in robots with sagittal reflection symmetry, or the front, hind, right, and left limbs of the Mini-Cheetah robot in fig. 4.1—eq. 4.9d imposes mass-equality constraints between the two bodies, together




Figure 4.4 Example diagram of a 3-legged robot in 2-dimensional space with morphological symmetry group given by the cyclic group of order 3, $\mathbb{G} = \mathbb{C}_3 = \{e, g, g^2\}$. The robot is composed of seven rigid bodies: the base and the two linked bodies of each leg, $\{B, l_{a,1}, l_{a,2}, l_{b,1}, l_{b,2}, l_{c,1}, l_{c,2}\}$. The body permutation induced by the generator g is denoted by σ_g and permutes the leg roles as $\sigma_g(l_{a,i}) = l_{b,i}$, $\sigma_g(l_{b,i}) = l_{c,i}$, and $\sigma_g(l_{c,i}) = l_{a,i}$, for $i \in \{1, 2\}$. The only rigid body left invariant by σ_g is the base B , whose mass distribution must therefore be invariant to rotations by 120° and 240° (see eq. 4.10). For bodies permuted by σ_g , the kinematic constraints in eq. 4.9b ensure that the Jacobians, and therefore the kinematic chains, of permuted bodies represent identical kinematic structures. Similarly, the inertial constraints in eq. 4.9d ensure that the mass distributions of corresponding leg bodies are identical. Hence, the first link of each leg is required to be identical to the first link of the other legs, and likewise for the second link of each leg. The diagram further illustrates that, although the kinematic structure is compatible with the larger dihedral symmetry group \mathbb{D}_3 , the unequal mass distribution of the first links breaks the reflection symmetry $h \in \mathbb{D}_3$, leaving only the cyclic subgroup \mathbb{C}_3 as a valid morphological symmetry group.

with an algebraic constraint relating their inertia matrices under the conjugate action of the isometry $g \in \mathbb{G}$.

An illustrative depiction of these constraints is shown in fig. 4.4 for a 3-legged robot in 2-dimensional space with morphological symmetry group $\mathbb{G} = \mathbb{C}_3$. In this example, the leg bodies are permuted under the action of the symmetry group, which constrains both their mass distributions and their Jacobians to satisfy eqs. 4.9b and 4.9d. In practice, the kinematic constraints on the Jacobians force the kinematic chains of each leg to be identical copies of one another, while the inertial constraints ensure that the leg bodies are identical across corresponding links. Crucially, the figure also shows that the kinematic structure satisfies the kinematic constraints of a larger symmetry group, the dihedral group of order 6, \mathbb{D}_3 (fig. 2.1), which contains the cyclic group \mathbb{C}_3 as a subgroup. However, the unequal mass distribution of the first links of the legs violates the inertial constraints of eq. 4.9d for states transformed by the reflection element $h \in \mathbb{D}_3$, breaking the dynamics equivalence between any state (q, \dot{q}) and its reflection $(h \triangleright q, h \triangleright \dot{q})$. Hence, the morphological symmetry group of the robot is \mathbb{C}_3 , and not \mathbb{D}_3 .

Summary 4.2

The existence of a morphological symmetry $g \in \mathbb{G}$ induces identifiable constraints on the robot's kinematic tree and inertial parameters. These can be summarized as follows: if, for some $g \in \mathbb{G}$, the group action on the robot's joint-space configuration manifold $\triangleright_{\mathcal{M}} : \mathbb{G} \times \mathcal{M} \rightarrow \mathcal{M}$ (eq. 4.7) relates the configuration of the i^{th} body of one kinematic branch to that of the $\sigma_g(i)^{\text{th}}$ body of a different kinematic branch, then the following constraints must hold for all $(q, \dot{q}) \in \mathcal{TQ}$:

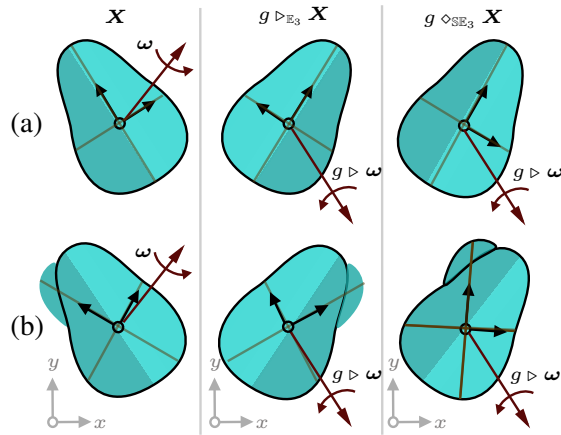


Figure 4.5 **Left:** A body with reflection-symmetric mass distribution can mimic the infeasible environmental reflection $\triangleright_{\mathbb{E}_3}$ (eq. 4.3b) through the feasible conjugate configuration $\diamond_{\mathbb{E}_3}$ (eq. 4.3c), while preserving the body's kinetic energy. **Right:** Breaking that mass-distribution symmetry destroys this equivalence, because the reconfiguration $\diamond_{\mathbb{E}_3}$ leads to a different body kinetic energy. Gray lines represent the body's principal axes of inertia, and the body frame is aligned with these axes. Figure adapted from Ordoñez-Apaez et al. (2025).

- * **Kinematic constraints:** The bodies' Jacobians must satisfy the matrix-equivariance constraint in eq. 4.9b.
- * **Inertial constraints:** The bodies' masses must be equal, and their inertia matrices must satisfy the conjugation constraint in eq. 4.9d.

4.3 Dynamics' harmonic analysis

Recall that in classical mechanics, the equations of motion of a body with mass m and inertia I in 3D space decouple along the orthogonal directions $\hat{x}, \hat{y}, \hat{z} \in \mathbb{R}^3$:

$$m\ddot{x} = f_{\hat{x}}, \quad m\ddot{y} = f_{\hat{y}}, \quad m\ddot{z} = f_{\hat{z}}, \quad (4.11)$$

where $(f_{\hat{x}}, f_{\hat{y}}, f_{\hat{z}})$ denotes the decomposition of the moving force into its orthogonal components along the \hat{x}, \hat{y} , and \hat{z} directions. This decomposition of the body dynamics follows from the orthogonality relations between the canonical axes, $\hat{x} \perp \hat{y} \perp \hat{z}$, which ensure that any component of the force $f_{\hat{x}}$ induces zero work along \hat{y} and \hat{z} , and thus does not contribute to the dynamics along those directions.

This section illustrates that morphological symmetries induce an analogous orthogonal decomposition of the robot's dynamics in the space of generalized coordinates. The key tool is a standard result from abstract harmonic analysis—namely the *isotypic decomposition* of a symmetric vector space (theorem 2.1), which generalizes the Fourier transform to arbitrary vector spaces and compact symmetry groups (Knapp, 1986)—and exposes orthogonal constraints between subspaces of the configuration space \mathcal{TQ} . To highlight the utility of identifying these orthogonal subspaces, we show that several relevant quadruped gaits, such as the trot, pace, bound, and pronk, can be modeled as dynamic motions constrained to evolve within one or a few of these subspaces rather than in the full configuration space (see fig. 4.6). This provides a direct route to identifying relevant subspaces for gait design and analysis.

Note that here we focus on the physical interpretation of the isotypic decomposition and do not delve into the representation-theoretic details underlying the analysis. Those details can be found in [theorem 2.1](#) and [Ordoñez-APraez et al. \(2025\)](#) and are referenced in the main text when needed.

Orthogonal configuration submanifolds induced by morphological symmetries

Recall that a robot with morphological symmetry group \mathbb{G} possesses a group action $\triangleright_{\mathcal{Q}} : \mathbb{G} \times \mathcal{Q} \rightarrow \mathcal{Q}$ (see eq. 4.3a). Once a basis is chosen to represent points $\mathbf{q} \in \mathcal{Q}$, this action can be written as a group representation $\rho_{\mathcal{Q}} : \mathbb{G} \mapsto \text{GL}(\mathbb{R}^{n_q})$, which maps each symmetry transformation $g \in \mathbb{G}$ to a linear invertible matrix $\rho_{\mathcal{Q}}(g) \in \mathbb{R}^{n_q \times n_q}$ defining how the coordinates of the robot's configuration manifold transform (see [definition 2.2](#)).¹

By [theorem 2.1](#), there exists a change of basis $\mathbf{T} : \mathbb{R}^{n_q} \mapsto \mathbb{R}^{n_q}$ that yields a block-diagonal decomposition of the representation $\rho_{\mathcal{Q}}$ and, equivalently, an orthogonal decomposition of the configuration space:

$$\rho_{\mathcal{Q}}^{\text{iso}}(g) = \mathbf{T}(\rho_{\mathcal{Q}}(g))\mathbf{T}^{\top} := \begin{bmatrix} \rho_{\mathcal{Q}_1^{\text{iso}}}(g) & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \rho_{\mathcal{Q}_2^{\text{iso}}}(g) & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \rho_{\mathcal{Q}_{n_{\text{iso}}}^{\text{iso}}}(g) \end{bmatrix} \iff \mathcal{Q} = \mathcal{Q}^{(1)} \oplus^{\perp} \mathcal{Q}^{(2)} \oplus^{\perp} \cdots \oplus^{\perp} \mathcal{Q}^{(n_{\text{iso}})} \quad (4.12)$$

Here, $\rho_{\mathcal{Q}}^{\text{iso}}$ denotes the group representation of \mathbb{G} in the *isotypic basis* ([definition 2.12](#)) of \mathcal{Q} . This characteristic basis of symmetric vector spaces exposes the orthogonal subspaces $\mathcal{Q}^{(k)}$, for $k \in \{1, \dots, n_{\text{iso}}\}$, each representing a configuration submanifold subject to orthogonality constraints with the others, much like the orthogonality relations between the \hat{x} , \hat{y} , and \hat{z} axes in \mathbb{R}^3 . Consequently, any generalized force acting along the directions of motion of one subspace $\mathcal{Q}^{(k)}$ does zero work along the directions of motion of any other subspace $\mathcal{Q}^{(k')}$ for $k' \neq k$.

To comprehend the implications and physical interpretations of this result in the context of robot dynamics, we must first characterize the structure of the isotypic decomposition, which depends entirely on the morphological symmetry group \mathbb{G} :

- * The number of orthogonal *isotypic subspaces*, $n_{\text{iso}} \in [1, |\mathbb{G}|]$, is determined by the number of real irreducible representations of the group, $\{\bar{\rho}_k\}_{k=1}^{n_{\text{iso}}}$, each characterizing a distinct subset of the symmetries of \mathbb{G} (see [definition 2.11](#) and [fig. 4.6](#)).
- * Each configuration isotypic subspace $\mathcal{Q}^{(k)}$ is constructed to feature only a subset of the symmetries of \mathbb{G} , making $\mathcal{Q}^{(k)}$ a space of symmetry-constrained robot configurations (see [fig. 4.6](#)). Formally, the effective symmetry group of $\mathcal{Q}^{(k)}$ is the quotient subgroup $\mathbb{G}^{(k)} := \mathbb{G}/\mathbb{N}_k$, where \mathbb{N}_k is the kernel of the irreducible representation of type k (see [definition 2.13](#)). For example, in [fig. 4.6](#), the first isotypic subspace $\mathcal{Q}^{(1)}$ defines the space of \mathbb{G} -invariant configurations, meaning that all legs share the same configuration. In this case, $\mathbb{N}_1 = \mathbb{K}_4$ and therefore $\mathbb{G}^{(1)} = \{e\}$. The second isotypic subspace $\mathcal{Q}^{(2)}$ defines configurations in which the left and right legs are synchronized, so $\mathbb{N}_2 = \{e, g_s\}$ and $\mathbb{G}^{(2)} \equiv \mathbb{C}_2$. Similarly, for $\mathcal{Q}^{(3)}$ we have $\mathbb{N}_3 = \{e, g_t\}$ and $\mathbb{G}^{(3)} \equiv \mathbb{C}_2$, whereas for $\mathcal{Q}^{(4)}$ we have $\mathbb{N}_4 = \{e, g_r\}$ and $\mathbb{G}^{(4)} \equiv \mathbb{C}_2$.
- * The orthogonality relations between the isotypic subspaces $\mathcal{Q}^{(k)}$ follow from Schur's orthogonality relations (see [lemma 2.1](#)). In our context, this has a direct physical interpretation:

¹Note that the group representation $\rho_{\mathcal{Q}}$ is well defined only when the robot orientation is represented in terms of Euler angles, which is implicitly assumed in this section.

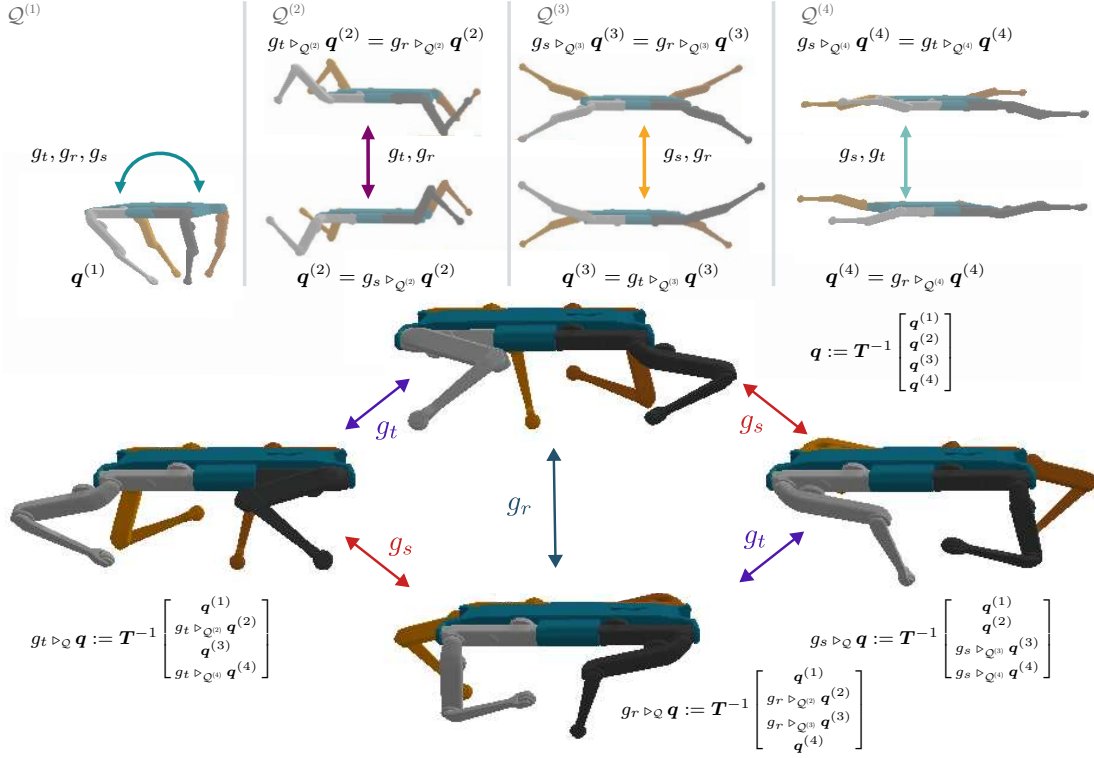


Figure 4.6 Illustration of the projection of the symmetry-related states of the Solo quadruped robot (Griminger et al., 2020), $\mathbb{K}_4 \mathbf{q} = \{g \triangleright_{\mathcal{Q}} \mathbf{q} | g \in \mathbb{K}_4\}$ (bottom-half), onto the isotypic subspaces $\{\mathcal{Q}^{(k)}\}_{k=1}^4$ of the robot's configuration manifold \mathcal{Q} (upper-half). For any state $\mathbf{q} \in \mathcal{Q}$, the matrix T encodes the change to the isotypic basis, in which the state decomposes into its projections onto the isotypic subspaces: $\mathbf{q} = T^{-1}(\mathbf{q}^{(1)} \oplus \mathbf{q}^{(2)} \oplus \mathbf{q}^{(3)} \oplus \mathbf{q}^{(4)})$, with $\mathbf{q}^{(k)} \in \mathcal{Q}^{(k)}$ for all $k \in [1, 4]$. Each isotypic subspace describes a family of symmetry-constrained robot configurations. For example, \mathcal{Q}^1 describes configurations in which all legs share the same configuration, \mathcal{Q}^2 those in which the left and right legs are synchronized, \mathcal{Q}^3 those in which the front and back legs are synchronized, and \mathcal{Q}^4 those in which diagonally opposed legs are synchronized. In a basis that exposes these subspaces, the group actions act block-wise on subsets of the configuration space while leaving the remaining subspaces invariant, as illustrated by the orbit $\mathbb{K}_4 \mathbf{q} = \{g \triangleright_{\mathcal{Q}} \mathbf{q} | g \in \mathbb{K}_4\}$ in the bottom half. Figure adapted from Ordoñez-Aperez et al. (2025).

any direction of displacement $\delta \mathbf{q}^{(k)} \in \mathcal{Q}^{(k)}$ is orthogonal to any direction of displacement $\delta \mathbf{q}^{(k')} \in \mathcal{Q}^{(k')}$ for $k' \neq k$, in the sense that no generalized force acting along $\delta \mathbf{q}^{(k)}$ can do work along $\delta \mathbf{q}^{(k')}$, and vice versa. This, in turn, enables us to derive independent equations of motion for each isotypic subspace, analogous to eq. 4.11, in which only inertial terms are shared across the subspaces equations of motion. We discuss this in more detail in the next section.

Decomposing motions into a superposition of normal configuration modes

The decomposition of the configuration manifold into orthogonal submanifolds $\mathcal{Q} = \mathcal{Q}^{(1)} \oplus^{\perp} \oplus^{\perp} \dots \oplus^{\perp} \mathcal{Q}^{(n_{\text{iso}})}$ implies that any robot configuration $\mathbf{q} \in \mathcal{Q}$ decomposes as a superposition of its projections onto the isotypic subspaces: $\mathbf{q} = T^{-1}(\mathbf{q}^{(1)} \oplus \mathbf{q}^{(2)} \oplus \dots \oplus \mathbf{q}^{(n_{\text{iso}})})$, where $\mathbf{q}^{(k)} \in \mathcal{Q}^{(k)}$ denotes the projection of \mathbf{q} onto the k^{th} isotypic subspace (see fig. 4.6). Given that for every $\mathbf{q} \in \mathcal{Q}$ the tangent space $\mathcal{T}_{\mathbf{q}} \mathcal{Q}$ inherits the group action $\triangleright_{\mathcal{T}_{\mathbf{q}} \mathcal{Q}}$ (see eq. 4.3d), the same decomposition applies

to the robot's velocity $\dot{\mathbf{q}}$, acceleration $\ddot{\mathbf{q}}$, and generalized forces $\boldsymbol{\tau}$, such that:

$$\mathbf{q} := \mathbf{T}^{-1} \begin{bmatrix} \mathbf{q}^{(1)} \\ \vdots \\ \mathbf{q}^{(n_{\text{iso}})} \end{bmatrix}, \quad \dot{\mathbf{q}} := \mathbf{T}^{-1} \begin{bmatrix} \dot{\mathbf{q}}^{(1)} \\ \vdots \\ \dot{\mathbf{q}}^{(n_{\text{iso}})} \end{bmatrix}, \quad \ddot{\mathbf{q}} := \mathbf{T}^{-1} \begin{bmatrix} \ddot{\mathbf{q}}^{(1)} \\ \vdots \\ \ddot{\mathbf{q}}^{(n_{\text{iso}})} \end{bmatrix}, \quad \text{and} \quad \boldsymbol{\tau} := \mathbf{T}^{-1} \begin{bmatrix} \boldsymbol{\tau}^{(1)} \\ \vdots \\ \boldsymbol{\tau}^{(n_{\text{iso}})} \end{bmatrix} \quad (4.13)$$

Where $\mathbf{q}^{(k)}$, $\dot{\mathbf{q}}^{(k)}$, $\ddot{\mathbf{q}}^{(k)}$, and $\boldsymbol{\tau}^{(k)}$ denote the projections of the generalized positions, velocities, accelerations, and forces onto the k^{th} isotypic subspace, of dimension $n_{k,j} = |\mathcal{M}_k^{\text{iso}}| \leq |\mathcal{M}| = n_{\text{js}}$ (see fig. 4.6 and animation). In the isotypic basis, each dimension of $\mathbf{T}\mathbf{q} \in \mathbb{R}^{n_{\text{js}}}$ no longer describes the position of a single DoF, but rather a synergy of DoF. Likewise, the dimensions of $\mathbf{T}\dot{\mathbf{q}}$, $\mathbf{T}\ddot{\mathbf{q}}$, and $\mathbf{T}\boldsymbol{\tau}$ describe the corresponding synergy velocities, accelerations, and generalized forces (see fig. 4.6).

Moreover, orthogonality between isotypic subspaces implies that generalized moving forces associated with one subspace do zero work along displacement directions in the others. Consequently, the instantaneous work decomposes into the sum of the work performed in each isotypic subspace:

$$W(\dot{\mathbf{q}}, \boldsymbol{\tau}) = \boldsymbol{\tau}^\top \dot{\mathbf{q}} = \sum_{k=1}^{n_{\text{iso}}} W_k := \sum_{k=1}^{n_{\text{iso}}} (\boldsymbol{\tau}^{(k)})^\top \dot{\mathbf{q}}^{(k)} \quad (4.14)$$

Normal configuration modes (NCMs) In molecular physics, a related symmetry-aware decomposition of states and dynamics is commonly described as a decomposition into *normal vibrational modes* (Dresselhaus et al., 2007). By analogy, we refer to the robot-motion decomposition in eq. 4.13 as a decomposition into **normal configuration modes (NCMs)** of motion. Each isotypic subspace $\mathcal{Q}^{(k)}$ then describes a family of symmetry-constrained robot configurations, namely one **NCM**.

NCMs of the dynamics of quadruped locomotion

To empirically assess the relevance of the isotypic decomposition and the associated **NCMs** for robot dynamics, Ordoñez-Appaez et al. (2024, 2025) study the decomposition of several Mini-Cheetah locomotion gaits and terrains into its $n_{\text{iso}} = 4$ **NCMs**. These modes match the structure of the **NCMs** of the Solo quadruped in fig. 4.6, since both robots share the same morphological symmetry group $\mathbb{G} = \mathbb{K}_4$.

Specifically, we take real-world motion trajectories $[(\mathbf{q}_t, \dot{\mathbf{q}}_t)]_{t \in \mathbb{T}}$, collected over a time span \mathbb{T} from the dataset of Lin et al. (2021), identify the change of basis \mathbf{T} to the isotypic basis of \mathcal{Q} using Dixon's reduction algorithm (Dixon, 1970), implemented in `symm_learning` (see theorem 2.1 and section 1.5.1), and project the trajectories into the isotypic subspaces $\{\mathcal{Q}^{(k)}\}_{k=1}^4$ via eq. 4.13. We then study the kinetic energy of each **NCM** as a proxy for its relevance in characterizing the dynamics of each gait and terrain.

The results, shown in fig. 4.7, demonstrate that decomposing the motion into **NCMs** readily reveals the lower-dimensional structure of the trotting, jumping, pronking, and galloping gaits of the Mini-Cheetah robot. Each gait primarily evolves in a 3- or 6-dimensional joint-space subspace composed of one or two **NCMs**. Furthermore, the less relevant subspaces for each gait appear to be excited only transiently during disturbance rejection, as if the controller actively penalized occupancy of those subspaces to keep the overall motion in a lower-dimensional space. Interestingly, even when the terrain changes, the distribution of kinetic energy across isotypic subspaces remains largely unchanged, as shown for trotting in fig. 4.7-left. These empirical results suggest that optimal control

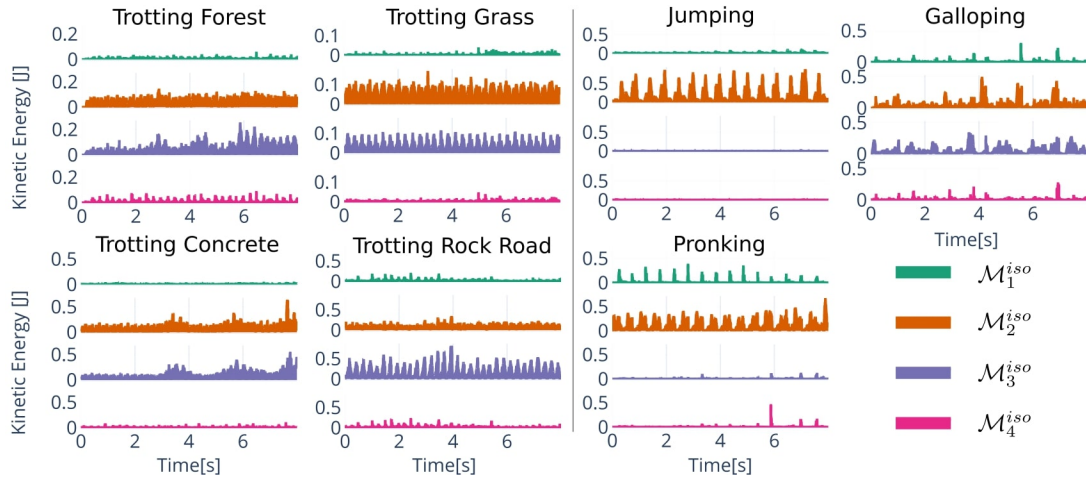


Figure 4.7 Decomposition of the joint space dynamics of legged locomotion of the Mini-Cheetah robot as a superposition of lower dimensional NCM. The dynamics of locomotion for different gait types and terrains are characterized by the projection of trajectories of motion onto the isotypic spaces $\mathcal{M}^{(1)}, \mathcal{M}^{(2)}, \mathcal{M}^{(3)}, \mathcal{M}^{(4)} \subseteq \mathbb{R}^3$. The relevance of each isotypic subspace, per gait type and terrain, is quantified by computing the joint space kinetic energy of the projected motion in the isotypic subspace. **Left:** The trotting gait dynamics on forest, grass, concrete, road with rocks terrains primarily evolve in the $\mathcal{M}^{(2)} \times \mathcal{M}^{(3)} \in \mathbb{R}^6$ subspace, as indicated by the kinetic energy distribution. **Right:** The jumping, pronking, gaits on asphalt terrain mainly evolve in $\mathcal{M}^{(3)}$ and $\mathcal{M}^{(1)} \times \mathcal{M}^{(2)} \in \mathbb{R}^6$ subspaces, respectively, while the galloping gait dynamics periodically evolve in all isotypic subspaces. Figure adapted from [Ordoñez-Apaez et al. \(2024\)](#).

strategies for quadrupedal locomotion in unstructured terrain converge to gaits with a low effective dimensionality of evolution relative to the full configuration space, and that NCMs provide a natural basis for describing this low-dimensional structure of the dynamics.

Chapter 5

Leveraging symmetries in robot learning

As discussed in [section 2.2](#) and [chapter 8](#) and [Thrun et al. \(2005\)](#), the core learning problems in robotics, namely dynamics modelling, state estimation, and optimal control, can in stochastic settings be formulated as learning problems whose goal is to approximate a target conditional probability distribution. For dynamics modelling and state estimation, the target distributions are $\mathbb{P}_{o|a\vec{o}}^{\Delta t}$ (eq. 2.15), the distribution of next-state observations conditioned on past observations and the present action, and $\mathbb{P}_{s|\vec{o}}^{\Delta t}$ (eq. 7.1), the distribution of the present state conditioned on past observations. Both can be approximated from system trajectories via supervised or self-supervised learning techniques (see eqs. 2.16 and 2.17). For optimal control, the target distribution is either an optimal policy π_* or an expert control distribution π_{exp} , which can be approximated via reinforcement or imitation learning.

In [chapters 6 to 8](#), we characterize how dynamics modelling, optimal control, and state estimation become symmetry-constrained learning problems whenever the robot features a spatial symmetry group \mathbb{G} (see [section 3.2](#) and [fig. 1.2](#)), thereby rendering the corresponding target conditional distributions \mathbb{G} -invariant (see [definition 2.17](#)). The core argument of this thesis is that, in such cases, one should leverage these symmetry priors through constrained optimization techniques to improve the quality of the target-distribution approximation, as well as the sample efficiency and generalization of the learning process.

Before formulating dynamics modelling, optimal control, and state estimation as symmetry-constrained learning problems in [chapters 6 to 8](#), this chapter presents the core theoretical result used to justify the exploitation of symmetry priors in these applications. Summarized in [prop. 5.1](#), this result states that if the target conditional distribution is \mathbb{G} -invariant, then the family of optimal parametric models approximating it is a subset of the family of \mathbb{G} -invariant conditional distributions. In other words, any optimal model is \mathbb{G} -invariant, and therefore the parametric approximation should be sought within the family of \mathbb{G} -invariant conditional distributions by means of constrained optimization.

Although this result may appear simple, it provides a clear criterion for determining when leveraging symmetry priors can improve the performance of learned models in practice. As discussed in [section 5.2](#), we argue that this criterion is satisfied by the vast majority of robot learning applications. Finally, in [section 5.3](#), we provide practical guidance on how to exploit these priors in practice.

5.1 Approximating \mathbb{G} -invariant conditional probability distributions

This section studies the learning problem of approximating a target conditional probability distribution $\mathbb{P}_{\mathbf{y}|\mathbf{x}} : \Sigma_{\mathcal{Y}} \times \mathcal{X} \rightarrow [0, 1]$ with a parametric model $\mathbb{P}_{\boldsymbol{\theta}} : \Sigma_{\mathcal{Y}} \times \mathcal{X} \rightarrow [0, 1]$, given two random variables \mathbf{y} and \mathbf{x} with marginal distributions $\mathbb{P}_{\mathbf{y}}$ and $\mathbb{P}_{\mathbf{x}}$. This problem can be formalized as:

$$\arg \min_{\boldsymbol{\theta}} \mathcal{L}_{\text{err}}(\mathbb{P}_{\boldsymbol{\theta}}) := \mathbb{E}_{\mathbf{x}} [D(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot|\mathbf{x}) \| \mathbb{P}_{\boldsymbol{\theta}}(\cdot|\mathbf{x}))]. \quad (5.1)$$

where $D : (\mathcal{Y}, \Sigma_{\mathcal{Y}}) \times (\mathcal{Y}, \Sigma_{\mathcal{Y}}) \rightarrow [0, \infty]$ is a divergence between probability measures on $(\mathcal{Y}, \Sigma_{\mathcal{Y}})$ —such as the Kullback-Leibler or other f-divergences—, with $\Sigma_{\mathcal{Y}}$ denoting the sigma-algebra on \mathcal{Y} , and $\mathbb{E}_{\mathbf{x}}$ denoting expectation with respect to the marginal distribution $\mathbb{P}_{\mathbf{x}}$.

Symmetry priors Crucially, we will study this problem under the presence of symmetry priors in the form of the \mathbb{G} -invariance of the marginal distribution $\mathbb{P}_{\mathbf{x}}$ (definition 2.16), i.e.,

$$\mathbb{P}_{\mathbf{x}}(\mathbb{X}) = \mathbb{P}_{\mathbf{x}}(g \triangleright_{\mathcal{X}} \mathbb{X}), \quad \forall g \in \mathbb{G}, \mathbb{X} \in \Sigma_{\mathcal{X}}. \quad \iff \quad \begin{array}{ccc} \Sigma_{\mathcal{X}} & \xrightarrow{\triangleright_{\mathcal{X}}} & \Sigma_{\mathcal{X}} \\ & \searrow \mathbb{P}_{\mathbf{x}} & \downarrow \mathbb{P}_{\mathbf{x}} \\ & & [0, 1] \end{array} \quad (5.2)$$

and the \mathbb{G} -invariance of the target conditional distribution $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ (definition 2.17), i.e.,

$$\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) = \mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_{\mathcal{Y}} \mathbb{Y} | g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \iff \quad \begin{array}{ccc} \Sigma_{\mathcal{Y}} \times \mathcal{X} & \xrightarrow{(\triangleright_{\mathcal{Y}}, \triangleright_{\mathcal{X}})} & \Sigma_{\mathcal{Y}} \times \mathcal{X} \\ & \searrow \mathbb{P}_{\mathbf{y}|\mathbf{x}} & \downarrow \mathbb{P}_{\mathbf{y}|\mathbf{x}} \\ & & [0, 1] \end{array} \quad (5.3)$$

where \mathbb{G} is a compact symmetry group (definition 2.1), and $\triangleright_{\mathcal{Y}}$ and $\triangleright_{\mathcal{X}}$ denote the group actions of \mathbb{G} on the vector spaces \mathcal{Y} and \mathcal{X} , respectively (see definition 2.2). An illustration of these symmetry priors for scalar random variables under the reflection symmetry group \mathbb{C}_2 is presented in fig. 5.1.

Remark: As we will see in chapters 6 to 8, in the context of robot dynamics modelling, state estimation, and optimal control, the random variable \mathbf{x} corresponds to the history of state observations and actions, and the random variable \mathbf{y} corresponds to the next-state observation, the present state, or the present action, respectively. For robotic systems with spatial symmetry group \mathbb{G} (section 3.2), the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ follows from the symmetries of classical mechanics, while the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ depends on the assumptions on the excitation of the system by a given control policy.

Considering these symmetry priors, let $\mathbb{P}_{\boldsymbol{\theta}}^{\mathbb{G}}$ denote a \mathbb{G} -invariant version of the parametric model $\mathbb{P}_{\boldsymbol{\theta}}$, constructed via group averaging as follows:

$$\mathbb{P}_{\boldsymbol{\theta}}^{\mathbb{G}}(\mathbb{Y}|\mathbf{x}) := \int_{\mathbb{G}} \mathbb{P}_{\boldsymbol{\theta}}(g \triangleright_{\mathcal{Y}} \mathbb{Y} | g \triangleright_{\mathcal{X}} \mathbf{x}) d\lambda(g), \quad (5.4)$$




Figure 5.1 Illustration of the symmetry priors in eqs. 5.2 and 5.3 for scalar random variables under the reflection symmetry group $\mathbb{C}_2 := \{e, g_r \mid g_r^2 = e\}$. The first two panels show \mathbb{C}_2 -invariant marginals \mathbb{P}_x and \mathbb{P}_y , while the last two panels show the symmetry-related conditional distributions $\mathbb{P}_{y|x}(\cdot | x_0)$ and $\mathbb{P}_{y|x}(\cdot | g_r \triangleright x_0)$.

such that $\mathbb{P}_\theta^{\mathbb{G}}$ satisfies the \mathbb{G} -invariance constraints of eq. 5.3 by construction. Here, λ denotes the normalized Haar measure on the compact group \mathbb{G} , which for finite symmetry groups reduces to the counting measure, reducing the integral in eq. 5.4 to a sum over group elements;

$$\mathbb{P}_\theta^{\mathbb{G}}(Y|x) := \frac{1}{|\mathbb{G}|} \sum_{g \in \mathbb{G}} \mathbb{P}_\theta(g \triangleright_y Y | g \triangleright_x x).$$

It can be shown, under mild natural assumptions on the divergence D , that the approximation error of the \mathbb{G} -invariant model $\mathcal{L}_{\text{err}}(\mathbb{P}_\theta^{\mathbb{G}})$ is always less or equal to that of the unconstrained model $\mathcal{L}_{\text{err}}(\mathbb{P}_\theta)$. This result is summarized in the following proposition.

Proposition 5.1. *Let (x, y) be two \mathbb{G} -invariant random variables (definition 2.16) with a \mathbb{G} -invariant conditional probability distribution $\mathbb{P}_{y|x} : \Sigma_y \times \mathcal{X} \rightarrow [0, 1]$ (definition 2.17). Consider any parametric model $\mathbb{P}_\theta : \Sigma_y \times \mathcal{X} \rightarrow [0, 1]$ and its corresponding symmetry-constrained version $\mathbb{P}_\theta^{\mathbb{G}}$ defined by eq. 5.4. If the approximation quality metric \mathcal{L}_{err} , as defined in eq. 5.1, is measured by a divergence D that is convex in the second argument, then:*

$$\mathcal{L}_{\text{err}}(\mathbb{P}_\theta^{\mathbb{G}}) \leq \mathcal{L}_{\text{err}}(\mathbb{P}_\theta), \quad (5.5)$$

Proof. Let $x \in \mathcal{X}$ be a realization of the random variable x . We begin by analyzing the pointwise approximation-quality term at an arbitrary conditioning point $x = x$, defined as the divergence between the true conditional distribution and the symmetry-constrained model:

$$\begin{aligned} \mathcal{L}_{\text{err},x}(\mathbb{P}_\theta^{\mathbb{G}}) &:= D\left(\mathbb{P}_{y|x}(\cdot | x) \parallel \mathbb{P}_\theta^{\mathbb{G}}(\cdot | x)\right). \\ &= D\left(\mathbb{P}_{y|x}(\cdot | x) \parallel \int_{\mathbb{G}} \mathbb{P}_\theta(g \triangleright_y \cdot | g \triangleright_x x) d\lambda(g)\right). \quad (\text{by eq. 5.4}) \end{aligned}$$

Since the divergence D is assumed to be convex in its second argument, Jensen's inequality yields an upper bound in which the divergence of the average is bounded by the average of the divergences:

$$\mathcal{L}_{\text{err},x}(\mathbb{P}_\theta^{\mathbb{G}}) \leq \int_{\mathbb{G}} D\left(\mathbb{P}_{y|x}(\cdot | x) \parallel \mathbb{P}_\theta(g \triangleright_y \cdot | g \triangleright_x x)\right) d\lambda(g).$$

For \mathbb{G} -invariant random variables (definition 2.16), the probability space remains invariant under symmetry transformations; therefore, the divergence D is invariant under the group action, i.e.,

$D(P(\cdot)\|Q(\cdot)) = D(P(g \triangleright_y \cdot)\|Q(g \triangleright_y \cdot))$. Hence,

$$\begin{aligned} D\left(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot|\mathbf{x}) \left\| \mathbb{P}_{\theta}(g \triangleright_y \cdot | g \triangleright_x \mathbf{x})\right.\right) &= D\left(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(g \triangleright_y \cdot | g \triangleright_x \mathbf{x}) \left\| \mathbb{P}_{\theta}(g \triangleright_y \cdot | g \triangleright_x \mathbf{x})\right.\right) \\ &= D\left(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot | g \triangleright_x \mathbf{x}) \left\| \mathbb{P}_{\theta}(\cdot | g \triangleright_x \mathbf{x})\right.\right). \quad (\text{by eq. 2.32}) \end{aligned}$$

Substituting this bound into the expected-loss definition in eq. 5.1 gives:

$$\begin{aligned} \mathcal{L}_{\text{err}}(\mathbb{P}_{\theta}^{\mathbb{G}}) &= \mathbb{E}_{\mathbf{x}}[\mathcal{L}_{\text{err},\mathbf{x}}(\mathbb{P}_{\theta}^{\mathbb{G}})] = \int_{\mathcal{X}} \mathcal{L}_{\text{err},\mathbf{x}}(\mathbb{P}_{\theta}^{\mathbb{G}}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \\ &\leq \int_{\mathcal{X}} \int_{\mathbb{G}} D(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot | g \triangleright_x \mathbf{x}) \left\| \mathbb{P}_{\theta}(\cdot | g \triangleright_x \mathbf{x})\right.) d\lambda(g) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}). \end{aligned}$$

Then, using the \mathbb{G} -invariance of the marginal measure $\mathbb{P}_{\mathbf{x}}$ (eq. 5.2), we obtain:

$$\begin{aligned} \mathcal{L}_{\text{err}}(\mathbb{P}_{\theta}^{\mathbb{G}}) &\leq \int_{\mathbb{G}} \left(\int_{\mathcal{X}} D(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot | g \triangleright_x \mathbf{x}) \left\| \mathbb{P}_{\theta}(\cdot | g \triangleright_x \mathbf{x})\right.) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \right) d\lambda(g) \\ &= \int_{\mathbb{G}} \left(\int_{\mathcal{X}} D(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot | \mathbf{x}) \left\| \mathbb{P}_{\theta}(\cdot | \mathbf{x})\right.) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \right) d\lambda(g) \\ &= \int_{\mathbb{G}} \mathcal{L}_{\text{err}}(\mathbb{P}_{\theta}) d\lambda(g) = \mathcal{L}_{\text{err}}(\mathbb{P}_{\theta}). \end{aligned}$$

This completes the proof. \square

It is fundamental to develop a strong intuition for [prop. 5.1](#), as it forms the core argument justifying symmetry exploitation in robot dynamics modelling, estimation, and control (see [chapters 6 to 8](#)). This proposition is, in a sense, restating *the very, very obvious*: if your goal is to approximate a \mathbb{G} -invariant conditional probability distribution, you should search for your parametric approximation within the family of \mathbb{G} -invariant conditional distributions. Given that, for any arbitrary model, its \mathbb{G} -invariant version is likely a better approximation. In other words, you should approach your learning problem as a constrained optimization problem of the form:

$$\begin{aligned} \arg \min_{\theta} \quad & \mathcal{L}_{\text{err}}(\mathbb{P}_{\theta}) := \mathbb{E}_{\mathbf{x}} [D(\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\cdot | \mathbf{x}) \left\| \mathbb{P}_{\theta}(\cdot | \mathbf{x})\right.)], \\ \text{s.t.} \quad & \mathbb{P}_{\theta}(\mathbb{Y} | \mathbf{x}) = \mathbb{P}_{\theta}(g \triangleright_y \mathbb{Y} | g \triangleright_x \mathbf{x}), \quad \forall g \in \mathbb{G}, \mathbf{x} \in \mathcal{X}, \mathbb{Y} \in \Sigma_{\mathbb{Y}}. \end{aligned} \quad (5.6)$$

5.2 When to leverage symmetry priors in robot learning?

In a nutshell, [prop. 5.1](#) states that if $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ and $\mathbb{P}_{\mathbf{x}}$ are \mathbb{G} -invariant, then approximating $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ becomes a symmetry-constrained optimization problem of the form of eq. 5.6. Therefore, the question of whether to leverage symmetry priors reduces to whether the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ and $\mathbb{P}_{\mathbf{x}}$ holds in practice.

[Chapters 6 to 8](#) provide strong arguments justifying that the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ holds in dynamics modelling, state estimation, and optimal control, as a consequence of the symmetries of classical mechanics, for every operational environment and control task. In practice, however, the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ depends entirely on the regime of operation of the robotic system. Therefore, the decision to

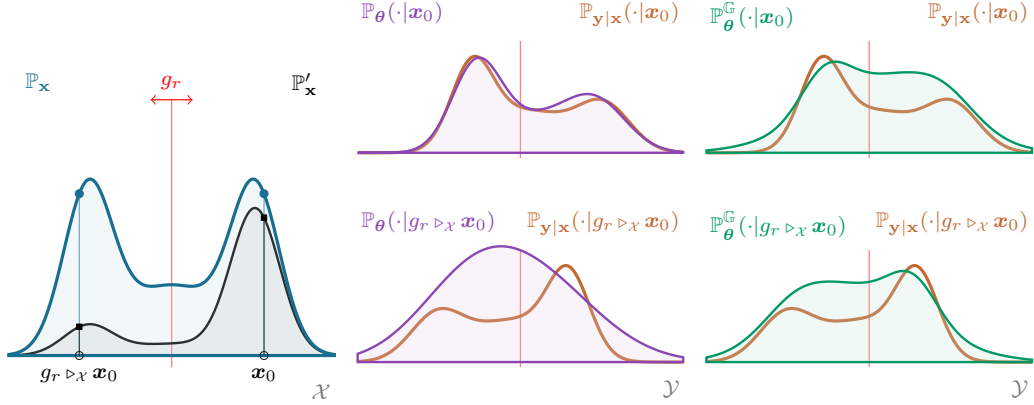


Figure 5.2 Illustration of a parametric model approximation under a symmetry prior. Left: The marginal distribution of conditioning variables \mathbb{P}_x and its non- \mathbb{G} -invariant empirical approximation \mathbb{P}'_x . Middle: A parametric conditional model \mathbb{P}_θ trained to approximate the ground-truth distribution (dashed) evaluated at two symmetry-related conditioning points x_0 and $g_r \triangleright_x x_0$. Right: The corresponding \mathbb{G} -invariant conditional model $\mathbb{P}_\theta^{\mathbb{G}}$, constructed by averaging the parametric model over the group orbits, which yields a significantly better approximation of the true conditional distribution across the orbit space.

leverage symmetry priors in robot learning hinges on whether the \mathbb{G} -invariance of \mathbb{P}_x is satisfied in practice.

To understand this point, we begin from a purely theoretical perspective, noting that the \mathbb{G} -invariance of \mathbb{P}_x is required in the last step of the proof of [prop. 5.1](#) to obtain the final bound $\mathcal{L}_{\text{err}}(\mathbb{P}_\theta^{\mathbb{G}}) \leq \mathcal{L}_{\text{err}}(\mathbb{P}_\theta)$, which is the key result of the proposition. Violating the \mathbb{G} -invariance of \mathbb{P}_x would imply that this bound is not guaranteed to hold, and therefore the invariant model $\mathbb{P}_\theta^{\mathbb{G}}$ may have a worse modelling error than the unconstrained model \mathbb{P}_θ .

To illustrate this fact intuitively, we present [fig. 5.2](#), which shows the performance of a parametric model \mathbb{P}_θ and its \mathbb{G} -invariant version $\mathbb{P}_\theta^{\mathbb{G}}$ in approximating a \mathbb{G} -invariant target conditional distribution $\mathbb{P}_{y|x}$ for the two random variables x and y of [fig. 5.1](#). In this figure, the parametric model is shown to have accurate approximations of the target distribution for conditioning points $x > 0$ and deteriorated approximation quality for conditioning points $x < 0$. Whenever the \mathbb{G} -invariance of the marginal distribution \mathbb{P}_x holds, the \mathbb{G} -invariant model $\mathbb{P}_\theta^{\mathbb{G}}$ features a lower modelling error ([eq. 5.5](#)), as the predictions on the left half of the conditioning space are improved by the symmetry constraint. However, if the \mathbb{G} -invariance of \mathbb{P}_x is violated, as shown in the left panel of [fig. 5.2](#) for \mathbb{P}'_x , and the likelihood of the negative conditioning events is negligible compared to the positive conditioning events, then the parametric model \mathbb{P}_θ , which features good approximation quality for the positive conditioning events, will yield a better overall approximation. This occurs due to the expectation over \mathbb{P}'_x in the definition of the approximation error \mathcal{L}_{err} in [eq. 5.1](#).

As we will see in [chapters 6 to 8](#), in the context of robot dynamics modelling, state estimation, and optimal control, the marginal distribution \mathbb{P}_x corresponds to the marginal distribution of past-state observations and actions, linked to the marginal distribution of robot states. Hence, the \mathbb{G} -invariance of \mathbb{P}_x can be justified if, in the expected operational regime of the robot, every world state $s \in \mathcal{S}$ is assumed to be equally likely to be observed as any element of its group orbit ([definition 2.4](#)), that is, as any of its symmetry-related states, $\mathbb{G} \triangleright_S s = \{g \triangleright_S s \mid g \in \mathbb{G}\}$. We posit that this is a natural assumption in the vast majority of robotic applications, but it inevitably depends on the chosen regime of operation. We present two application examples that violate this assumption:

1. **Legged locomotion with a pre-defined forward direction:** Consider the control application of legged locomotion with the MiniCheetah robot (fig. 1.3). The robot is capable of locomoting with the torso facing forward/backward, left/right, and upwards/downwards. For this control task, the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ would entail that, during normal operational conditions, we expect the robot to be equally likely to locomote in any of these torso configurations. However, if the practitioner chooses to only operate the robot with the torso in a single of these configurations, the empirical distribution of states $\mathbb{P}'_{\mathbf{x}}$ would violate the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$. Symmetry exploitation in this case would enable the dynamics/state-estimation/control model to operate comparably well across these 8 modes of locomotion, but this performance will never be used in practice, as the practitioner will only operate the robot in one of these modes. Furthermore, leveraging the symmetry priors might hamper performance in the chosen operational mode.
2. **Ambidextrous and non-ambidextrous manipulation:** Consider the control application of bimanual robot manipulation with the dual-arm robot in figs. 8.1 and 8.12. For this application, the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ implies that the robot is expected to tackle ambidextrous manipulation, meaning that any manipulation task and its reflected counterpart are equally likely to be observed in practice (see fig. 8.1). If, however, the practitioner wants to learn a non-ambidextrous control policy, say, learning to write only with the right hand, then the empirical distribution of states $\mathbb{P}'_{\mathbf{x}}$ would violate the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$. In this case, symmetry exploitation would enable the control model to perform well for both right-handed and left-handed writing, but only the right-handed writing performance will be used in practice. Furthermore, leveraging the symmetry priors might hamper performance in the chosen operational mode.

5.3 How to leverage symmetry priors in robot learning?

If in your control application, the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ is a natural assumption, there exist two main approaches to leverage symmetry priors in the constrained optimization problem in eq. 5.6, which are not mutually exclusive and can be combined for even better results:

- * **Data augmentation:** To empirically approximate \mathcal{L}_{err} , you should use a dataset of pairs of realizations of \mathbf{x} and \mathbf{y} sampled from their joint distribution, $\mathbb{D} := \{(\mathbf{x}_i, \mathbf{y}_i) \sim \mathbb{P}_{\mathbf{xy}}\}_{i=1}^N$. If both $\mathbb{P}_{\mathbf{x}}$ and $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ are \mathbb{G} -invariant, then so is their joint distribution $\mathbb{P}_{\mathbf{xy}}$ (see prop. 2.3). You can exploit this by improving the empirical coverage of your dataset via data augmentation: apply all group elements to each dataset sample, increasing the effective sample size from N to (approximately) $|\mathbb{G}|N$, yielding the \mathbb{G} -augmented dataset/minibatch:

$$\mathbb{D}^{\mathbb{G}} := \{(g \triangleright_{\mathbf{x}} \mathbf{x}_i, g \triangleright_{\mathbf{y}} \mathbf{y}_i) \sim \mathbb{P}_{\mathbf{xy}}\}_{i=1}^N, g \in \mathbb{G}. \quad (5.7)$$

By leveraging data-augmentation in every minibatch you use to approximate \mathcal{L}_{err} empirically, you obtain an empirical estimator with strictly lower variance (Wang et al., 2022b). In other words, you get better loss and gradient estimates.

- * **Invariance constraint enforcement:** In addition to data-augmentation, the \mathbb{G} -invariance constraint in eq. 5.6 can be enforced via hard/explicit or soft/regularized constrained optimization (or a combination of both):

- ◆ **Explicit constraint enforcement:** The \mathbb{G} -invariance constraints can be explicitly enforced via projection methods (e.g., projected gradient descent (Nordenfors et al., 2023)) or

by reparametrization methods (e.g., by using invariant/equivariant NN architectures (Bronstein et al., 2021; Cesa et al., 2021; Kondor and Trivedi, 2018; Pertigkiozoglou et al., 2026; Ordoñez-Apaez et al., 2025, 2026; Weiler et al., 2021)), such that the model is guaranteed to satisfy \mathbb{G} -invariance *globally* by construction, even for out of distribution conditioning events.

Remark: Note that for a \mathbb{G} -equivariant/invariant model, data-augmentation is redundant, and should not provide any additional performance benefits.

- ◆ **Regularized constraint enforcement:** Alternatively, the \mathbb{G} -invariance constraint can be enforced via regularization (e.g., penalty, barrier and relaxation methods (Finzi et al., 2021; Pertigkiozoglou et al., 2024, 2026; Xie and Smidt, 2025)), such that the model is *encouraged* to satisfy the \mathbb{G} -invariance constraint *on the support of the training distribution*, but has no guarantees of satisfying the constraint for out of distribution conditioning events.

For these approaches one computes a metric of constraint violation, such as:

$$\mathcal{P}_{\text{inv-err}}(\mathbb{P}_\theta) := \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_x} \left[\int_{\mathbb{G}} \mathcal{D}(\mathbb{P}_{y|x}(\cdot|\mathbf{x}) \parallel \mathbb{P}_\theta(g \triangleright_y \cdot |g \triangleright_x \mathbf{x})) d\lambda(g) \right], \quad (5.8)$$

which for finite groups reduces to

$$\mathcal{P}_{\text{inv-err}}(\mathbb{P}_\theta) := \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_x} \left[\frac{1}{|\mathbb{G}|} \sum_{g \in \mathbb{G}} \mathcal{D}(\mathbb{P}_{y|x}(\cdot|\mathbf{x}) \parallel \mathbb{P}_\theta(g \triangleright_y \cdot |g \triangleright_x \mathbf{x})) \right], \quad (5.9)$$

and then optimize the regularized loss $\mathcal{L}_{\text{err}}(\mathbb{P}_\theta; \gamma) := \mathcal{L}_{\text{err}}(\mathbb{P}_\theta) + \gamma \mathcal{P}_{\text{inv-err}}(\mathbb{P}_\theta)$, where $\gamma > 0$ is a constraint violation penalty coefficient.

A practitioner’s guide to leveraging symmetry priors in robot learning

Although leveraging symmetry priors results in improved empirical performance of the learned models, different forms of symmetry exploitation require different levels of implementation effort and expertise. Therefore, we recommend the following order of approaches to leverage symmetry priors in practice:

1. **Data augmentation:** Begin with data augmentation, as it is the simplest approach to implement and only requires you to define the group action on \mathcal{Y} and \mathcal{X} . This will help you eliminate errors arising from incorrectly specifying the symmetry group and its action. Note that `morpho_symm` and `symm_learning` (see section 1.5.1) provide group actions for a wide variety of robotic systems, along with documentation and example code for implementing data augmentation in NN training pipelines.
2. **Constraint-violation regularization:** Once you observe improved model performance via data augmentation, proceed with regularized constraint enforcement. Its implementation is straightforward once the group actions are well defined.
3. **Explicit constraint enforcement:** Finally, leverage explicit constraint enforcement via equivariant NN architectures, which ensure the optimality constraint is satisfied *globally* by construction.

`symm_learning` provides a wide variety of equivariant **NN** modules and ready-to-use equivariant **NN** architectures, such as a BERT-type transformer (see [fig. 8.5](#)).

In the next three chapters, we will apply these symmetry constrained optimization techniques to the problems of robot dynamics modelling, state estimation, and optimal control, showing both theoretically and empirically the impact of leveraging symmetry priors in these applications.

Chapter 6

Morphological symmetries in dynamics modelling

This chapter studies the problem of learning stochastic robot dynamics in realistic contact-rich operational environments in the presence of state/morphological symmetries. Its main result is [prop. 6.1](#), which provides a theoretical justification for formulating dynamics modelling in this setting as a symmetry-constrained learning problem. In doing so, we provide a clear theoretical motivation for symmetry-exploitation in contact-rich robot dynamics modelling.

To establish this result, we first characterize in [section 6.1](#) how spatial symmetries of the underlying deterministic robot dynamics translate into symmetry constraints on the stochastic world dynamics that govern the robot’s evolution in contact-rich environments. Then, in [section 6.2](#), we characterize how symmetries of the partially observable world dynamics transfer to state-observation dynamics, which are the target of our modelling efforts. Finally, we establish conditions under which approximating the observation stochastic dynamics with a parametric learnable model is a constrained optimization problem, where the model must satisfy the \mathbb{G} -invariance constraints induced by the symmetries of the underlying dynamics.

6.1 Symmetric stochastic dynamics

As detailed in [section 2.2](#), epistemic and measurement uncertainty in contact-rich operational environments motivate modelling the world dynamics as a controlled state-dependent diffusion process of the form:

$$\mathbf{s}_{t+\Delta t} = \underbrace{F_{\Delta t}(\mathbf{s}_t, \mathbf{a}_t)}_{\text{Controlled deterministic drift}} + \underbrace{\mathbf{D}(\mathbf{s}_t)\sqrt{\Delta t}\boldsymbol{\xi}_t}_{\text{State-dependent diffusion}}, \quad \text{with } \boldsymbol{\xi}_t \sim \mathcal{N}(\mathbf{0}, \text{Id}_{|\mathcal{S}|}). \quad (6.1)$$

Here, $\mathbf{s} \in \mathcal{S}$ and $\mathbf{a} \in \mathcal{A}$ denote the world state and action variables, respectively (see [eq. 2.12](#)). The map $F_{\Delta t} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is a deterministic drift term that captures the known robot dynamics under observable forcing terms such as gravity, control torques, and nonlinear coupling (see [eq. 2.5](#)). The map $\mathbf{D} : \mathcal{S} \rightarrow \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ is a state-dependent diffusion matrix that models the uncertainty of partially observable forcing terms, including contact forces and friction (see [section 2.2](#) for details).

For a robotic system with morphological symmetry group \mathbb{G} (see [chapter 3](#)), the world state space \mathcal{S} and action space \mathcal{A} inherit symmetry transformations induced by the morphological group action on the robot state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$ and by the environmental group action on the world interaction state $\mathbf{z} \in \mathcal{Z}$ (see [eq. 4.1](#)):

$$\begin{aligned} \triangleright_{\mathcal{S}} : \mathbb{G} \times \mathcal{S} &\mapsto \mathcal{S}, & \triangleright_{\mathcal{A}} : \mathbb{G} \times \mathcal{A} &\mapsto \mathcal{A}, \\ (g, \mathbf{s}) &\mapsto g \triangleright_{\mathcal{S}} \mathbf{s} := \begin{bmatrix} g \triangleright_{\mathcal{Q}} \mathbf{q} \\ g \triangleright_{\mathcal{T}_q \mathcal{Q}} \dot{\mathbf{q}} \\ g \triangleright_{\mathcal{Z}} \mathbf{z} \end{bmatrix} & (g, \mathbf{a}) &\mapsto g \triangleright_{\mathcal{A}} \mathbf{a} \end{aligned} \quad (6.2)$$

Furthermore, the \mathbb{G} -equivariance of the robot's equations of motion ([corollary 3.1](#)) constrains the stochastic dynamics in [eq. 6.1](#). In particular, the deterministic drift $F_{\Delta t}$ and the state-dependent diffusion matrix \mathbf{D} become \mathbb{G} -equivariant maps satisfying:

$$g \triangleright_{\mathcal{S}} F_{\Delta t}(\mathbf{s}, \mathbf{a}) = F_{\Delta t}(g \triangleright_{\mathcal{S}} \mathbf{s}, g \triangleright_{\mathcal{A}} \mathbf{a}) \quad \text{and} \quad g \triangleright_{\mathcal{S}} \mathbf{D}(\mathbf{s})_{S \triangleleft g^{-1}} = \mathbf{D}(g \triangleright_{\mathcal{S}} \mathbf{s}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}. \quad (6.3)$$

The equivariance of $F_{\Delta t}$ follows directly from [corollary 3.1](#), while the matrix-equivariance of \mathbf{D} follows from the symmetry relation of the uncertainty distributions of forcing terms at symmetry-related states—that is, $\mathbf{D}(\mathbf{s}_t) \sqrt{\Delta t} \boldsymbol{\xi}_t$ and $\mathbf{D}(g \triangleright_{\mathcal{S}} \mathbf{s}_t) \sqrt{\Delta t} \boldsymbol{\xi}_t$, for all $g \in \mathbb{G}$. To understand this, consider [fig. 6.1-left](#): by the symmetries of Newtonian mechanics, the uncertainty over ground reaction forces at state \mathbf{s} has an equivalent uncertainty distribution at any symmetry-related state.

Together, the two equivariance constraints in [eq. 6.3](#) ensure that the state-conditional transition distribution $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ is \mathbb{G} -invariant under the symmetry transformations of the state and action spaces. This invariance is the defining property of symmetric stochastic dynamical systems, as captured in the following definition.

Definition 6.1 (Symmetric stochastic dynamical system). *Let $\mathbf{s} \in \mathcal{S}$ and $\mathbf{a} \in \mathcal{A}$ denote the state and action variables of a stochastic dynamical system with state-transition distribution $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$. The system dynamics are said to possess a symmetry group \mathbb{G} if $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ is invariant under the group action on the state and action spaces, that is, if it satisfies:*

$$\begin{aligned} \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(g \triangleright_{\mathcal{S}} \mathbb{S} | g \triangleright_{\mathcal{S}} \mathbf{s}, g \triangleright_{\mathcal{A}} \mathbf{a}) = \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbb{S} | \mathbf{s}, \mathbf{a}) \\ \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbb{S} \subseteq \mathcal{S}, \mathbf{a} \in \mathcal{A}. \end{aligned} \quad \iff \quad \begin{array}{ccc} \Sigma_{\mathcal{S}} \times \mathcal{S} \times \mathcal{A} & \xrightarrow{(\triangleright_{\mathcal{S}}, \triangleright_{\mathcal{S}}, \triangleright_{\mathcal{A}})} & \Sigma_{\mathcal{S}} \times \mathcal{S} \times \mathcal{A} \\ & \searrow \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t} & \downarrow \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t} \\ & & [0, 1] \end{array} \quad (6.4)$$

In a nutshell, [eq. 6.4](#) can be interpreted as the generalization of the equivariance of the robot's equations of motion (see [corollary 3.1](#)) to the flow of world-state distributions, as depicted in [fig. 6.1](#).

In the context of robot operation in contact-rich environments, following the stochastic dynamics in [eq. 6.1](#), the state-conditional probability distribution $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ is assumed to be Gaussian (see [section 2.2](#)). Then, [eq. 6.4](#) implies equivariance constraints on the drift and diffusion terms that parameterize the mean and covariance of the Gaussian distribution:

Corollary 6.1 (Equivariant conditional Gaussian transition distribution). *Given that $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ for [eq. 6.1](#) is Gaussian, invariance of the conditional distribution is equivalent to equivariance of the conditional mean and matrix-equivariance of the conditional covariance, that is:*

$$\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(\cdot | \mathbf{s}, \mathbf{a}) = \mathcal{N}(F_{\Delta t}(\mathbf{s}, \mathbf{a}), \mathbf{D}(\mathbf{s})\mathbf{D}(\mathbf{s})^{\top} \Delta t) \quad \text{s.t.} \quad F_{\Delta t} \text{ and } \mathbf{D} \text{ satisfy } \text{eq. 6.3}. \quad (6.5)$$

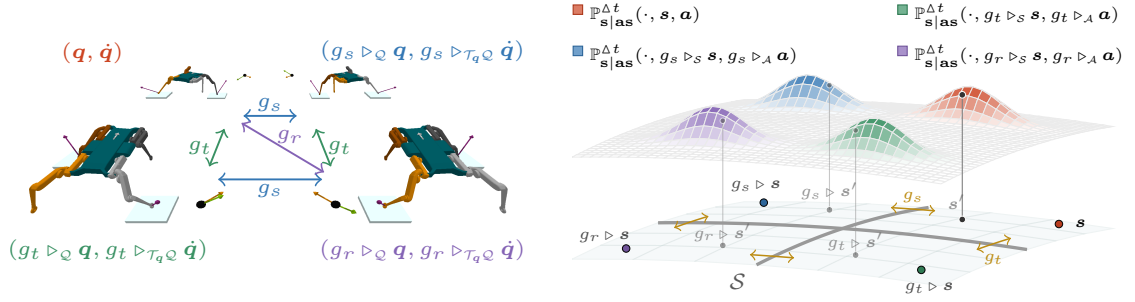


Figure 6.1 Illustration of the equivariant stochastic dynamics of the quadruped robot Solo (Griminger et al., 2020) featuring the Klein-four morphological symmetry group $\mathbb{K}_4 = \{e, g_s, g_r, g_t | g_s^2 = g_r^2 = g_t^2 = e, g_s \circ g_t = g_r\}$. The figure shows four symmetry-related world states (left), each acted upon by symmetry-related actions, and the resulting state-conditional transition distributions $\mathbb{P}_{s|as}^{\Delta t}$ for each state-action pair (right). The \mathbb{K}_4 -equivariant nature of the robot’s dynamics (corollary 3.1), together with the symmetry relations between uncertainty distributions of ground reaction contact forces, results in \mathbb{K}_4 -invariance of the transition distribution $\mathbb{P}_{s|as}^{\Delta t}$ under symmetry transformations of the state and action spaces (eq. 6.4), which implies that the transition distribution generated by (s, \mathbf{a}) is equivalent—up to a symmetry transformation of the domain $\mathcal{S} \times \mathcal{A}$ —to the transition distribution generated by $(g \triangleright_s s, g \triangleright_{\mathcal{A}} \mathbf{a})$ for any $g \in \mathbb{G}$.

Corollary 6.1 captures the fact that the next-state conditional Gaussian distribution given $(s, \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$ is equivalent to the distribution generated by the symmetry-transformed conditioning $(g \triangleright_s s, g \triangleright_{\mathcal{A}} \mathbf{a})$, for all $g \in \mathbb{G}$. This invariance property is illustrated in fig. 6.1-right for a quadruped robot, which features the Klein-four morphological symmetry group.

Symmetries under partial observability

Having established how symmetries of the underlying deterministic robot dynamics induce symmetry constraints on the stochastic world dynamics, we now characterize their implications under partial observability of the world state $s \in \mathcal{S}$.

Until now, we have studied stochastic world dynamics with uncertainty in forcing terms associated with partially observable contact forces and friction (eq. 6.1). In that setting, for any world state $s \in \mathcal{S}$ and sufficiently small Δt , the robot is assumed to be in a known contact mode (a set of contact surfaces with corresponding contact forces), and forcing-term uncertainty captures epistemic and measurement uncertainty in contact points/areas and contact wrenches for that mode. In practice, however, the world state is only partially observable, since the contact mode is usually not directly observable and must be inferred from sensor measurements from onboard sensors.

The set of proprioceptive and exteroceptive sensor measurements defines the observation space \mathcal{O} , sampled at each timestep t by an observation sensor model $\mathbb{P}_{o|s} : \Sigma_{\mathcal{O}} \times \mathcal{S} \mapsto [0, 1]$ that specifies the conditional distribution of state observations given world states (see eq. 2.14). From a modelling perspective, our aim is to model state-observation dynamics using only sensor measurements—i.e., present and past observations. Formally, this implies approximating the ground-truth observation dynamics defined by the following conditional observation transition distribution (see section 2.2):

$$\begin{aligned}
\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathbf{o}}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \vec{\mathbf{o}}_t) &= \int_{\mathcal{S}} \underbrace{\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \mathbf{s})}_{\text{Ground-truth observation dynamics}} \underbrace{\mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t}(d\mathbf{s}|\vec{\mathbf{o}}_t)}_{\text{Optimal state estimation}}, \quad \forall \quad \mathbb{O} \subseteq \mathcal{O}, \mathbf{a}_t \in \mathcal{A}, t \in \mathbb{R}, \\
&\quad \vec{\mathbf{o}}_t := (\mathbf{o}_0, \mathbf{o}_1, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}. \\
&= \int_{\mathcal{S}} \left(\underbrace{\int_{\mathcal{S}} \mathbb{P}_{\mathbf{o}|\mathbf{s}}(\mathbb{O}|\mathbf{s}')}_{\text{Observation model}} \underbrace{\mathbb{P}_{\mathbf{s}'|\mathbf{a}\mathbf{s}}^{\Delta t}(d\mathbf{s}'|\mathbf{a}_t, \mathbf{s})}_{\text{Ground-truth state dynamics}} \right) \mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t}(d\mathbf{s}|\vec{\mathbf{o}}_t)
\end{aligned} \tag{6.6}$$

Here, $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathbf{o}}}^{\Delta t} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \vec{\mathcal{O}} \rightarrow [0, 1]$ defines state-observation dynamics conditioned on the current action $\mathbf{a}_t \in \mathcal{A}$ and the history of past state observations $\vec{\mathbf{o}}_t \in \vec{\mathcal{O}}$ —with $\vec{\mathcal{O}} := \bigcup_{t \geq 0} \mathcal{O}^{t+1}$ denoting the space of past state-observation trajectories. This model is composed of the ground-truth observation dynamics $\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}$ —defining the distribution of next-state observations conditioned on the current partially observable state and action (see eq. 2.14), and an *optimal* state-estimation model, $\mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t}$, that infers the distribution of the current world state $\mathbf{s}_t \in \mathcal{S}$ from the full history of past state observations $\vec{\mathbf{o}}_t \in \vec{\mathcal{O}}$. Refer to section 2.2 for details.

We now study the conditions under which $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathbf{o}}}^{\Delta t}$ inherits the \mathbb{G} -invariance constraints of the underlying dynamics $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ (eq. 6.4). For this to happen, both $\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}$ and $\mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t}$ must also satisfy \mathbb{G} -invariance constraints. For ground-truth observation dynamics $\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}$ to inherit the \mathbb{G} -invariance constraints from the underlying dynamics $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ (eq. 6.4), we make the reasonable assumption that the observation model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ also satisfies \mathbb{G} -invariance constraints. Because this assumption is crucial for our theoretical results, we state it formally and discuss scenarios where it may not hold in practice.

Assumption 6.1 (\mathbb{G} -invariant observation model). *Let \mathcal{S} and \mathcal{O} be the state and observation spaces of a symmetric stochastic dynamical system with symmetry group \mathbb{G} (see definition 6.1). We assume that, for all $g \in \mathbb{G}$, the probability of observing any $\mathbf{o} \in \mathbb{O}$ at state \mathbf{s} equals the probability of observing any $\mathbf{o} \in g \triangleright_{\mathcal{O}} \mathbb{O}$ at the symmetry-transformed state $g \triangleright_{\mathcal{S}} \mathbf{s}$, that is:*

$$\begin{aligned}
\mathbb{P}(\mathbf{o} \in \mathbb{O} | \mathbf{s} = \mathbf{s}) = \mathbb{P}(\mathbf{o} \in g \triangleright_{\mathcal{O}} \mathbb{O} | \mathbf{s} = g \triangleright_{\mathcal{S}} \mathbf{s}) \quad \iff \quad \mathbb{P}_{\mathbf{o}|\mathbf{s}}(\mathbb{O}|\mathbf{s}) = \mathbb{P}_{\mathbf{o}|\mathbf{s}}(g \triangleright_{\mathcal{O}} \mathbb{O} | g \triangleright_{\mathcal{S}} \mathbf{s}) \\
\forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbb{O} \subseteq \mathcal{O}.
\end{aligned}$$

This is equivalent to assuming that \mathcal{O} inherits a group action $\triangleright_{\mathcal{O}} : \mathbb{G} \times \mathcal{O} \rightarrow \mathcal{O}$ (definition 2.2) and that the observation sensor model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ is \mathbb{G} -invariant (definition 2.17).

Justification: *In robotics, the observation space \mathcal{O} is defined by proprioceptive and exteroceptive sensor measurements composed of points, vectors, and higher-order geometric structures living in either 3-dimensional Euclidean space or the robot’s state space \mathcal{TQ} . Consequently, it is natural to assume that \mathcal{O} inherits the symmetry transformations of the world state space \mathcal{S} via the morphological and environmental group actions defined in eq. 4.1. Furthermore, if the observation sensor model is unbiased—a standard assumption in robotics—or at least has bias invariant under symmetry transformations, then the distributions of \mathbf{s} and \mathbf{o} satisfy prop. 2.3.*

Limitations: *This assumption is violated whenever the group action $\triangleright_{\mathcal{O}}$ is not properly defined. For instance, this occurs when \mathcal{O} contains RGB images from camera frames not aligned with the robot’s morphological symmetries, in which case rotated/reflected/translated images cannot be computed via invertible linear transformations of the original image. This limitation can be mitigated if the perception system outputs point clouds in \mathbb{R}^3 , which can be transformed via \mathbb{R}^3 (see eq. 4.1).*

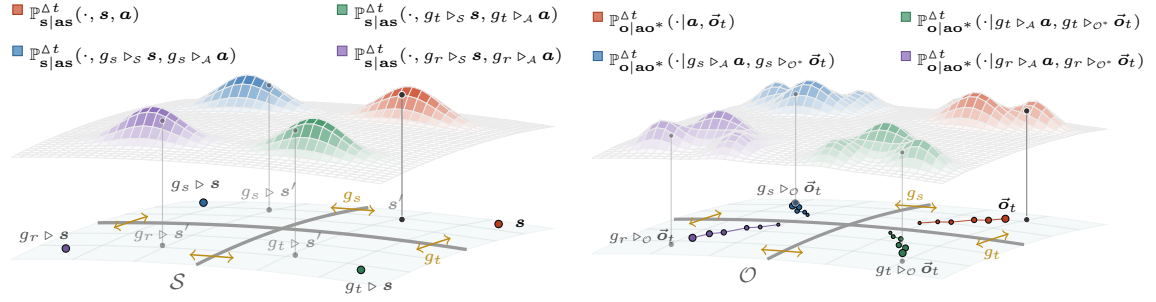


Figure 6.2 Illustration of symmetry transfer from the state transition distribution to the observation transition distribution for a quadruped with Klein-four morphological symmetry group \mathbb{K}_4 . Left: \mathbb{K}_4 -invariant state-conditional transition distributions $\mathbb{P}_{s|as}^{\Delta t}$ for symmetry-related state-action pairs (cf. fig. 6.1). Right: the corresponding attainable conditional observation transition distributions $\mathbb{P}_{o|ao}^{\Delta t}(\cdot | \mathbf{a}, \vec{o}_t)$, which are generally non-Gaussian and multimodal (see eq. 2.15). Under assumption 6.1, \mathbb{K}_4 -invariance of $\mathbb{P}_{s|as}^{\Delta t}$ together with \mathbb{G} -invariance of the observation model $\mathbb{P}_{o|s}$ implies \mathbb{K}_4 -invariance of $\mathbb{P}_{o|ao}^{\Delta t}$ (see eq. 6.8). Note that \mathcal{S} and \mathcal{O} are plotted as similar manifolds only for visualization purposes, these need not have the same dimensionality.

Under assumption 6.1, the ground-truth observation stochastic dynamics (see eq. 6.6) inherit, by construction, the \mathbb{G} -invariance of the state transition distribution $\mathbb{P}_{s|as}^{\Delta t}$. That is, they satisfy:

$$\mathbb{P}_{o|as}^{\Delta t}(\mathcal{O} | \mathbf{s}, \mathbf{a}) = \mathbb{P}_{o|as}^{\Delta t}(g \triangleright_{\mathcal{O}} \mathcal{O} | g \triangleright_{\mathcal{S}} \mathbf{s}, g \triangleright_{\mathcal{A}} \mathbf{a}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}, \mathcal{O} \subseteq \mathcal{O}. \quad (6.7)$$

In chapter 7, we will show that the optimal estimation model $\mathbb{P}_{s|\vec{o}}^{\Delta t}$ also satisfies \mathbb{G} -invariance constraints (see eq. 7.3), which together with eq. 6.7 implies that the full observation transition distribution $\mathbb{P}_{o|ao}^{\Delta t}$ is \mathbb{G} -invariant under symmetry transformations of the state and action spaces, as described in the following equation and depicted in fig. 6.2-right:

$$\mathbb{P}_{o|ao}^{\Delta t}(\mathcal{O} | \mathbf{a}, \vec{o}) = \mathbb{P}_{o|ao}^{\Delta t}(g \triangleright_{\mathcal{O}} \mathcal{O} | g \triangleright_{\mathcal{A}} \mathbf{a}, g \triangleright_{\vec{\mathcal{O}}} \vec{o}) \quad \Leftrightarrow \quad \begin{array}{ccc} \Sigma_{\mathcal{O}} \times \mathcal{A} \times \vec{\mathcal{O}} & \xrightarrow{(\triangleright_{\mathcal{O}}, \triangleright_{\mathcal{A}}, \triangleright_{\vec{\mathcal{O}}})} & \Sigma_{\mathcal{O}} \times \mathcal{A} \times \vec{\mathcal{O}} \\ & \searrow \mathbb{P}_{o|ao}^{\Delta t} & \downarrow \mathbb{P}_{o|ao}^{\Delta t} \\ & & [0, 1] \end{array} \quad (6.8)$$

Where the group action $\triangleright_{\vec{\mathcal{O}}} : \mathbb{G} \times \vec{\mathcal{O}} \rightarrow \vec{\mathcal{O}}$ is defined as the natural extension of $\triangleright_{\mathcal{O}}$ to the space of past observation trajectories $\vec{\mathcal{O}}$, that is, $g \triangleright_{\vec{\mathcal{O}}} \vec{o}_t := (g \triangleright_{\mathcal{O}} \mathbf{o}_0, g \triangleright_{\mathcal{O}} \mathbf{o}_1, \dots, g \triangleright_{\mathcal{O}} \mathbf{o}_t)$ for all $g \in \mathbb{G}$ and $\vec{o}_t \in \vec{\mathcal{O}}$ (see fig. 6.2).

Given that eq. 6.8 is the target conditional distribution we aim to approximate using numerical stochastic models (fig. 6.3), its \mathbb{G} -invariance represents a relevant inductive bias, which renders world dynamics modelling a symmetry-constrained optimization problem, as discussed in the next section.

6.2 Modelling symmetric dynamics under partial observability

Recall from section 2.2 that a parametric model of a stochastic dynamical system of the form stated in eq. 6.1 is defined as a parametric family of conditional probability distributions $\Phi_{\theta} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \mathcal{O}^* \mapsto [0, 1]$ that is learned to approximate the true observation transition distribution $\mathbb{P}_{o|ao}^{\Delta t}$ (eq. 6.8).

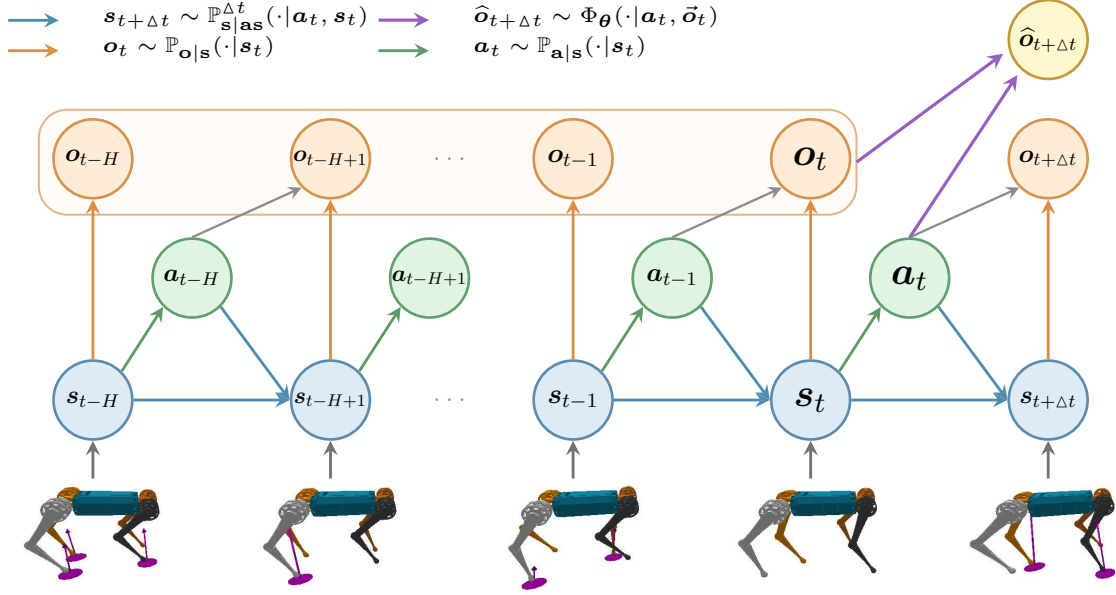


Figure 6.3 Graphical depiction of a stochastic dynamics model Φ_θ of robot dynamics. Given a world state s_t , a state observation o_t is sampled from the observation model $\mathbb{P}_{o|s}$ (2.14), and an action a_t is sampled from a given control policy $\mathbb{P}_{a|s}$. The next predicted observation is sampled from the stochastic dynamics model Φ_θ , conditioned on the history of past observations $\vec{o}_t = (o_{t_0}, \dots, o_t)$ and the current action a_t , where $t_0 = \max\{0, t - H + 1\}$ for some fixed history horizon $H \in \mathbb{N}$. Gray links from each past action to the next observation indicate that an observation at time t can carry information about the previous action a_{t-1} .

In this context, learning a dynamics model from trajectory data can be formulated as the following optimization problem:

$$\arg \min_{\theta} \mathcal{L}_{\text{err}}(\Phi_\theta) := \mathbb{E}_{\vec{o}, a \sim \mathbb{P}_{\vec{o}a}} \left[\mathcal{D} \left(\mathbb{P}_{o|a\vec{o}}^{\Delta t}(\cdot|a, \vec{o}) \parallel \Phi_\theta(\cdot|a, \vec{o}) \right) \right], \quad (6.9)$$

where $\mathcal{D} : \Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{O}} \rightarrow \mathbb{R}_{\geq 0}$ is a convex divergence metric—such as the KL divergence or Wasserstein distance—that quantifies the approximation error for each conditioning observation-history/action pair $(\vec{o}_t, a_t) \sim \mathbb{P}_{\vec{o}a}$. Here, $\mathbb{P}_{\vec{o}a}$ denotes the joint distribution over observation trajectories and actions induced by a control policy or a family of control policies. Importantly, this marginal distribution weights the modelling error of world state-action pairs according to their likelihood of being visited under nominal operation, thereby downweighting the error for pairs that are rarely encountered. Since our objective is to learn dynamics models that are effective for control, we assume that $\mathbb{P}_{\vec{o}a}$ is induced either by the family of all control policies the robot may execute in its operational environment or by a sufficiently rich, representative subset of such policies.

In the vast majority of control tasks involving robots with an environmental or morphological symmetry group \mathbb{G} (chapter 4), the robot is expected, during nominal operation, to execute both a given controlled trajectory and its symmetry-related counterparts with comparable or identical frequency. For example, in quadruped or humanoid loco-manipulation tasks with robots exhibiting bilateral morphological symmetry (see figs. 4.1 and 8.9), the robot is expected to walk in any direction and turn either way with similar frequency. From a modelling perspective, this motivates seeking models that perform consistently across all symmetry-related trajectories. Formally, this is equivalent

to assuming that joint distribution $\mathbb{P}_{\vec{\sigma}\mathbf{a}}$ is \mathbb{G} -invariant (definition 2.17), which can be interpreted as assuming that modelling error is evaluated under \mathbb{G} -equivariant excitation. This assumption is fundamental to the theoretical results and is therefore stated and discussed in detail.

Assumption 6.2 (\mathbb{G} -equivariant system excitation). *Let $\mathbb{P}_{\vec{\sigma}\mathbf{a}} : \Sigma_{\vec{\sigma}} \times \Sigma_{\mathcal{A}} \rightarrow [0, 1]$ be the joint distribution over observation trajectories and actions used to weight the modelling error in eq. 6.9 for different regions of $\mathcal{A} \times \vec{\mathcal{O}}$. We assume that $\mathbb{P}_{\vec{\sigma}\mathbf{a}}$ is \mathbb{G} -invariant, that is:*

$$\mathbb{P}_{\vec{\sigma}\mathbf{a}}(g \triangleright_{\vec{\sigma}} \mathbb{O}, g \triangleright_{\mathcal{A}} \mathbb{A}) = \mathbb{P}_{\vec{\sigma}\mathbf{a}}(\mathbb{O}, \mathbb{A}) \quad \forall g \in \mathbb{G}, \mathbb{O} \subseteq \vec{\mathcal{O}}, \mathbb{A} \subseteq \mathcal{A}.$$

As a result, the modelling error metric $\mathcal{L}_{\text{err}}(\Phi_{\theta})$ in eq. 6.9 is \mathbb{G} -invariant, meaning it assigns equal weight to the modelling error of any given conditioning pair $(\vec{\sigma}_t, \mathbf{a}_t)$ and to all symmetry-related conditioning pairs $(g \triangleright_{\vec{\sigma}} \vec{\sigma}_t, g \triangleright_{\mathcal{A}} \mathbf{a}_t)$ for every $g \in \mathbb{G}$.

Justification: *If we aim to learn a world dynamics models that perform well across arbitrary actuation for one or multiple control tasks, \mathbb{G} -equivariant excitation is a natural assumption. In the vast majority of robot control tasks, this is a natural assumption given that during nominal operation, a control trajectory and all its symmetry-related trajectories are expected to occur with comparable frequency. In loco-manipulation, this means expecting symmetry-related tasks to occur with similar frequency, as is typical in locomotion (figs. 3.2, 4.1, 8.9 and 8.11) and manipulation (figs. 8.1 and 8.12).*

Exceptions: *This assumption may be violated when one symmetric mode of operation is arbitrarily favored. For example, in bimanual dexterous manipulation (see figs. 8.1 and 8.12), if nominal operation requires handwriting exclusively with the right hand, then $\mathbb{P}_{\vec{\sigma}\mathbf{a}}$ would not be reflection-invariant, and \mathcal{L}_{err} would assign zero weight to modelling errors on left-hand writing trajectories. Similarly, for quadruped locomotion (fig. 1.3), although the robot can locomote with torso facing forwards or backwards and upwards or downwards (see animation), one torso configuration may be **arbitrarily** preferred. In such cases, the training distribution will cover only a subset of the support of the \mathbb{G} -invariant $\mathbb{P}_{\vec{\sigma}\mathbf{a}}$. The philosophy of this work is to avoid such arbitrary selection of a preferred symmetric mode of operation.*

The relevance of \mathbb{G} -equivariant excitation is two fold:

1. When empirically approximating \mathcal{L}_{err} using a dataset or batch of observation-history trajectories, $\mathbb{D} := \{(\vec{\sigma}_i, \mathbf{a}_i) \sim \mathbb{P}_{\vec{\sigma}\mathbf{a}}\}_{i=1}^N$, the \mathbb{G} -invariance of $\mathbb{P}_{\vec{\sigma}\mathbf{a}}$ motivates the use of data augmentation (see chapter 5). By applying all group elements to each sample, the effective sample size increases from N to $|\mathbb{G}|N$ (Wang et al., 2022b), leading to the estimator:

$$\widehat{\mathcal{L}}_{\text{err}}(\Phi_{\theta}) := \frac{1}{|\mathbb{G}|N} \sum_{i=1}^N \sum_{g \in \mathbb{G}} \mathcal{D}(\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\sigma}}^{\Delta t}(\cdot | g \triangleright_{\mathcal{A}} \mathbf{a}_i, g \triangleright_{\vec{\sigma}} \vec{\sigma}_i) \| \Phi_{\theta}(\cdot | g \triangleright_{\mathcal{A}} \mathbf{a}_i, g \triangleright_{\vec{\sigma}} \vec{\sigma}_i)) \quad (6.10)$$

2. If the modelling error is weighted uniformly across symmetry-related state-action trajectories, then for any parametric model Φ_{θ} , we can construct a \mathbb{G} -invariant model $\Phi_{\theta}^{\mathbb{G}}$ with equal or lower modelling error, i.e., $\mathcal{L}_{\text{err}}(\Phi_{\theta}^{\mathbb{G}}) \leq \mathcal{L}_{\text{err}}(\Phi_{\theta})$. For finite morphological symmetry groups, this invariant model is obtained by group averaging:

$$\Phi_{\theta}^{\mathbb{G}}(\mathbb{O}|\mathbf{a}, \vec{\sigma}) := \frac{1}{|\mathbb{G}|} \sum_{g \in \mathbb{G}} \Phi_{\theta}(g \triangleright_{\mathcal{O}} \mathbb{O} | g \triangleright_{\mathcal{A}} \mathbf{a}, g \triangleright_{\vec{\sigma}} \vec{\sigma}). \quad (6.11)$$

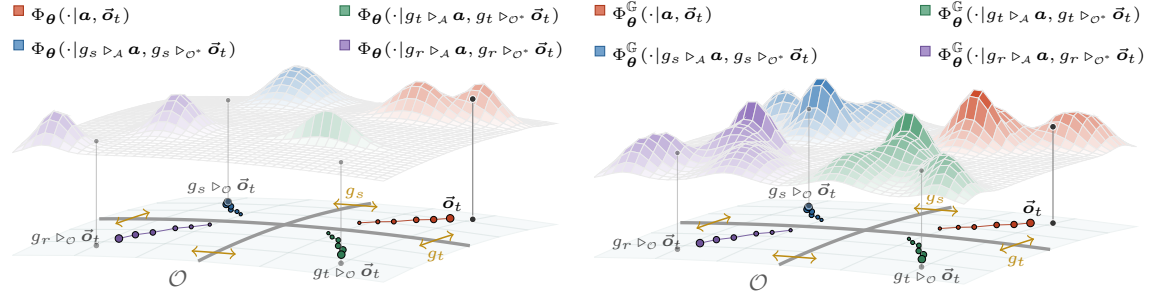


Figure 6.4 Illustration of model symmetrization under the Klein-four symmetry group. Left: an unconstrained model Φ_θ can assign non-equivalent conditional distributions across symmetry-related conditioning pairs $(\mathbf{a}, \vec{\sigma})$, $(g_s \triangleright_{\mathcal{A}} \mathbf{a}, g_s \triangleright_{\mathcal{O}} \vec{\sigma})$, $(g_t \triangleright_{\mathcal{A}} \mathbf{a}, g_t \triangleright_{\mathcal{O}} \vec{\sigma})$, and $(g_r \triangleright_{\mathcal{A}} \mathbf{a}, g_r \triangleright_{\mathcal{O}} \vec{\sigma})$. Right: after group averaging, the model $\Phi_\theta^{\mathbb{G}}$, constructed via eq. 6.11, yields equivalent conditional distributions for symmetry-related conditioning pairs, as a consequence of its \mathbb{G} -invariance constraint. Crucially, the modelling error of the symmetrized model is equal or lower than that of the original model, i.e., $\mathcal{L}_{\text{err}}(\Phi_\theta^{\mathbb{G}}) \leq \mathcal{L}_{\text{err}}(\Phi_\theta)$, whenever the modelling error is weighted uniformly across symmetry-related conditioning pairs (assumption 6.2).

This group averaging leverages the symmetry prior encoded in the invariance of the target observation transition kernel $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\sigma}}^{\Delta t}$, ensuring that $\Phi_\theta^{\mathbb{G}}$ is \mathbb{G} -invariant by construction (see fig. 6.4). Crucially, since a model can always be improved (or matched) by enforcing \mathbb{G} -invariance, the set of optimal solutions to eq. 6.9 is contained within the family of \mathbb{G} -invariant conditional distributions. This motivates reformulating world-dynamics modelling as a symmetry-constrained optimization problem. The result, which provides a clear theoretical justification for exploiting symmetry in contact-rich robot dynamics modelling, is formalized in the following proposition.

Proposition 6.1 (Symmetry-constrained partially observable dynamics modelling). *Let \mathcal{S} , \mathcal{A} , and \mathcal{O} denote the state, action, and observation spaces of a partially observable controlled state-dependent diffusion process (eq. 6.1), modelling robot dynamics in contact-rich environments. Let $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\sigma}}^{\Delta t} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \vec{\mathcal{O}} \rightarrow [0, 1]$ (eq. 6.6) be the observation stochastic dynamics to be approximated by a parametric model $\Phi_\theta : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \vec{\mathcal{O}} \rightarrow [0, 1]$. Assume that:*

1. *The world dynamics are symmetric under the action of a compact symmetry group \mathbb{G} , so the partially observable state dynamics $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ is \mathbb{G} -invariant (definition 6.1).*
2. *The observation model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ is \mathbb{G} -invariant (assumption 6.1). Consequently $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\sigma}}^{\Delta t}$ is \mathbb{G} -invariant (eq. 6.8).*
3. *The modelling error is measured via*

$$\mathcal{L}_{\text{err}}(\Phi_\theta) := \mathbb{E}_{\vec{\sigma}, \mathbf{a} \sim \mathbb{P}_{\vec{\sigma}\mathbf{a}}} \left[\mathcal{D} \left(\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\sigma}}^{\Delta t}(\cdot|\mathbf{a}, \vec{\sigma}) \parallel \Phi_\theta(\cdot|\mathbf{a}, \vec{\sigma}) \right) \right],$$

for some divergence metric $\mathcal{D} : \Sigma_{\mathcal{O}} \times \Sigma_{\mathcal{O}} \rightarrow [0, \infty)$ —assumed to be convex in its second argument—and averaged under a \mathbb{G} -invariant distribution of observation-history/action pairs $\mathbb{P}_{\vec{\sigma}\mathbf{a}}$ (assumption 6.2).

Then, for any parametric model Φ_θ , its symmetrized version $\Phi_\theta^{\mathbb{G}}$ defined in eq. 6.11 satisfies

$$\mathcal{L}_{\text{err}}(\Phi_\theta^{\mathbb{G}}) \leq \mathcal{L}_{\text{err}}(\Phi_\theta).$$

Therefore, modelling the world dynamics of such a symmetric partially observable stochastic dynamical system must be reformulated as the following symmetry-constrained learning problem:

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg \min_{\boldsymbol{\theta}} \mathcal{L}_{err}(\Phi_{\boldsymbol{\theta}}) \\ \text{s.t. } \Phi_{\boldsymbol{\theta}}(\mathbb{O}|\mathbf{a}, \vec{\boldsymbol{o}}) &= \Phi_{\boldsymbol{\theta}}(g_{\triangleright_{\mathcal{O}}} \mathbb{O} | g_{\triangleright_{\mathcal{A}}} \mathbf{a}, g_{\triangleright_{\vec{\mathcal{O}}}} \vec{\boldsymbol{o}}) \quad \forall g \in \mathbb{G}, \mathbf{a} \in \mathcal{A}, \mathbb{O} \subseteq \mathcal{O}, \vec{\boldsymbol{o}} \in \vec{\mathcal{O}}. \end{aligned} \quad (6.12)$$

Proof. The result follows directly from [prop. 5.1](#) by identifying $\mathbf{y} = \mathbf{o}_{t+\Delta t}$ and $\mathbf{x} = (\mathbf{a}_t, \vec{\boldsymbol{o}}_t)$, with group action on \mathbf{x} defined as $g_{\triangleright_{\mathcal{X}}} \mathbf{x} := (g_{\triangleright_{\mathcal{A}}} \mathbf{a}_t, g_{\triangleright_{\vec{\mathcal{O}}}} \vec{\boldsymbol{o}}_t)$ for all $g \in \mathbb{G}$. Note that the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ follows from [assumption 6.2](#), and the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ follows from [eq. 6.8](#). \square

As discussed in [chapter 5](#) the symmetry priors of [prop. 6.1](#) can be leveraged either by using data augmentation to improve the empirical coverage of the dataset of observation-history/action pairs used to approximate \mathcal{L}_{err} or by enforcing the \mathbb{G} -invariance constraint on the model, via soft penalization methods, or via hard architectural constraints.

Truncated history models Note that [prop. 6.1](#) makes few assumptions about the parametrization of the dynamics model $\Phi_{\boldsymbol{\theta}}$, and therefore applies to virtually any stochastic model parametrization. The result also holds for dynamics models that depend on a truncated history of past observations, as is the case in most practical settings (see [section 8.2](#) and [Chi et al. \(2024\)](#)). For an H -truncated history, the target distribution becomes $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\boldsymbol{o}}}^{\Delta t} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \mathcal{O}^H \rightarrow [0, 1]$, and correspondingly, the dynamics model becomes $\Phi_{\boldsymbol{\theta}} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \mathcal{O}^H \rightarrow [0, 1]$, where \mathcal{O}^H is the space of observation trajectories of length H (see [section 8.2](#)).

Remark: Section 6.2

This thesis does not empirically address the problem of world dynamics modelling directly; this remains an avenue for future work. The theoretical results in this section, however, provide a rigorous justification for exploiting symmetry in the modelling of contact-rich robot dynamics, and are included to offer a complete theoretical characterization of the implications of morphological symmetries for the core learning problems in robotics: modelling, estimation ([chapter 7](#)), and control ([chapter 8](#)). In particular, symmetry priors in estimation and control rely on [assumptions 6.1](#) and [6.2](#) and [prop. 6.1](#) as their theoretical foundation.

Notably, in [section 9.3.2](#), we address the problem of modelling the *closed-loop* dynamics of quadruped locomotion via approximations of the closed-loop transfer/Koopman operator ([definition 9.3](#)). In this context, [prop. 6.1](#) is used to justify the adoption of symmetry priors for operator approximation, demonstrating one of the first examples of robot dynamics modelling through the Koopman operator ([Ordoñez-Appraez et al., 2024](#)). Here, symmetry exploitation yields substantial empirical benefits in model prediction accuracy, data efficiency, and algorithmic stability.

Chapter 7

Morphological symmetries in state estimation

This chapter studies the implications of spatial symmetries for state estimation in contact-rich robotics. In [section 7.1](#), we cast state estimation as conditional-distribution learning under partial observability and characterize the symmetry priors inherited from the underlying robot dynamics and observation model. The core result is [prop. 7.1](#), which shows that, under natural assumptions, optimal state-estimation models lie within the family of \mathbb{G} -invariant conditional probability distributions. This yields state estimation as a symmetry-constrained learning problem.

To characterize how to leverage these symmetry priors in practice, in [section 7.2.1](#) we discuss deterministic inference, covering supervised regression and classification. Then, in [section 7.3](#), we address probabilistic inference and uncertainty quantification.

7.1 Symmetry-priors in optimal state estimation

As discussed in detail in [section 2.2](#) and [chapter 6](#), when modelling the dynamics of a robot in a contact-rich environment, our goal is to capture the evolution of the *world state* $s \in \mathcal{S}$ rather than just the robot's state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$. This is because the robot's evolution depends not only on its own configuration but also on the contact modes, contact forces, and the specific control task. Thus, the world state space $\mathcal{S} := \mathcal{TQ} \times \mathcal{Z}$ comprises both the robot's state space and the robot–environment state space \mathcal{Z} , which encodes contact information and control task state (see [section 2.2](#) and [chapter 8](#)).

By construction, the world state $s \in \mathcal{S}$ is only partially observable, for two main reasons:

- * The robot's state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$ is subject to measurement uncertainty in joint positions and velocities. Additionally, some state variables, such as the global position of a floating-base robot, may not be directly measured by sensors.

- ✱ The robot–environment state variable $z \in \mathcal{Z}$ is not directly measurable, as no sensor can reliably determine the current contact mode, including contact surfaces, contact forces, friction regimes, material deformation, and related quantities.

Thus, the problem of *state estimation*¹ can be formulated, in its most general form, as the approximation of the true conditional probability distribution of states given the history of past state observations, $\mathbb{P}_{\mathbf{s}|\vec{\sigma}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \rightarrow [0, 1]$, defined as:

$$\mathbb{P}_{\mathbf{s}|\vec{\sigma}}^{\Delta t}(\mathbb{S}|\vec{\sigma}_t) := \mathbb{P}(\mathbf{s}_t \in \mathbb{S} | \vec{\sigma}_t = \vec{\sigma}_t) \quad \forall \vec{\sigma}_t := (\mathbf{o}_0, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}, \mathbb{S} \subseteq \mathcal{S}, t \in \mathbb{R}, \quad (7.1)$$

Here, $\vec{\sigma}_t \in \vec{\mathcal{O}}$ denotes the history of past state observations up to time t . Each state observation $\mathcal{O} \ni \mathbf{o}_t \sim \mathbb{P}_{\mathbf{o}|\mathbf{s}}(\cdot | \mathbf{s}_t)$ consists of the collection of proprioceptive and exteroceptive sensor measurements available from onboard sensors at time t , assumed to be sampled from an observation model $\mathbb{P}_{\mathbf{o}|\mathbf{s}} : \Sigma_{\mathcal{O}} \times \mathcal{S} \rightarrow [0, 1]$. \mathcal{O} denotes the state-observation space, and $\vec{\mathcal{O}}$ is the space of all possible observation trajectories (see eq. 2.15).

Consequently, state estimation can be cast as a conditional distribution learning problem, where we approximate $\mathbb{P}_{\mathbf{s}|\vec{\sigma}}^{\Delta t}$ with a parametric model $\Upsilon_{\theta} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \rightarrow [0, 1]$:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{se}}(\Upsilon_{\theta}) := \arg \min_{\theta} \mathbb{E}_{\vec{\sigma} \sim \mathbb{P}_{\vec{\sigma}}} \left[\mathcal{D}(\mathbb{P}_{\mathbf{s}|\vec{\sigma}}^{\Delta t}(\cdot | \vec{\sigma}) \| \Upsilon_{\theta}(\cdot | \vec{\sigma})) \right], \quad (7.2)$$

Here, $\mathcal{D} : (\mathcal{S}, \Sigma_{\mathcal{S}}) \times (\mathcal{S}, \Sigma_{\mathcal{S}}) \rightarrow [0, \infty)$ is a convex divergence measure between probability distributions, such as the Kullback–Leibler divergence, and \mathcal{L}_{se} denotes the corresponding average divergence under the marginal distribution of state-observation histories $\mathbb{P}_{\vec{\sigma}}$. In practice, this marginal weights the estimation errors of Υ_{θ} for different conditioning observation histories $\vec{\sigma}$ according to their likelihood of occurrence during operation. That is, errors in estimating world states that are highly likely to occur are heavily penalized, while states never observed during operation are ignored in this metric.

Importantly, $\mathbb{P}_{\vec{\sigma}}$ is induced by the control policy and the true underlying dynamics, modeled by the state-transition probability distribution $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ (see eq. 2.13). This kernel gives the conditional distribution of next states at time $t + \Delta t$ when the robot takes action $\mathbf{a}_t \in \mathcal{A}$ in world state $\mathbf{s}_t \in \mathcal{S}$ at time t :

$$\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbb{S} | \mathbf{a}_t, \mathbf{s}_t) := \mathbb{P}(\mathbf{s}_{t+\Delta t} \in \mathbb{S} | \mathbf{a}_t = \mathbf{a}_t, \mathbf{s}_t = \mathbf{s}_t) \quad \forall \mathbb{S} \subseteq \mathcal{S}, \mathbf{a}_t \in \mathcal{A}, \mathbf{s}_t \in \mathcal{S}, t \in \mathbb{R},$$

State estimation of symmetric dynamical systems As detailed in [chapter 6](#), when modelling the contact-rich dynamics of a robot with an environmental or morphological symmetry group \mathbb{G} (see [chapter 4](#)), the state-transition distribution $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ is a \mathbb{G} -invariant conditional probability distribution ([definition 6.1](#)), satisfying:

$$\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(g \triangleright_{\mathcal{S}} \mathbf{s}' | g \triangleright_{\mathcal{S}} \mathbf{s}, g \triangleright_{\mathcal{A}} \mathbf{a}) = \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbf{s}' | \mathbf{s}, \mathbf{a}) \quad \forall g \in \mathbb{G}, \mathbf{s}, \mathbf{s}' \in \mathcal{S}, \mathbf{a} \in \mathcal{A},$$

where $\triangleright_{\mathcal{S}} : \mathbb{G} \times \mathcal{S} \rightarrow \mathcal{S}$ and $\triangleright_{\mathcal{A}} : \mathbb{G} \times \mathcal{A} \rightarrow \mathcal{A}$ denote the group actions of \mathbb{G} on the state and action spaces, respectively (see [fig. 6.1](#) and [eq. 6.2](#)). If, in addition, the observation model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ is also

¹In robotics, “state estimation” often refers to determining the robot’s position in a global/world reference frame. Here, we use the term in the classical dynamical-systems sense: estimating unobserved or partially observed state quantities.




Figure 7.1 Illustration of invariance of the conditional probability of world states given past observation histories, $\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}$ (eq. 7.3), for a quadruped robot with the morphological symmetry group $\mathbb{K}_4 := \{e, g_s, g_t, g_r | g_s^2 = g_t^2 = g_r^2 = e, g_s g_t = g_r, g_t g_s = g_r\}$, composed of two orthogonal reflections g_s and g_t and a 180° rotation g_r . For a given set of symmetry related noisy past observation histories $\mathbb{G}\vec{\sigma}_t := \{g \triangleright_{\vec{\sigma}} \vec{\sigma}_t | g \in \mathbb{K}_4\}$, the state probability distributions $\{\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}(\cdot | g \triangleright_{\vec{\sigma}} \vec{\sigma}_t) | g \in \mathbb{K}_4\}$ are related by the reflections and rotation of \mathbb{K}_4 to the state distribution $\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}(\cdot | \vec{\sigma}_t)$ induced by the original observation history $\vec{\sigma}_t$ (right). Here, $\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}(\cdot | \vec{\sigma}_t)$ denotes the uncertainty distribution over robot state *and* contact states and contact forces acting on the robot at state s_t (orange). The relation between the symmetry-related state distributions models the relationship between uncertainties of the model state and contact mode and forces at the symmetry related states (green, blue, purple).

\mathbb{G} -invariant—a natural assumption in robotics, discussed in detail in [assumption 6.1](#)—then $\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}$ is itself a \mathbb{G} -invariant conditional probability distribution (see [fig. 7.1](#)), satisfying:

$$\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}(g \triangleright_S \mathbb{S} | g \triangleright_{\vec{\sigma}} \vec{\sigma}) = \mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}(\mathbb{S} | \vec{\sigma}) \quad \forall g \in \mathbb{G}, \mathbb{S} \subseteq \mathcal{S}, \vec{\sigma} \in \vec{\mathcal{O}}, \quad (7.3)$$

This \mathbb{G} -invariance follows directly from the symmetries of the underlying robot dynamics ([chapter 3](#)), and formalizes the intuitive idea that the uncertainty distribution over the current world state $s_t \in \mathcal{S}$, conditioned on a history of past observations $\vec{\sigma}_t \in \vec{\mathcal{O}}$, transforms under a symmetry $g \in \mathbb{G}$ to the uncertainty distribution over the symmetry-related world state $g \triangleright_S s_t$, given the symmetry-transformed observation history $g \triangleright_{\vec{\sigma}} \vec{\sigma}_t$. For example, consider the world state $s_t \in \mathcal{S}$ of a quadruped robot, as illustrated in [fig. 7.1](#). The uncertainty in contact states and forces is modeled by the state distribution $\mathbb{P}_{\mathcal{S}|\vec{\sigma}}^{\Delta t}(\cdot | \vec{\sigma}_t)$, induced by a history of proprioceptive and exteroceptive observations $\vec{\sigma}_t$. For any symmetry-related state, such as $g_s \triangleright_S s_t$, the corresponding contact state and contact forces are the reflected versions of the original. Thus, [eq. 7.3](#) states that if we were to observe the reflected observation history $g_s \triangleright_{\vec{\sigma}} \vec{\sigma}_t$, the resulting uncertainty distribution over $g_s \triangleright_S s_t$ would be the reflection of the original distribution. This reasoning applies to any robotic system and any environmental or morphological symmetry group \mathbb{G} (see [figs. 3.1, 3.2 and 7.1](#)).

Similarly, as in the problem of modelling stochastic world dynamics ([chapter 6 and prop. 6.1](#)), the target conditional probability distribution to be modeled is \mathbb{G} -invariant due to the underlying symmetries of the robot dynamics and the observation model. Furthermore, we assume that the marginal distribution of observation histories $\mathbb{P}_{\vec{\sigma}}$ is \mathbb{G} -invariant—a natural assumption in the context of dynamics learning and state estimation, as discussed in detail in [assumption 6.2](#), since for the vast majority of robot tasks, we are interested in learning state estimation models that perform identically or equivalently across symmetry-related modes of operation. For example, in quadruped or humanoid loco-manipulation tasks with robots exhibiting bilateral reflection morphological symmetry (see

figs. 4.1 and 8.9), reflection invariance of $\mathbb{P}_{\vec{o}}$ implies that we aim to learn models that perform equally well for controlled trajectories of walking and turning in any direction (Mittal et al., 2024; Su et al., 2024; Ordoñez-Apaez et al., 2025; Yu et al., 2018).

The \mathbb{G} -invariance of both the target conditional distribution $\mathbb{P}_{\vec{o}}^{\Delta t}$ and the marginal distribution of observation histories $\mathbb{P}_{\vec{o}}$ (which weights the state estimation errors in eq. 7.2) justifies formulating state estimation for \mathbb{G} -symmetric dynamical systems as a constrained optimization problem. We state this result in the following proposition:

Proposition 7.1 (Symmetry-constrained state estimation learning). *Let \mathcal{S} , \mathcal{A} , and \mathcal{O} denote the state, action, and observation spaces of a partially observable stochastic dynamical system modelling robot dynamics in contact-rich environments. Let $\mathbb{P}_{\vec{o}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \rightarrow [0, 1]$ be the true conditional distribution of states given observation histories (eq. 7.1), and let $\Upsilon_{\theta} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \rightarrow [0, 1]$ be any stochastic state estimation model, parameterized by θ (eq. 7.2). Assume that:*

1. *The robot possesses an environmental or morphological compact symmetry group \mathbb{G} (chapter 4), which renders the state-transition distribution $\mathbb{P}_{\vec{o}}^{\Delta t}$ \mathbb{G} -invariant (definition 6.1).*
2. *The observation sensor model $\mathbb{P}_{\mathcal{O}|\mathcal{S}}$ is \mathbb{G} -invariant (assumption 6.1). Consequently, the target distribution $\mathbb{P}_{\vec{o}}^{\Delta t}$ is \mathbb{G} -invariant (eq. 7.3).*
3. *The state estimation error is measured via the loss:*

$$\mathcal{L}_{se}(\Upsilon_{\theta}) \leq \mathcal{L}_{se}(\Upsilon_{\theta}^{\mathbb{G}}) := \mathbb{E}_{\vec{o} \sim \mathbb{P}_{\vec{o}}} \left[\mathcal{D} \left(\mathbb{P}_{\vec{o}}^{\Delta t}(\cdot | \vec{o}) \parallel \Upsilon_{\theta}(\cdot | \vec{o}) \right) \right],$$

for some divergence metric $\mathcal{D} : (\mathcal{S}, \Sigma_{\mathcal{S}}) \times (\mathcal{S}, \Sigma_{\mathcal{S}}) \rightarrow [0, \infty)$ —that is be convex in its second argument.

4. *The state estimation error is averaged under a \mathbb{G} -invariant marginal distribution of observation histories $\mathbb{P}_{\vec{o}}$ (assumption 6.2).*

Then, the family of optimal state estimation models is a subset of the family of \mathbb{G} -invariant conditional probability distributions of states given observation histories. Hence, state estimation must be formulated as the following symmetry-constrained learning problem:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} \mathcal{L}_{se}(\Upsilon_{\theta}) \\ \text{s.t. } &\Upsilon_{\theta}(\mathbb{S} | \vec{o}) = \Upsilon_{\theta}(g \triangleright_{\mathcal{S}} \mathbb{S} | g \triangleright_{\vec{\mathcal{O}}} \vec{o}) \quad \forall g \in \mathbb{G}, \mathbb{S} \subseteq \mathcal{S}, \vec{o} \in \vec{\mathcal{O}}. \end{aligned} \tag{7.4}$$

Proof. The result follows directly from prop. 5.1 by identifying $\mathbf{y} = \mathbf{s}$ and $\mathbf{x} = \vec{o}$, with group action on \mathbf{x} defined as $g \triangleright_{\mathcal{X}} \mathbf{x} := g \triangleright_{\vec{\mathcal{O}}} \vec{o}_t$ for all $g \in \mathbb{G}$. The \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{x}}$ follows from (4), and the \mathbb{G} -invariance of $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ follows from (1). \square

As discussed in chapter 5 the symmetry priors of eq. 7.4 can be leveraged either by using data augmentation to improve the empirical coverage of the dataset of state and state-observation pairs used to approximate \mathcal{L}_{se} or by enforcing the \mathbb{G} -invariance constraint on the model, via soft penalization methods, or via hard architectural constraints.

Next, we discuss the implications of these symmetry priors in the case of deterministic state estimation (i.e., regression and classification of unobservable quantities) and probabilistic state estimation (i.e., conditional probability estimation and uncertainty quantification) for specific model parameterizations.

7.2 Deterministic inference

For many unobservable quantities of the world state $\mathbf{s}_t \in \mathcal{S}$ of a robotic system operating in contact-rich environments, it is often assumed that the conditional distribution of the target variables given the observation history is (or can be well approximated by) a Gaussian or another well-defined parametric probability distribution (e.g., a categorical distribution for classification tasks).

Under this assumption, the state estimation learning problem in eq. 7.2 reduces to a standard supervised regression or classification task. This section presents both theoretical and empirical implications of incorporating state and morphological symmetry priors in supervised regression and classification of unobservable components of the robot’s state. The analysis and results are adapted from our previous work [Ordoñez-Appraez et al. \(2023, 2025\)](#).

7.2.1 Supervised regression

Here, the target random variable of interest is a subset of the world state, $\mathbf{y} \subset \mathbf{s}$, taking values $\mathbf{y} \in \mathcal{Y} \subseteq \mathcal{S}$, and composed of unobservable world state quantities whose conditional distribution, given state observations, can reasonably be assumed to be approximately Gaussian:

$$\mathbb{P}_{\mathbf{y}|\vec{\sigma}} := \mathbb{P}(\mathbf{y} \in \cdot | \vec{\sigma} = \vec{\sigma}) \approx \mathcal{N}(\mathbf{f}(\vec{\sigma}), \Sigma(\vec{\sigma})),$$

where the mean $\mathbf{f} : \mathcal{O}^H \rightarrow \mathcal{Y}$ and covariance $\Sigma : \mathcal{O}^H \rightarrow \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$ are assumed to be non-linear functions of the truncated observation history $\vec{\sigma} := (\mathbf{o}_{t_0}, \dots, \mathbf{o}_t) \in \mathcal{O}^H$ (see [fig. 6.3](#)), and where $H \in \mathbb{N}$ is the truncated history length and $t_0 := \max(0, t - H)$.

A classical example of such an unobservable variable is the robot’s position and orientation in global coordinates, $\mathbf{y} = \mathbf{X}_B \in \mathbb{S}\mathbb{E}_3$, estimated from past IMU linear acceleration and angular velocity, lidar/camera measurements, and joint positions and velocities. Another common example is the ground reaction forces acting on the feet of a quadruped robot with point feet, $\mathbf{y} = [\mathbf{f}_{FR}, \mathbf{f}_{FL}, \mathbf{f}_{HR}, \mathbf{f}_{HL}] \in \mathbb{R}^{12}$ ([fig. 1.3](#)), which are commonly assumed to be Gaussian-distributed in the literature for short time discretizations Δt ([Camurri et al., 2017](#)) (although in reality they often exhibit skewed and multimodal distributions; see [section 7.3](#)).

In this setting, the problem of state estimation—namely, predicting the most likely target variable $\mathbf{y} \in \mathcal{Y}$ from the Gaussian conditional distribution $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}$ —reduces to a mean-seeking supervised regression problem of approximating the conditional mean function $\mathbf{f} : \mathcal{O}^H \rightarrow \mathcal{Y}$ with a parametric model $\mathbf{f}_\theta : \mathcal{O}^H \rightarrow \mathcal{Y}$; that is,

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{mse}}(\mathbf{f}_\theta) := \arg \min_{\theta} \mathbb{E}_{\vec{\sigma} \sim \mathbb{P}_{\vec{\sigma}}} [\|\mathbf{f}(\vec{\sigma}) - \mathbf{f}_\theta(\vec{\sigma})\|^2], \quad (7.5)$$

Symmetry priors in supervised regression As discussed in [section 7.1](#), for robotic systems with an environmental or morphological symmetry group \mathbb{G} , the conditional distribution of the target variable \mathbf{y} given the observation history $\vec{\sigma}$ is \mathbb{G} -invariant (eq. 7.3). Under the working assumption that $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}$ is approximately Gaussian, the \mathbb{G} -invariance of the conditional distribution implies the

\mathbb{G} -equivariance of both the conditional mean function $\mathbf{f} : \mathcal{O}^H \rightarrow \mathcal{Y}$ and the conditional covariance² $\Sigma : \mathcal{O}^H \rightarrow \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$ (see [definition 2.14](#) and [figs. 6.1](#) and [7.1](#)), namely:

$$\begin{aligned} \mathbb{P}_{\mathbf{y}|\vec{\sigma}}(g \triangleright_{\mathcal{Y}} \mathbb{S} | g \triangleright_{\mathcal{O}^H} \vec{\sigma}) &= \mathbb{P}_{\mathbf{y}|\vec{\sigma}}(\mathbb{S} | \vec{\sigma}) & \iff & \quad g \triangleright_{\mathcal{Y}} \mathbf{f}(\vec{\sigma}) = \mathbf{f}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}) \quad \text{and} \\ \mathcal{N}(\mathbf{f}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}), \Sigma(g \triangleright_{\mathcal{O}^H} \vec{\sigma})) &= \mathcal{N}(\mathbf{f}(\vec{\sigma}), \Sigma(\vec{\sigma})) & & \quad g \triangleright_{\mathcal{Y}} \Sigma(\vec{\sigma})_{\mathcal{Y} \triangleleft g} = \Sigma(g \triangleright_{\mathcal{O}^H} \vec{\sigma}) \quad (7.6) \\ & & & \quad \forall g \in \mathbb{G}, \mathbb{S} \subseteq \mathcal{Y}, \vec{\sigma} \in \mathcal{O}^H. \end{aligned}$$

As detailed in [chapter 5](#), the \mathbb{G} -equivariance of \mathbf{f} in the supervised regression problem (eq. [7.5](#)), together with the \mathbb{G} -invariance of the marginal distribution of observation histories $\mathbb{P}_{\vec{\sigma}}$ ([assumption 6.2](#)), justifies reformulating the learning problem as a \mathbb{G} -equivariant regression problem of the form:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} \mathcal{L}_{\text{mse}}(\mathbf{f}_{\theta}) \\ \text{s.t.} \quad & g \triangleright_{\mathcal{Y}} \mathbf{f}_{\theta}(\vec{\sigma}) = \mathbf{f}_{\theta}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}) \quad \forall g \in \mathbb{G}, \vec{\sigma} \in \mathcal{O}^H. \end{aligned} \quad (7.7)$$

As previously explained, these symmetry priors can be leveraged either via data augmentation or by constrained optimization methods (such as using \mathbb{G} -equivariant NN architectures ([Bronstein et al., 2021](#); [Cohen et al., 2019](#); [Wang et al., 2022b](#))), or via regularization methods. In the latter case, the stochastic \mathbb{G} -invariance violation penalty term in eq. [5.8](#) reduces to a \mathbb{G} -equivariance penalty term of the form:

$$\mathcal{P}_{\text{eq}}(\mathbf{f}_{\theta}) := \mathbb{E}_{\vec{\sigma} \sim \mathbb{P}_{\vec{\sigma}}} \left[\frac{1}{|\mathbb{G}|} \sum_{g \in \mathbb{G}} \|g \triangleright_{\mathcal{Y}} \mathbf{f}_{\theta}(\vec{\sigma}) - \mathbf{f}_{\theta}(g \triangleright_{\mathcal{O}^H} \vec{\sigma})\|^2 \right]. \quad (7.8)$$

Next, we present an empirical analysis of the impact of leveraging these symmetry priors in supervised regression problems in robotics.

Experimental validation

To validate the impact of symmetry in supervised regression problems in robotics, we present a synthetic experiment on predicting a robot's CoM linear ($\mathbf{l} \in \mathbb{R}^3$) and angular ($\mathbf{k} \in \mathbb{R}^3$) momenta from noisy state observations. Specifically, the target function is given by:

$$\mathbf{y}_t := \begin{bmatrix} \mathbf{l}_t \\ \mathbf{k}_t \end{bmatrix} = \mathbf{f}(\mathbf{o}_t) := \mathbf{A}_C(\hat{\mathbf{q}}_t) \left(\hat{\dot{\mathbf{q}}}_t \right) \quad \forall t \in \mathbb{R}, \quad \text{with} \quad \begin{aligned} H &= 1, \\ \vec{\sigma}_t &:= \mathbf{o}_t = \begin{bmatrix} \hat{\mathbf{q}}_t := \mathbf{q}_t + \epsilon_{\mathcal{Q}}(t) \\ \hat{\dot{\mathbf{q}}}_t := \dot{\mathbf{q}}_t + \epsilon_{\mathcal{T}_q \mathcal{Q}}(t) \end{bmatrix} \end{aligned} \quad (7.9)$$

In this setting, the target random variable $\mathbf{y} \in \mathbb{R}^6$ consists of the stacked linear and angular momentum components. The function $\mathbf{A}_C : \mathcal{Q} \rightarrow \mathbb{R}^{6 \times n_q}$ is the state-dependent centroidal momentum matrix, introduced by [Orin et al. \(2013\)](#). For each configuration $\mathbf{q} \in \mathcal{Q}$, $\mathbf{A}_C(\mathbf{q})$ maps generalized velocities to the center of mass linear and angular momenta. Since \mathbf{A}_C depends only on the current robot state $(\mathbf{q}_t, \dot{\mathbf{q}}_t) \in \mathcal{T}\mathcal{Q}$, we set $H = 1$ in this experiment and reduce the observation history to the present observation $\mathbf{o}_t := (\hat{\mathbf{q}}_t, \hat{\dot{\mathbf{q}}}_t) \in \mathbb{R}^{2 \cdot n_q}$, which consists of sensor measurements of the generalized coordinates and velocities at time t , corrupted by zero-mean Gaussian noise $\epsilon_{\mathcal{Q}}$ and $\epsilon_{\mathcal{T}_q \mathcal{Q}}$, respectively.

As detailed in ([Orin et al., 2013](#); [Ordoñez-Apaez et al., 2023, 2025](#)), for robotic systems with an environmental or morphological symmetry group \mathbb{G} , the centroidal momentum matrix function is a

²Note that the equivariance of the covariance matrix imposes strict constraints on the degrees of freedom of the Gaussian uncertainty (see [prop. 2.2](#)). These constraints are of relevant interest for symmetry aware sensor-fusion frameworks, common in robot pose estimation from IMU data and other sensor modalities.




Figure 7.2 Sample efficiency plots (final test set loss achieved vs. training dataset size) for regression of the center of mass linear and angular momenta (see eq. 7.9) for the Atlas humanoid and Solo quadruped robots. We compare MLP, MLPaug, and eMLP models across varying model capacities and symmetry groups. Left and Middle: Sample efficiency for Solo and Atlas tasks, comparing model variants with different numbers of neurons in the hidden layers of the NN (hidden channels $hc=128/256/512/1024$). Right: Sample efficiency for Solo with models at $hc=512$, leveraging either $\mathbb{G} = \mathbb{K}_4$ (sagittal and transversal symmetries) or $\mathbb{G} = \mathbb{C}_2$ (sagittal symmetry only). Results show mean and standard deviation across 10 random training seeds.

\mathbb{G} -equivariant matrix-valued function, satisfying

$$g \triangleright_{\epsilon} \mathbf{A}_C(\mathbf{q}) \tau_{qQ} \triangleleft g^{-1} = \mathbf{A}_C(g \triangleright_Q \mathbf{q}) \quad \forall g \in \mathbb{G}, \mathbf{q} \in Q,$$

where the group actions $\triangleright_{\epsilon}$ and $\triangleright_{\tau_Q Q}$, defined in eq. 4.3d, ensure that the target function \mathbf{f} is \mathbb{G} -equivariant, as illustrated in fig. 6.2 for the case of a quadruped robot with $\mathbb{G} = \mathbb{K}_4$. In this figure, the orange and green vectors next to each robot visually represent the robot’s linear and angular center of mass momenta for that particular robot state, and $\triangleright_{\epsilon}$ represents the group action relating these symmetry-related center of mass momenta.

The experiment consists of a center of mass regression task for two robots: Atlas, a 32-DoF humanoid robot with $\mathbb{G} = \mathbb{C}_2$ (see figs. 2.1 and 4.1), and Solo, a 12-DoF quadruped robot with $\mathbb{G} = \mathbb{K}_4$ (see figs. 2.1 and 6.2). The baseline model is a standard Multi-Layer Perceptron (MLP) NN architecture, which is compared against two alternative approaches that exploit symmetry: (i) a MLP trained with data augmentation (MLPaug), and (ii) an Equivariant MLP (eMLP) architecture.

In fig. 7.2-left and -middle, we study the impact of symmetry exploitation on the models’ sample efficiency (final test loss vs. the size of the training dataset) and generalization (final test loss). To achieve this, we compare the performance of the model variants when leveraging the robots’ entire symmetry groups— \mathbb{K}_4 and \mathbb{C}_2 , respectively—and varying the width of the hidden layers of the underlying 3-layer deep MLP NNs. Across both robots and all model capacities, eMLP and MLPaug outperform MLP in terms of sample efficiency (better generalization with fewer data) and robustness to overfitting when training data is limited and the model is over-parameterized. Among the eMLP and MLPaug variants, models with the lowest number of trainable parameters exhibit similar sample complexity and performance; however, as model capacity increases, the eMLPs models reach consistently better test losses, highlighting the relevance of leveraging constrained optimization methods beyond mere data augmentation.

Additionally, fig. 7.2-right compares sample efficiency and optimality for the regression task of the Solo robot when leveraging the entire symmetry group \mathbb{K}_4 (sagittal and transversal symmetries) and

when leveraging only a subgroup of the true symmetry group $\mathbb{C}_2 = \{e, g[s]\} \subset \mathbb{K}_4$ (sagittal symmetry only). The results indicate a strong improvement in sample efficiency and generalization capacity with the increased number of symmetries leveraged, with the optimality gaps between **MLPaug** and **eMLP** being even more pronounced when leveraging the larger symmetry group \mathbb{K}_4 . This motivates the importance of leveraging the largest possible symmetry group in state estimation problems, and the relevance of constrained optimization methods beyond data augmentation when leveraging larger symmetry groups.

7.2.2 Supervised classification

Here, the unobservable target random variable of interest $\mathbf{y} \subset \mathbf{s}$ is categorical, taking values from a discrete set of classes $\mathbf{y} \in \mathcal{Y}$. This means its conditional distribution, given state observations, can reasonably be assumed to be a categorical distribution:

$$\mathbb{P}_{\mathbf{y}|\vec{\sigma}} := \mathbb{P}(\mathbf{y} = \mathbf{y}_i | \vec{\sigma} = \vec{\sigma}) \approx \text{Categorical}(\mathbf{f}(\vec{\sigma})),$$

where the probability vector function $\mathbf{f} : \mathcal{O}^H \rightarrow \Delta^{|\mathcal{Y}|-1}$ maps the truncated observation history $\vec{\sigma} \in \mathcal{O}^H$ to the probability simplex of the possible categorical states.

A classical example of such a variables in legged robotics is the binary contact state of a quadruped robot with point feet. The categorical target variable $\mathbf{y} \in \{0, 1\}^4$, representing the contact state of the four feet (e.g., in a static-friction-regime, or in a slipping or non-contact state), takes one of $|\mathcal{Y}| = 16$ possible binary combinations (Lin et al., 2021). These categorical estimated measurements are fundamental in state estimation pipelines dependent on leg odometry (e.g., Camurri et al., 2017) and legged locomotion control (e.g., Corbères et al., 2024; Lee et al., 2020; Mastalli et al., 2020b).

In this setting, the problem of state estimation reduces to a standard mean-seeking supervised classification problem of approximating the conditional class probability function $\mathbf{f} : \mathcal{O}^H \rightarrow \Delta^{|\mathcal{Y}|-1}$ with a parametric model (such as a neural network) $\mathbf{f}_\theta : \mathcal{O}^H \rightarrow \Delta^{|\mathcal{Y}|-1}$, typically trained by minimizing the expected cross-entropy loss:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{ce}}(\mathbf{f}_\theta) := \arg \min_{\theta} \mathbb{E}_{\vec{\sigma} \sim \mathbb{P}_{\vec{\sigma}}} \left[- \sum_{i=1}^{|\mathcal{Y}|} y_i \log \left(\mathbf{f}_\theta^{(i)}(\vec{\sigma}) \right) \right], \quad (7.10)$$

where \mathbf{y} is the one-hot encoded ground truth categorical state.

Symmetry priors in supervised classification Following section 7.1, under the assumption that $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}$ is a categorical distribution, the \mathbb{G} -invariance of the conditional distribution implies that the underlying probability function \mathbf{f} must be \mathbb{G} -equivariant:

$$g \triangleright_{\mathcal{Y}} \mathbf{f}(\vec{\sigma}) = \mathbf{f}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}) \quad \forall g \in \mathbb{G}, \vec{\sigma} \in \mathcal{O}^H.$$

In this case, the group action on the categorical variable's probability simplex, $\triangleright_{\mathcal{Y}}$, is simply defined as a subset of permutations acting on the categorical classes. These permutations represent how the categorical values transform under the symmetry group. For instance, for the quadruped robot in fig. 7.1 the binary contact states of the legs of the robot at each symmetry related state are permutations of the original contact state.




Figure 7.3 Static-friction-regime contact state classification results for the Mini-Cheetah quadruped robot. We compare the baseline Convolutional NN (CNN) architecture from Lin et al. (2021) with its symmetry-aware variants: CNN trained with data-augmentation (CNNaug) and eCNN. Left and Middle: Sample efficiency plots showing the final test set cross-entropy loss and the average continuous leg F1-score versus the size of the training dataset. Right: Contact state classification metrics evaluated on the test set using models trained with the entire dataset, reporting overall binary contact-state accuracy (Acc) and individual leg F1-scores. Due to the sagittal and transversal symmetries of the robot ($\mathbb{G} = \mathbb{K}_4 \times \mathbb{C}_2$), the F1-scores of symmetric leg pairs (e.g., LF/RF and LH/RH) are expected to be well-balanced for the symmetry-aware models. Results show mean and standard deviation across 8 random training seeds.

Experimental validation

To validate the impact of symmetry in supervised classification problems in robotics, we present an experiment on classifying the contact state variables $\mathbf{y} \in \{0, 1\}^4$ of the Mini-Cheetah quadruped robot from a history of proprioceptive sensor measurements $\vec{\sigma}$. For this system, the symmetry group corresponds to $\mathbb{G} = \mathbb{K}_4 \times \mathbb{C}_2$.

The dataset $\mathbb{D} = \{(\vec{\sigma}_i, \mathbf{y}_i)\}_{i=1}^N$, first introduced by Lin et al. (2021), comprises samples containing the one-hot encoded categorical target variable $\mathbf{y} \in \{0, 1\}^{16}$. This variable represents the ground truth system’s contact state among the 16 different combinations of each of the 4 legs’ possible binary contact states. Each sample also includes a truncated observation history $\vec{\sigma} \in \mathcal{O}^H$ of the past $H = 150$ frames of proprioceptive sensor measurements, defined by:

$$\mathbf{o}_t := (\mathbf{q}_{js} + \epsilon_{\mathcal{M}}(t), \dot{\mathbf{q}}_{js} + \epsilon_{\mathcal{T}_q\mathcal{M}}(t), \ddot{\mathbf{r}}_B + \epsilon_{\dot{\mathbf{r}}}(t), \boldsymbol{\omega}_B + \epsilon_{\boldsymbol{\omega}}(t), \mathbf{r}_{\text{feet}} + \epsilon_{\mathbf{r}}(t), \dot{\mathbf{r}}_{\text{feet}} + \epsilon_{\dot{\mathbf{r}}}(t),)$$

Here, $\mathbf{q}_{js} \in \mathcal{M} \subseteq \mathbb{R}^{12}$ and $\dot{\mathbf{q}}_{js} \in \mathcal{T}_q\mathcal{M} \subseteq \mathbb{R}^{12}$ denote the robot’s joint space position and velocity generalized coordinates, $\ddot{\mathbf{r}}_B$ and $\boldsymbol{\omega}_B \in \mathbb{R}^3$ represent the robot’s base linear acceleration and angular velocity measured by an onboard IMU, and $\mathbf{r}_{\text{feet}} = [\mathbf{r}_{\text{RF}}, \mathbf{r}_{\text{LF}}, \mathbf{r}_{\text{RH}}, \mathbf{r}_{\text{LH}}] \in \mathcal{O}_{\text{feet}} \subseteq \mathbb{R}^{12}$ and $\dot{\mathbf{r}}_{\text{feet}} = [\dot{\mathbf{r}}_{\text{RF}}, \dot{\mathbf{r}}_{\text{LF}}, \dot{\mathbf{r}}_{\text{RH}}, \dot{\mathbf{r}}_{\text{LH}}] \in \mathcal{O}_{\text{feet}} \subseteq \mathbb{R}^{12}$ are the positions and velocities of each of the four legs’ feet, respectively, referenced to the robot’s base frame. All these observation measurements are perturbed by zero-mean Gaussian noise modelling the respective sensor model uncertainty.

The baseline model is a standard 1D-CNN architecture utilized in (Lin et al., 2021), which is compared against two alternative approaches that exploit symmetry: (i) a data-augmented CNNaug, and (ii) an equivariant CNN (eCNN) architecture. In fig. 7.3-left and -middle, we study the impact of symmetry exploitation on the models’ sample efficiency (final test cross-entropy loss vs. the size of the training dataset) and generalization (average continuous leg F1-score). Across all data regimes, eCNN and CNNaug outperform CNN in terms of sample efficiency (achieving lower test loss and higher F1-score with fewer data) and robustness to dataset biases when training data is limited. Among the eCNN

and CNNaug variants, the fully equivariant model eCNN reaches consistently better test metrics and accuracy. Additionally, fig. 7.3-right compares the classification metrics of the models on the test set. Here, a closer look at the individual legs' F1-scores illustrates how exploiting symmetries mitigates suboptimal asymmetries commonly learned from finite datasets. While the unconstrained CNN favors one leg over its symmetric counterpart (e.g., learning a better predictor for the LF leg than for the RF leg), the symmetry-aware models, particularly eCNN, maintain a balanced predictive performance across symmetric morphological components, directly improving the overall classification reliability.

7.3 Probabilistic inference

For most target random variables $\mathbf{y} \subset \mathbf{s}$, we cannot safely assume that the conditional distribution $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}$, given the past observation history, belongs to a simple parametric family (e.g., Gaussian, Poisson, or Mixture of Gaussians) for all $\vec{\sigma} \in \mathcal{O}^H$, since in realistic scenarios, epistemic uncertainty in contact modes and forces induces complex multimodality in the observation dynamics and in the true conditional state distribution $\mathbb{P}_{\mathbf{s}|\vec{\sigma}}^{\Delta t}$ (see figs. 6.2 and 7.1).

In such cases, state estimation cannot be reduced to a standard supervised regression problem, given that, for most conditional distributions, the conditional mean (or *conditional expectation*),

$$\mathbf{y}_{\text{mean}} := \mathbb{E}[\mathbf{y}|\vec{\sigma}=\vec{\sigma}_t] = \int_{\mathbf{y}} \mathbf{y} \mathbb{P}_{\mathbf{y}|\vec{\sigma}}(d\mathbf{y}|\vec{\sigma}_t),$$

is not sufficiently informative for state estimation—that is, it does not allow us to sample likely values of \mathbf{y} according to $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}$ (see fig. 7.4). In this scenarios, state estimation should be addressed in its full stochastic form, as the approximation of a conditional probability distribution (see eqs. 7.2 and 7.4). However, estimating such conditional distributions remains challenging in high-dimensional settings without strong inductive biases (Izbicki and B. Lee, 2017; Nagler and Czado, 2016; Scott, 1991).

We differentiate between two classes of conditional probability modelling: methods that aim to approximate the sampling process of the conditional distribution via generative modelling techniques, and approaches that provide point- or set-wise estimates of the conditional likelihood of events of interest (e.g., the probability that a contact force is within a specific range, the probability a disturbance in my left arm is present). In summary:

- ✱ **Generative modelling:** Most state of the art generative modelling methods—such as diffusion and flow-matching Ho et al. (2020); Lipman et al. (2022)—model first-order information of the probability, such as the gradient of the likelihood with respect to the target variable, $\nabla_{\mathbf{y}} \mathbb{P}_{\mathbf{y}|\vec{\sigma}}(\mathbf{y}|\vec{\sigma}_t)$. This is used to approximate sampling from the conditional distribution via a generative process mapping samples from a base distribution over \mathbf{y} to $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}(\mathbf{y}|\vec{\sigma}_t)$. By construction, these approaches do not estimate the likelihood of events and therefore cannot truly tell from two generated samples which one is more likely, rendering them not directly suited to solve the *maximal a posteriori* state estimation problem. Although recent efforts have




Figure 7.4 Example uninformative conditional expectations.

begun to adapt generative modelling to maximum a posteriori estimation (Tiofack et al., 2025; Willingham et al., 2025), these approaches remain out of the scope of this work.

Note: In sections 2.4 and 8.3.1, we discuss the implications of our symmetry priors in the context of flow-matching based imitation learning. The results transfer to state estimation by changing the target from control actions \mathbf{a}_t to state variables \mathbf{y}_t .

✱ **Point/set-wise estimation:** We interpret these methods as approximating the zero-th order information of the conditional probability, $\mathbb{P}_{\mathbf{y}|\vec{\sigma}}(\cdot | \vec{\sigma}_t)$, at specific points or sets of interest (Izbicki and B. Lee, 2017; Kostic et al., 2024; Nagler and Czado, 2016). This means that they can estimate a scalar value that approximates the likelihood of specific events of interest, or that it’s proportional to the likelihood. These are the focus of this section as they enable the robust decision-making, planning and control.

7.3.1 Conditional probability modelling and uncertainty quantification

The problem of conditional probability modelling and parametric uncertainty quantification remains an open challenge in high-dimensional settings in which one cannot assume that the target conditional distribution belongs to a simple parametric family (Izbicki and B. Lee, 2017; Nagler and Czado, 2016; Scott, 1991). Recent advances in Deep Learning (DL) and deep representation learning have shown promise tackling conditional probability modelling—and therefore reliable uncertainty quantification—in these high-dimensional and unstructured settings by training non-linear representations of the target and conditioning variables within the framework of contrastive representation learning (Ryu et al., 2024; Tosh et al., 2021; Turri et al., 2025).

Specifically, Kostic et al. (2024) proposed Neural Conditional Probability (NCP) (discussed in detail in chapter 9), a contrastive representation learning algorithm that rely on the pre-trained representation to provide, via linear regression, estimates of conditional probabilities of events of the form $\mathbf{y} \in \mathbb{Y}$ for sets $\mathbb{Y} \subset \mathcal{Y}$, via

$$\mathbb{P}(\mathbf{y} \in \mathbb{Y} | \vec{\sigma} = \vec{\sigma}_t) \approx \hat{\mathbb{P}}_{\theta}(\mathbf{y} \in \mathbb{Y} | \vec{\sigma}_t) := \hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{Y}}(\mathbf{y})] + \phi_{\theta}(\vec{\sigma}_t)^{\top} \mathbf{E}_{\theta} \hat{\mathbb{E}}_{\mathbf{y}}[\psi_{\theta}(\mathbf{y}) \otimes \mathbb{1}_{\mathbb{Y}}(\mathbf{y})], \quad (7.11)$$

where $\phi_{\theta} : \mathcal{O}^H \rightarrow \mathbb{R}^n$ and $\psi_{\theta} : \mathcal{Y} \rightarrow \mathbb{R}^n$ are learned, nonlinear, high-dimensional representations of the conditioning variable $\vec{\sigma}$ and the target variable \mathbf{y} , respectively, and $\mathbf{E}_{\theta} \in \mathbb{R}^{n \times n}$ is a linear map between the two representation spaces (see fig. 9.3). Moreover, $\mathbb{1}_{\mathbb{Y}}(\cdot)$ denotes the indicator function of the event $\mathbf{y} \in \mathbb{Y}$, and $\hat{\mathbb{E}}_{\mathbf{y}}[\psi_{\theta}(\mathbf{y}) \otimes \mathbb{1}_{\mathbb{Y}}(\mathbf{y})] \in \mathbb{R}^n$ is the vector of basis-expansion coefficients of $\mathbb{1}_{\mathbb{Y}}(\cdot)$ in the representation space induced by ψ_{θ} . Consequently, for sets \mathbb{Y} for which we can reliably estimate the marginal probability $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{Y}}(\mathbf{y})] \approx \mathbb{P}(\mathbf{y} \in \mathbb{Y})$, eq. 7.11 provides a valuable estimate of $\mathbb{P}(\mathbf{y} \in \mathbb{Y} | \vec{\sigma} = \vec{\sigma}_t)$ for a wide range of conditioning values $\vec{\sigma}_t \in \mathcal{O}^H$.

In the next section, we study how eq. 7.11 can be leveraged for parametric uncertainty quantification of unobservable state variables in robotics, using any NN architecture to parameterize ϕ_{θ} and ψ_{θ} . Crucially, this method works in high-dimensional settings and without assumptions on the form of the conditional distribution; hence, it can be used in foundational tasks of robust decision-making, planning, and control under uncertainty.

Details on the theoretical underpinnings of this deep NN-based estimate (eq. 7.11) require an extensive operator-theoretic treatment of conditional distributions and are therefore provided in chapter 9. The rest of this section focuses on the practical use of eq. 7.11 for uncertainty quantification in robotics.

Conditional quantile regression One of the main approaches to conditional uncertainty quantification is conditional quantile regression. In this setting, given a *scalar* target random variable $y \subset \mathbf{y} \subset \mathbf{s}$, the goal is to predict a confidence interval with coverage level α for y , conditioned on the observation history $\vec{\sigma}$. For example, we may wish to predict the $\alpha = 95\%$ confidence interval for the z -component of a contact force. To this end, we regress the lower and upper conditional quantiles, $c_{\alpha/2} : \mathcal{O}^H \rightarrow \mathbb{R}$ and $c_{1-\alpha/2} : \mathcal{O}^H \rightarrow \mathbb{R}$, respectively. These quantiles are derived from the **conditional Cumulative Distribution Function (cCDF)** of the target variable y , defined as

$$\text{cCDF}(y \mid \vec{\sigma}) := \mathbb{P}(y \leq y \mid \vec{\sigma} = \vec{\sigma}). \quad (7.12)$$

To approximate this function, we discretize the marginal range of y into N_b bins and define the sets $\mathbb{A}_n := \{y \in \mathbb{R} : y \leq b_n\}$, where b_n denotes the upper bound of the n -th bin.³ This allows us to leverage $\widehat{\mathbb{P}}_\theta$ (eq. 7.11) to estimate a discretized version of the cCDF. For all conditioning events $\vec{\sigma} \in \mathcal{O}^H$ and all bins $n \in [N_b]$, we define

$$\widehat{\text{cCDF}}_\theta(y \mid \vec{\sigma}) := \widehat{\mathbb{P}}_\theta(y \in \mathbb{A}_n \mid \vec{\sigma}) = \widehat{\mathbb{E}}_y[\mathbb{1}_{\mathbb{A}_n}(y)] + \phi_\theta(\vec{\sigma})^\top \mathbf{E}_\theta \widehat{\mathbb{E}}_y[\psi_\theta(y) \otimes \mathbb{1}_{\mathbb{A}_n}(y)]. \quad (7.13)$$

A key advantage of this formulation is that the marginal probabilities $\widehat{\mathbb{E}}_y[\mathbb{1}_{\mathbb{A}_n}(y)]$ can be estimated easily and reliably from training data, making $\widehat{\mathbb{P}}_\theta$ a practical estimator of the cCDF for arbitrary scalar random variables y . Moreover, for a discretization with N_b bins, the full estimate $\widehat{\text{cCDF}}_\theta$ can be obtained using a single forward pass through the NNs ϕ_θ and ψ_θ , followed by a linear layer that maps $\phi_\theta(\vec{\sigma}) \in \mathbb{R}^n$ to \mathbb{R}^{N_b} . This makes the approach computationally efficient and well suited to real-time applications. Importantly, when extending the prediction of the cCDF to a vector-valued target $\mathbf{y} = [y_1, \dots] \subset \mathbf{s}$, the additional computational cost arises only in the final linear head, which is now defined by a matrix $\mathbb{R}^{n \times (|\mathcal{Y}| \times N_b)}$ composed of the basis expansion coefficients of the indicator functions $\mathbb{1}_{\mathbb{A}_{i,n}}(\cdot)$ for all dimensions i and bins n .

Symmetry priors in conditional quantile regression In the context of uncertainty quantification for robotic systems with an environmental or morphological symmetry group \mathbb{G} , the conditional distribution of the target vector-valued random variable is \mathbb{G} -invariant (see eq. 7.3):

$$\mathbb{P}_{\mathbf{y}|\vec{\sigma}}(\mathbf{y} \in \mathbb{Y} \mid \vec{\sigma} = \vec{\sigma}) = \mathbb{P}_{\mathbf{y}|\vec{\sigma}}(g \triangleright_{\mathbf{y}} \mathbf{y} \in \mathbb{Y} \mid \vec{\sigma} = g \triangleright_{\mathcal{O}^H} \vec{\sigma}) \quad \forall g \in \mathbb{G}, \vec{\sigma} \in \mathcal{O}^H, \mathbb{Y} \subseteq \mathcal{Y}.$$

whenever there is a well-defined group action $\triangleright_{\mathbf{y}}$ on the target variable space \mathcal{Y} . For example, when estimating the contact forces of a quadruped robot, the forces at each foot $\mathbf{y} \in \mathbb{R}^{12}$ for symmetry-related robot states transform according to $\triangleright_{\mathbf{y}}$, which permutes the legs' contact forces and applies the appropriate reflections or rotations to the vectors (see fig. 7.1).⁴

³To obtain a discretization adapted to the marginal distribution, one may use methods such as quantile binning (Pedregosa et al., 2011), which allocates bins so as to more uniformly cover the support of the marginal distribution \mathbb{P}_y .

⁴The discretization of the target variable's range into bins for cCDF regression in this scenario can be most conveniently defined in a hyperspherical coordinate frame, instead in the canonical Cartesian coordinates, as

In such cases, as discussed in [prop. 7.1](#), uncertainty quantification can be rendered as a symmetry-constrained optimization problem in which the estimator $\hat{\mathbb{P}}_\theta$ must be \mathbb{G} -invariant. As discussed in detail in [section 7.1](#), to leverage the symmetry priors we can resort to data-augmentation, regularization, or reparametrization techniques via \mathbb{G} -equivariant architectures. In the following section we discuss how to In a nutshell, to enforce the \mathbb{G} -invariance constraint of $\hat{\mathbb{P}}_\theta$ implies parameterizing the NN representations ϕ_θ and ψ_θ with \mathbb{G} -equivariant architectures, and the linear map E_θ with a \mathbb{G} -equivariant linear layer. [chapter 9](#).

Experimental validation

We present two experiments on uncertainty quantification via conditional quantile regression: (i) a synthetic experiment illustrating the empirical underpinnings of [NCP](#) and [Equivariant Neural Conditional Probability \(eNCP\)](#), and (ii) a simulated robotics state-estimation experiment in which we predict confidence intervals for relevant observables in quadruped legged locomotion.

To assess the impact of symmetry exploitation, we compare the baselines [NCP](#) and [Conditional Quantile Regression \(CQR\)](#) ([Feldman et al., 2023](#)) with their respective equivariant adaptations, [eNCP](#) and [Equivariant CQR \(eCQR\)](#), proposed in [Ordoñez-Apaez et al. \(2026\)](#). Here, [NCP](#) and [eNCP](#) are trained to approximate the conditional distribution of the target variable, $\mathbb{P}_{y|\delta}$ (see [eq. 7.11](#) and [chapter 9](#)). By contrast, the baseline [CQR](#) ([Feldman et al., 2023](#)) and its equivariant adaptation [eCQR](#) directly regress quantiles for a fixed coverage level (e.g., $c_{\alpha/2}$ and $c_{1-\alpha/2}$ for a coverage level of α), which are not reusable across coverage levels and require retraining for different coverage values. All models are parameterized by NNs with equivalent architectural footprints, so performance differences can be attributed to contrastive versus regression losses and symmetry exploitation, rather than differences in model capacity. Refer to ([Ordoñez-Apaez et al., 2026](#)) for extensive details on the experimental setup.

Conditional quantile regression of skewed, symmetric, and bimodal distributions This synthetic experiment is intended to demonstrate how the [NCP](#) and [eNCP](#) frameworks can approximate conditional distributions that are highly skewed and multimodal, for which the mean (or conditional expectation) is not informative (see [fig. 7.4](#)). For such applications, we rely on accurate approximation of the conditional distribution to perform uncertainty quantification by regressing the [cCDF](#) of a one-dimensional target random variable $y \in \mathbb{R}$ conditioned on a one-dimensional variable $x \in \mathbb{R}$.

As illustrated in [fig. 7.5-top](#), the target conditional distribution $\mathbb{P}_{y|x}$ has three distinct regimes, all affected by heteroscedastic noise: (i) a skewed distribution with an exponential tail, (ii) a symmetric distribution, and (iii) a bimodal distribution. Crucially, the conditional distribution of the target variable is \mathbb{C}_2 -invariant, with a trivial group action on $\mathcal{Y} \in \mathbb{R}$ and a group action on $\mathcal{X} \in \mathbb{R}$ defined such that the reflection action of \mathbb{C}_2 maps x to $-x$, that is:

$$\mathbb{P}_{y|x}(y \in \mathbb{Y} \mid x = x) = \mathbb{P}_{y|x}(y \in \mathbb{Y} \mid x = -x) \quad \forall x \in \mathbb{R}, \mathbb{Y} \subseteq \mathbb{R}.$$

The task consists of predicting conditional quantiles at diverse coverage levels for the target variable y over the full range of x values. To achieve this, we discretize the range of y into $N_b = 200$ bins using `quantile_transform` from [Pedregosa et al. \(2011\)](#), which allocates higher bin resolution

described in [Ordoñez-Apaez et al. \(2026\)](#), such that the discretization covers each representative of the coset space \mathcal{Y}/\mathbb{G} or equivalently.

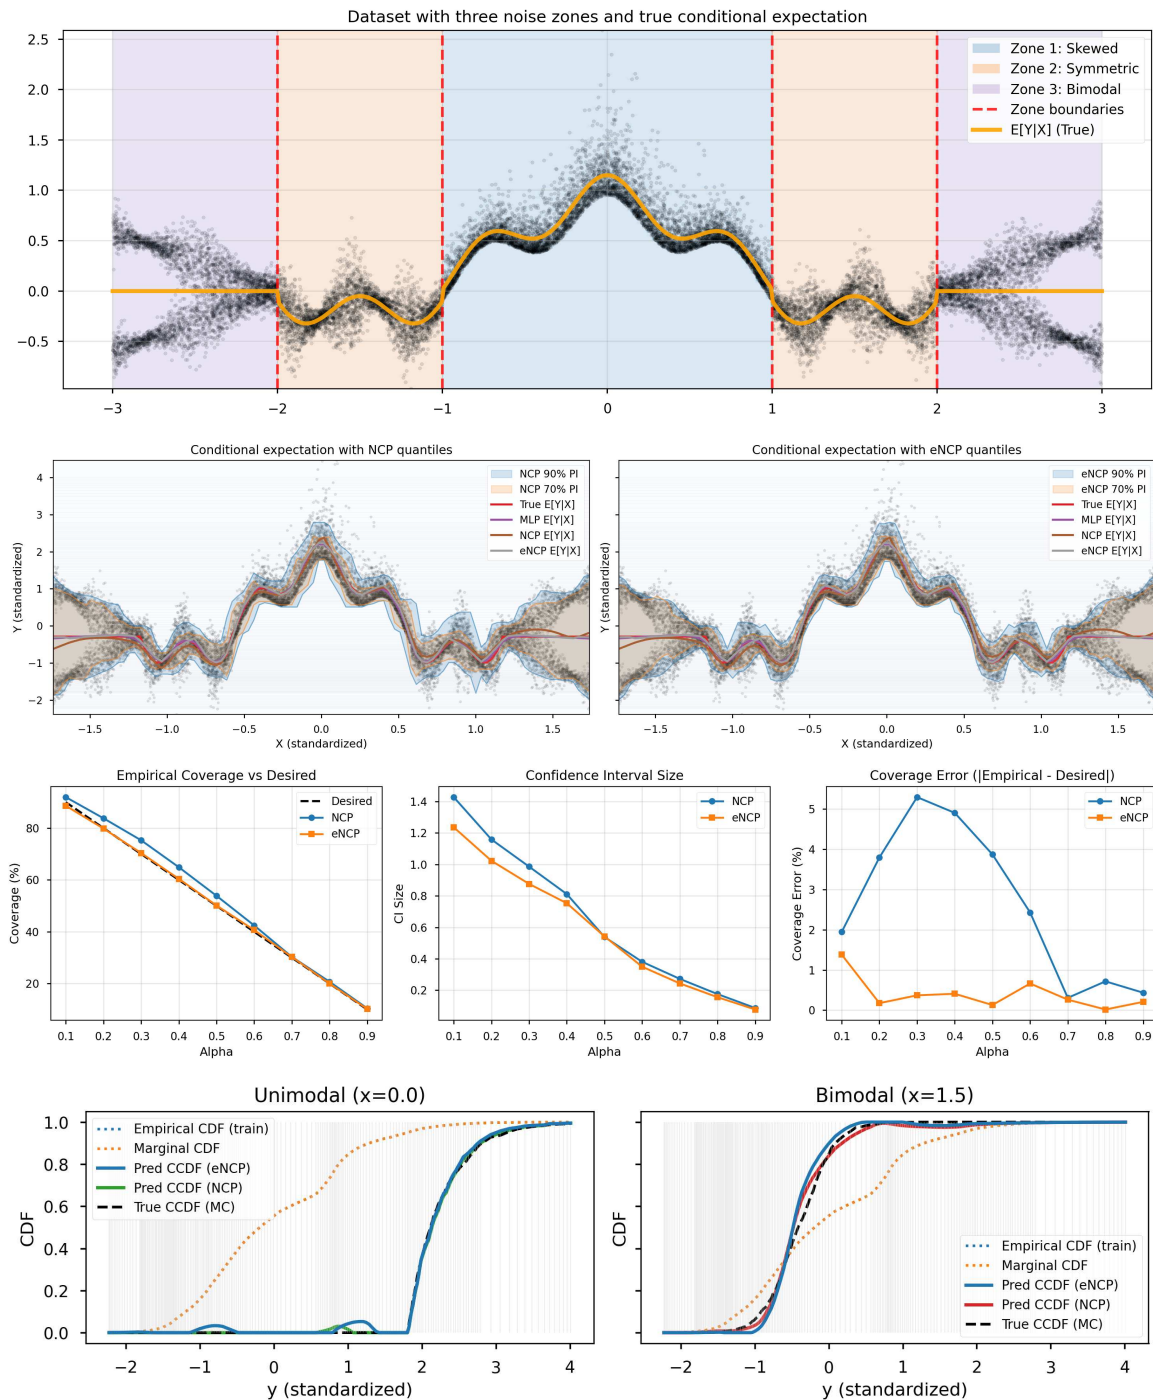


Figure 7.5 **1st row**: Target conditional probability distribution of y given x , featuring three distinct regimes: (i) a skewed distribution with an exponential tail, (ii) a symmetric distribution, and (iii) a bimodal distribution. **2nd row**: Illustration of the predicted confidence intervals for coverage levels $\alpha = \{70\%, 90\%\}$ for NCP (left) and eNCP (right). **3rd row**: Empirical coverage tracking and confidence interval size for coverage levels from 10% to 90%. **4th row**: Example cCDFs predicted by NCP and eNCP in the skewed and bimodal regimes, with vertical gray lines indicating the bins used to discretize the range of y for cCDF regression.

| | Validation | | | Test | | |
|------|-----------------------|---------------------|---|-----------------------|---------------------|---|
| | r-Coverage \uparrow | Coverage \uparrow | Set Size \downarrow | r-Coverage \uparrow | Coverage \uparrow | Set Size \downarrow |
| eNCP | 99.3 \pm 0.0% | 94.1 \pm 0.4% | 2.4 \pm 0.4 \times 10 ¹⁰ | 99.5 \pm 0.1% | 95.0 \pm 0.4% | 4.3 \pm 3.6 \times 10 ⁹ |
| NCP | 96.4 \pm 0.0% | 56.9 \pm 0.1% | 3.9 \pm 4.5 \times 10 ¹⁰ | 99.5 \pm 0.0% | 56.9 \pm 0.3% | 2.6 \pm 1.4 \times 10 ¹⁰ |
| eCQR | 70.7 \pm 0.6% | 7.3 \pm 1.7% | 3.7 \pm 2.6 \times 10 ⁸ | 84.2 \pm 0.7% | 6.7 \pm 1.2% | 1.7 \pm 1.7 \times 10 ⁷ |
| CQR | 67.6 \pm 1.8% | 7.6 \pm 0.4% | 2.5 \pm 2.4 \times 10 ⁹ | 80.5 \pm 3.7% | 8.5 \pm 0.9% | 1.4 \pm 0.1 \times 10 ⁸ |

Table 7.1 Validation and test metrics for 95% Confidence Intervals (CIs) on quadruped robot observables while traversing rough terrain. Metrics: (i) relaxed coverage (r-Coverage), defined as the probability that *any* scalar random variable v_i lies inside the predicted confidence interval; (ii) coverage, defined as the probability that the entire vector-valued target variable \mathbf{y} lies inside the predicted confidence interval; and (iii) set size, the volume of the confidence region. Best results are shown in blue. While eCQR and CQR produce smaller confidence intervals, they fail to achieve the expected 95% coverage on both validation and test sets. The eNCP model achieves the best overall coverage for reliable uncertainty quantification. Importantly, eNCP and NCP can provide CIs at any coverage level **without retraining**, whereas CQR and eCQR require retraining for each level.

to regions of the support of \mathbf{y} with higher probability mass, and regress the cCDF of \mathbf{y} using eq. 7.13. Then, the quantiles of any desired coverage level can be obtained by applying a linear search over the predicted cCDF (see Kostic et al. (2024)). Example cCDF predictions for the NCP and eNCP models are shown in fig. 7.5-bottom for the skewed and bimodal regimes, with the vertical gray lines indicating the bins used to discretize the range of \mathbf{y} .

The results for this low-dimensional example show that both NCP and eNCP can learn to approximate the conditional distribution $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ across the three regimes and produce parametric confidence intervals at different coverage levels that are well calibrated (see fig. 7.5-second-and-third-row). However, even in this low-dimensional setting, leveraging symmetry priors via eNCP yields smaller confidence intervals that match the empirical coverage across all coverage levels $\alpha \in \{10\%, \dots, 90\%\}$, while NCP produces conservative confidence intervals with empirical coverage above the desired levels.

Uncertainty quantification for quadruped locomotion We demonstrate the practical impact of symmetry exploitation in robotics uncertainty quantification by providing robust uncertainty estimates for unavailable yet crucial state observables in robot control and state estimation (Bledt et al., 2018; Maravagakis et al., 2023). Specifically, the task is to use a history of proprioceptive sensor observations to provide 90% CIs for the robot’s Ground Reaction Forces (GRF) $\boldsymbol{\tau}_{\text{grf}} \in \mathbb{R}^{12}$, the instantaneous work exerted by or absorbed by the robot $U(\mathbf{q}, \dot{\mathbf{q}}, \boldsymbol{\tau}) \in \mathbb{R}$ (eq. 2.6), and the robot’s kinetic energy $T(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}$ (eq. 2.3a) while the robot traverses rough terrain; i.e., $\mathbf{y} = [U(\mathbf{q}, \dot{\mathbf{q}}, \boldsymbol{\tau}), T(\mathbf{q}, \dot{\mathbf{q}}), \boldsymbol{\tau}_{\text{grf}}]$. Reliable probabilistic estimates of these quantities are crucial for optimal control (Bledt et al., 2018), contact detection (Maravagakis et al., 2023), state estimation (Nisticò et al., 2025), and system identification (Gautier, 1997).

The history of past observations $\vec{\mathbf{o}}_t \in \mathcal{O}^H$ includes $H = 5$ state observations defined by:

$$\mathbf{o}_t = (\mathbf{q}_t + \epsilon_{\mathcal{Q}}, \dot{\mathbf{q}}_t + \epsilon_{\mathcal{T}_q \mathcal{Q}}, \mathbf{a}_t + \epsilon_{\mathbf{a}}, \dot{\mathbf{r}}_t + \epsilon_{\dot{\mathbf{r}}}, \dot{\mathbf{r}}_{t,\text{err}} + \epsilon_{\dot{\mathbf{r}}}, \boldsymbol{\omega}_t + \epsilon_{\boldsymbol{\omega}}, \boldsymbol{\omega}_{t,\text{err}} + \epsilon_{\boldsymbol{\omega}}, \mathbf{g}_t + \epsilon_{\mathbf{g}}, \dot{\mathbf{p}}_{t,\text{feet}} + \epsilon_{\text{feet}}, \boldsymbol{\tau}_t^{\text{cmd}}) \quad (7.14)$$

where $\mathbf{q}_t \in \mathbb{R}^{n_q}$ and $\dot{\mathbf{q}}_t \in \mathbb{R}^{n_q}$ are the joint positions and velocities, respectively; $\mathbf{a}_t \in \mathbb{R}^3$ is the linear acceleration of the robot’s base frame measured by the IMU; $\mathbf{v}_t \in \mathbb{R}^3$ is the base linear velocity, while $\mathbf{v}_{t,\text{err}} \in \mathbb{R}^3$ is the base linear-velocity command error; $\boldsymbol{\omega}_t \in \mathbb{R}^3$ and $\boldsymbol{\omega}_{t,\text{err}} \in \mathbb{R}^3$ are the base angular velocity and its command error; $\mathbf{g}_t \in \mathbb{R}^3$ is the gravity vector expressed in the base frame;

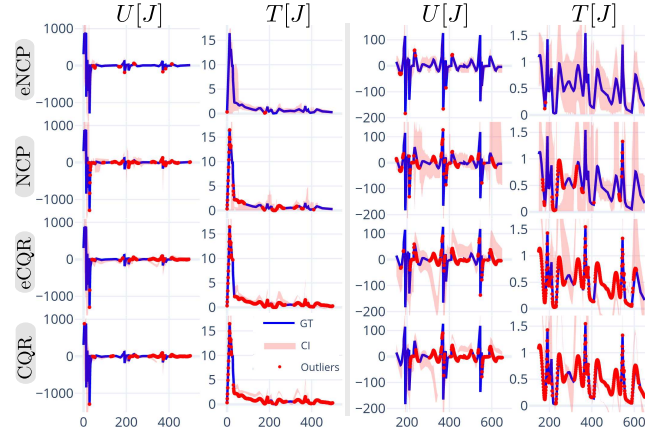


Figure 7.6 Prediction of 90% CIs for \mathbb{G} -invariant observables in quadruped locomotion: instantaneous work U and kinetic energy T . We compare eNCP, NCP, eCQR, and CQR. Confidence intervals are shown as light red areas; samples outside the predicted interval are shown in red, and samples inside are shown in blue.

$\dot{\mathbf{p}}_{t,\text{feet}} \in \mathbb{R}^{12}$ stacks the linear velocities of the four feet (three components each); and $\boldsymbol{\tau}_t^{\text{cmd}} \in \mathbb{R}^{n_q}$ contains the commanded joint torques.

In [fig. 7.6](#), we compare model performance for uncertainty quantification of the \mathbb{G} -invariant scalar observables U and T , while in [fig. 7.7](#), we compare model performance for uncertainty quantification of the \mathbb{G} -equivariant vector-valued observable $\boldsymbol{\tau}_{\text{grf}}$. Furthermore, [table 7.1](#) reports empirical coverage on the validation and test sets, as well as the average size of the predicted confidence intervals, which are crucial metrics for assessing practical reliability. The results identify eNCP as the only model capable of robust uncertainty quantification, as it is the only model with test-set empirical coverage close to the desired value, rendering the others unreliable for practical applications. The baseline models CQR and eCQR produce severely overconfident intervals that fit the training data but fail to generalize to the test set. Although eCQR outperforms CQR in empirical coverage, it still fails to achieve the desired test coverage level, pointing to a limitation of this regression-based approach to uncertainty quantification in high-dimensional settings.

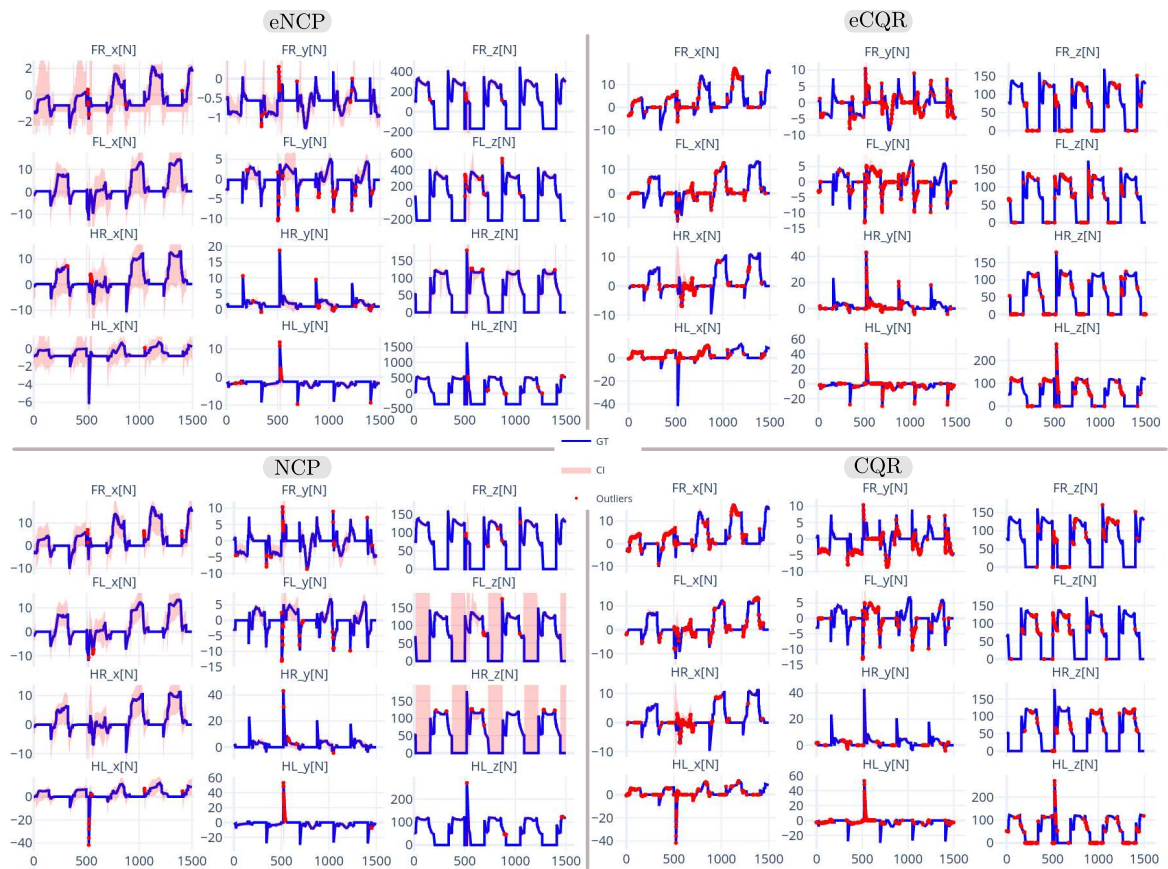


Figure 7.7 Prediction of 90% CIs for the \mathbb{G} -equivariant ground-reaction forces $\tau_{\text{grf}} \in \mathbb{R}^{12}$. Intervals are shown per leg (FR, FL, HR, HL) and Cartesian component (x , y , z). Confidence intervals are shown as light red areas; samples outside the predicted interval are shown in red, and samples inside are shown in blue.

Chapter 8

Morphological symmetries in data-driven optimal control

This chapter studies the implications of spatial symmetries for optimal control of robots in contact-rich environments. In [section 8.1](#), we analyze symmetry priors for stochastic control tasks modelled as MDP, extending [Ravindran and Barto \(2001\)](#); [Zinkevich and Balch \(2001\)](#) to continuous state and action spaces and compact symmetry groups. The core results establish that, for any \mathbb{G} -symmetric MDP, the optimal quality and value functions are \mathbb{G} -invariant ([theorem 8.1](#) and [corollary 8.1](#)), and identify symmetry constraints on optimal stochastic and deterministic policy parametrizations ([corollaries 8.2](#) and [8.3](#)). In [section 8.2](#), we transfer this analysis to [Partially Observable Markov Decision Process \(POMDPs\)](#) via the belief-MDP formulation and recover analogous invariance results for partially observable control settings, including real-world robotic control tasks in contact-rich environments.

Finally, in [section 8.3](#), we leverage these results in two practical settings:

- * **Imitation learning from expert demonstrations:** we characterize why behaviour cloning becomes a symmetry-constrained learning problem for robots with morphological symmetries ([eq. 8.15](#)), and present symmetry-exploitation strategies for flow-matching policy parametrizations ([Lipman et al., 2022](#)). We study Push-T and bimanual mobile manipulation under reflection symmetry ([Siebenborn et al., 2026](#)).
- * **Model-free reinforcement learning (RL):** we characterize how morphological symmetries induce constraints for the family of [Proximal Policy Optimization \(PPO\)](#) algorithms, and study symmetry-exploitation strategies in quadruped locomanipulation with \mathbb{C}_2 symmetry groups ([Su et al., 2024](#)) and dexterous multi-arm manipulation with \mathbb{C}_2 and \mathbb{C}_4 symmetry groups ([Li et al., 2025](#)).

8.1 Symmetry priors in optimal control

As discussed in detail in [section 2.2](#) and [chapter 6](#), the discrete-time stochastic world dynamics of a robot operating in a contact-rich environment are modelled by a state-transition probability distribution $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, which defines the conditional distribution of next states given current

states and actions, such that

$$\mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}(\mathbb{S} \mid \mathbf{a}_t, \mathbf{s}_t) := \mathbb{P}(\mathbf{s}_{t+\Delta t} \in \mathbb{S} \mid \mathbf{a}_t = \mathbf{a}_t, \mathbf{s}_t = \mathbf{s}_t) \quad \forall \mathbb{S} \subseteq \mathcal{S}, \mathbf{a}_t \in \mathcal{A}, \mathbf{s}_t \in \mathcal{S}, t \in \mathbb{R},$$

quantifies the probability of transitioning to a state in the set \mathbb{S} after a Δt time step, given that the robot takes action \mathbf{a}_t in world state \mathbf{s}_t at time t .

Here, the world state space \mathcal{S} comprises the robot state $(\mathbf{q}, \dot{\mathbf{q}}) \in \mathcal{TQ}$ and the robot–environment interaction state $\mathbf{z} \in \mathcal{Z}$ —which defines the control task and determines the robot’s current contact mode—while the action space $\mathcal{A} \subseteq \mathbb{R}^{n_{ctrl}}$ is the space of control forces applied to the robot’s actuated DoF. For instance, for any manipulation task (see fig. 8.1), \mathbf{s}_t captures the robot’s state, its contact mode, and the manipulation control task metrics, such as the configuration of a target object to be manipulated.

In this context, any robot control task can be modelled as an MDP, defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}, r, \mathbb{P}_{\mathcal{S}}^0, \mathbb{A}_{\pi}, \gamma)$, where $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a *bounded* control-task reward function, $\mathbb{P}_{\mathcal{S}}^0 : \Sigma_{\mathcal{S}} \rightarrow [0, 1]$ is the probability distribution of initial states, $\mathbb{A}_{\pi} : \mathcal{S} \rightarrow \Sigma_{\mathcal{A}}$ denotes the admissible action map that maps each state to its set of admissible actions, and $\gamma \in [0, 1)$ is a constant discount factor. The objective of optimal control is to find a parametric stochastic control policy $\pi_{\theta} : \Sigma_{\mathcal{A}} \times \mathcal{S} \mapsto [0, 1]$ that maximizes the expected discounted reward, which can be expressed as the following optimization problem:

$$\arg \max_{\theta} \mathbb{E}_{\mathbf{s}_t \sim \mathbb{P}_{\mathcal{S}}^0} \left[\mathbb{E}_{\pi_{\theta}} \left(\sum_{k=0}^{\infty} \gamma^k r(\mathbf{s}_{t+k\Delta t}, \mathbf{a}_{t+k\Delta t}) \right) \right] = \mathbb{E}_{\mathbf{s}_t \sim \mathbb{P}_{\mathcal{S}}^0} \left[V_{\pi_{\theta}}(\mathbf{s}_t) \right] \quad (8.1)$$

where $\mathbb{E}_{\pi_{\theta}}(\cdot)$ denotes the expectation over the stochastic trajectories generated by the world dynamics $\mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}$ when following policy π_{θ} .

To analyze such control problems and study the implications of spatial symmetries, we define two functions of interest: the *value function* $V_{\pi_{\theta}} : \mathcal{S} \rightarrow \mathbb{R}$ and the *action-value*, or *quality*, function $Q_{\pi_{\theta}} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. These quantify the expected cumulative reward of following policy π_{θ} starting from a given state \mathbf{s} and from a given state-action pair (\mathbf{s}, \mathbf{a}) , respectively. Formally, the quality function of policy π_{θ} is defined recursively as:

$$Q_{\pi_{\theta}}(\mathbf{s}_t, \mathbf{a}_t) = r(\mathbf{s}_t, \mathbf{a}_t) + \gamma \int_{\mathcal{S}} V_{\pi_{\theta}}(\mathbf{s}_{t+\Delta t}) \mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}(d\mathbf{s}_{t+\Delta t} \mid \mathbf{s}_t, \mathbf{a}_t) \quad \forall \mathbf{s}_t \in \mathcal{S}, \mathbf{a}_t \in \mathcal{A}. \quad (8.2)$$

Consequently, the value function of policy π_{θ} is defined as:

$$V_{\pi_{\theta}}(\mathbf{s}_t) = \int_{\mathcal{A}} Q_{\pi_{\theta}}(\mathbf{s}_t, \mathbf{a}_t) \pi_{\theta}(d\mathbf{a}_t \mid \mathbf{s}_t) \quad \forall \mathbf{s}_t \in \mathcal{S}. \quad (8.3)$$

Our aim is to show that, in control tasks for robots featuring a state/morphological symmetry group \mathbb{G} , the unique optimal value and quality functions (V^* and Q^*) are constrained to be \mathbb{G} -invariant, and that as a consequence, any optimal stochastic control policy π_* is also constrained to be \mathbb{G} -invariant conditional distribution (definition 2.17). To achieve this, we first define the symmetry priors of the control problem, reflected as structural properties of the MDP.

Symmetry priors in the control problem. As detailed in chapter 6, when a robotic system features a state/morphological symmetry group \mathbb{G} , the state and action spaces are endowed with group actions




Figure 8.1 Reflection-related box-lifting tasks for the Tiago mobile bimanual manipulator with sagittal reflection symmetry $g_s \in \mathbb{C}_2$. The initial states s_0 and $g_s \triangleright_S s_0$ correspond to left- and right-table variants of the same task. Task equivalence implies reward invariance, $r(s, \mathbf{a}) = r(g_s \triangleright_S s, g_s \triangleright_A \mathbf{a})$ for all $s \in \mathcal{S}$, $\mathbf{a} \in \mathcal{A}$, so an optimal trajectory for one variant maps to an optimal trajectory for the reflected variant by applying g_s to the state-action trajectory. Figure adapted from Siebenborn et al. (2026).

$\triangleright_S : \mathbb{G} \times \mathcal{S} \rightarrow \mathcal{S}$ and $\triangleright_A : \mathbb{G} \times \mathcal{A} \rightarrow \mathcal{A}$, defining symmetry transformations of the world state and robot control actions, respectively. Since the world state $s \in \mathcal{S}$ captures the robot state, its contact mode, and the control-task metrics, a symmetry-transformed world state $g \triangleright_S s$ describes the symmetry-transformed robot state, symmetry-transformed contact mode, and the symmetry-transformed control task. Importantly, both the original and symmetry-transformed control tasks are equivalent as a consequence of the \mathbb{G} -invariance of the world dynamics, which implies that $\mathbb{P}_{s|\mathbf{a}\mathbf{s}}^{\Delta t}$ is a \mathbb{G} -invariant conditional probability distribution (see eq. 6.4). A depiction of this principle is provided in fig. 8.1 for the case of two symmetry-related manipulation control tasks.

In this context, the MDP modelling the control of a robot inherits symmetry constraints that we state in the following definition and discuss in detail:

Definition 8.1 (Symmetric MDP). *An MDP is said to possess a symmetry group \mathbb{G} , if the state and action spaces admit group actions $\triangleright_S : \mathbb{G} \times \mathcal{S} \rightarrow \mathcal{S}$ and $\triangleright_A : \mathbb{G} \times \mathcal{A} \rightarrow \mathcal{A}$, such that the transition dynamics $\mathbb{P}_{s|\mathbf{a}\mathbf{s}}^{\Delta t}$, the bounded reward function r , and initial state distributions \mathbb{P}_s^0 are \mathbb{G} -invariant, and the admissible action map \mathbb{A}_π is a \mathbb{G} -equivariant map (definition 2.14). That is, if for all $g \in \mathbb{G}$, $s \in \mathcal{S}$, and $\mathbf{a} \in \mathcal{A}$ the following conditions hold:*

$$\mathbb{P}_{s|\mathbf{a}\mathbf{s}}^{\Delta t}(g \triangleright_S ds | g \triangleright_A \mathbf{a}, g \triangleright_S s) = \mathbb{P}_{s|\mathbf{a}\mathbf{s}}^{\Delta t}(ds | \mathbf{a}, s), \quad (8.4a)$$

$$r(g \triangleright_S s, g \triangleright_A \mathbf{a}) = r(s, \mathbf{a}) < \infty, \quad (8.4b)$$

$$\mathbb{P}_s^0(g \triangleright_S ds) = \mathbb{P}_s^0(ds) \quad (8.4c)$$

$$\mathbb{A}_\pi(g \triangleright_S s) = g \triangleright_A \mathbb{A}_\pi(s). \quad (8.4d)$$

It is straightforward to verify that control tasks for robots with a state/morphological symmetry group \mathbb{G} satisfy definition 8.1. Eq. 8.4a follows directly from the \mathbb{G} -equivariance of the underlying deterministic robot dynamics (see eq. 6.4), and eq. 8.4b follows from the equivalence of symmetry-related control tasks (see fig. 8.1). Eq. 8.4c is a natural operational assumption when the robot is equally likely to face a task or any of its symmetry-related variants—for instance, turning left or right while navigating, manipulating an object with the left or right arm (fig. 8.1), or locomoting forward or backward (fig. 4.1). Finally, eq. 8.4d follows from the definition of spatial symmetries (see chapter 3).

These MDP symmetry priors induce constraints on optimal control policies and on optimal value and quality functions. To see this, we begin by proving that the optimal quality function Q^* is \mathbb{G} -invariant.

Theorem 8.1 (\mathbb{G} -invariant optimal quality function). *For any \mathbb{G} -symmetric MDP satisfying definition 8.1, the optimal quality function Q^* is \mathbb{G} -invariant, that is:*

$$Q^*(g \triangleright_S \mathbf{s}, g \triangleright_A \mathbf{a}) = Q^*(\mathbf{s}, \mathbf{a}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}. \quad (8.5)$$

Proof. The proof follows the same rationale as Ravindran and Barto (2001, Theorem 1), generalized here to continuous state and action spaces. Define, for any $k \in \mathbb{N}$, the k -step discounted optimal quality function as:

$$Q_k^*(\mathbf{s}_0, \mathbf{a}_0) := r(\mathbf{s}_0, \mathbf{a}_0) + \gamma \int_{\mathcal{S}} \sup_{\mathbf{a}_1 \in \mathbb{A}_\pi(\mathbf{s}_1)} [Q_{k-1}^*(\mathbf{s}_1, \mathbf{a}_1)] \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(d\mathbf{s}_1 | \mathbf{s}_0, \mathbf{a}_0)$$

with $Q_0^*(\mathbf{s}, \mathbf{a}) := r(\mathbf{s}, \mathbf{a}) \quad \forall \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}.$

Here, the supremum is taken over the compact admissible-action set $\mathbb{A}_\pi(\mathbf{s}_1)$; hence, it is well defined under the standing assumptions. Note that, for notational simplicity, we use indices 0 and 1 instead of t and $t + \Delta t$, without loss of generality.

Let $(\mathbf{s}_0, \mathbf{a}_0) \in \mathcal{S} \times \mathcal{A}$ and define its symmetry-transformed pair as $(g \triangleright_S \mathbf{s}_0, g \triangleright_A \mathbf{a}_0)$ for any $g \in \mathbb{G}$. For $k = 0$, we have $Q_0^*(\mathbf{s}_0, \mathbf{a}_0) = r(\mathbf{s}_0, \mathbf{a}_0) = r(g \triangleright_S \mathbf{s}_0, g \triangleright_A \mathbf{a}_0) = Q_0^*(g \triangleright_S \mathbf{s}_0, g \triangleright_A \mathbf{a}_0)$ for all $g \in \mathbb{G}$, by reward \mathbb{G} -invariance (eq. 8.4b). Moreover, since the admissible-action map is \mathbb{G} -equivariant (eq. 8.4d), the set of optimal actions according to $Q_0^*(\mathbf{s}, \cdot)$ is symmetry-related to the set of optimal actions according to $Q_0^*(g \triangleright_S \mathbf{s}, \cdot)$, since:

$$\sup_{\mathbf{a} \in \mathbb{A}_\pi(g \triangleright_S \mathbf{s})} Q_0^*(g \triangleright_S \mathbf{s}, \mathbf{a}) = \sup_{\mathbf{a}' \in \mathbb{A}_\pi(\mathbf{s})} Q_0^*(g \triangleright_S \mathbf{s}, g \triangleright_A \mathbf{a}') = \sup_{\mathbf{a}' \in \mathbb{A}_\pi(\mathbf{s})} Q_0^*(\mathbf{s}, \mathbf{a}'). \quad \text{with } \mathbf{a} := g \triangleright_A \mathbf{a}'$$

For $k = 1$, we combine the result for $k = 0$ with the \mathbb{G} -invariance of the reward and transition kernel; and corollary 2.1 to obtain

$$\begin{aligned} Q_1^*(g \triangleright_S \mathbf{s}_0, g \triangleright_A \mathbf{a}_0) &= r(g \triangleright_S \mathbf{s}_0, g \triangleright_A \mathbf{a}_0) + \gamma \int_{\mathcal{S}} \sup_{\mathbf{a}_1 \in \mathbb{A}_\pi(\mathbf{s}_1)} [Q_0^*(\mathbf{s}_1, \mathbf{a}_1)] \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(d\mathbf{s}_1 | g \triangleright_S \mathbf{s}_0, g \triangleright_A \mathbf{a}_0) \\ &= r(\mathbf{s}_0, \mathbf{a}_0) + \gamma \int_{\mathcal{S}} \sup_{\mathbf{a}_1 \in \mathbb{A}_\pi(g^{-1} \triangleright_S \mathbf{s}_1)} [Q_0^*(g^{-1} \triangleright_S \mathbf{s}_1, \mathbf{a}_1)] \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(g^{-1} \triangleright_S d\mathbf{s}_1 | \mathbf{s}_0, \mathbf{a}_0) \\ &= r(\mathbf{s}_0, \mathbf{a}_0) + \gamma \int_{\mathcal{S}} \sup_{\mathbf{a}'_1 \in \mathbb{A}_\pi(\mathbf{s}_1)} [Q_0^*(\mathbf{s}_1, \mathbf{a}'_1)] \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(d\mathbf{s}_1 | \mathbf{s}_0, \mathbf{a}_0) = Q_1^*(\mathbf{s}_0, \mathbf{a}_0) \end{aligned}$$

Applying the same argument inductively, invariance of Q_{k-1}^* implies invariance of Q_k^* . Hence, Q_k^* is \mathbb{G} -invariant for all $k \in \mathbb{N}$. Since the reward is bounded, $\{Q_k^*\}_{k=0}^\infty$ converges to Q^* , which is therefore \mathbb{G} -invariant. \square

As a direct result of the \mathbb{G} -invariance of the optimal quality function, we obtain the \mathbb{G} -invariance of the optimal value function:

Corollary 8.1 (\mathbb{G} -invariant optimal value function). *For any \mathbb{G} -symmetric MDP satisfying definition 8.1, the optimal value function V^* is \mathbb{G} -invariant, that is:*

$$V^*(g \triangleright_S \mathbf{s}) = V^*(\mathbf{s}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}. \quad (8.6)$$

Proof. Define the k -step discounted optimal value function as:

$$V_k^*(\mathbf{s}) := \sup_{\mathbf{a} \in \mathbb{A}_\pi(\mathbf{s})} [Q_k^*(\mathbf{s}, \mathbf{a})] \quad \forall k \in \mathbb{N}, \mathbf{s} \in \mathcal{S}.$$

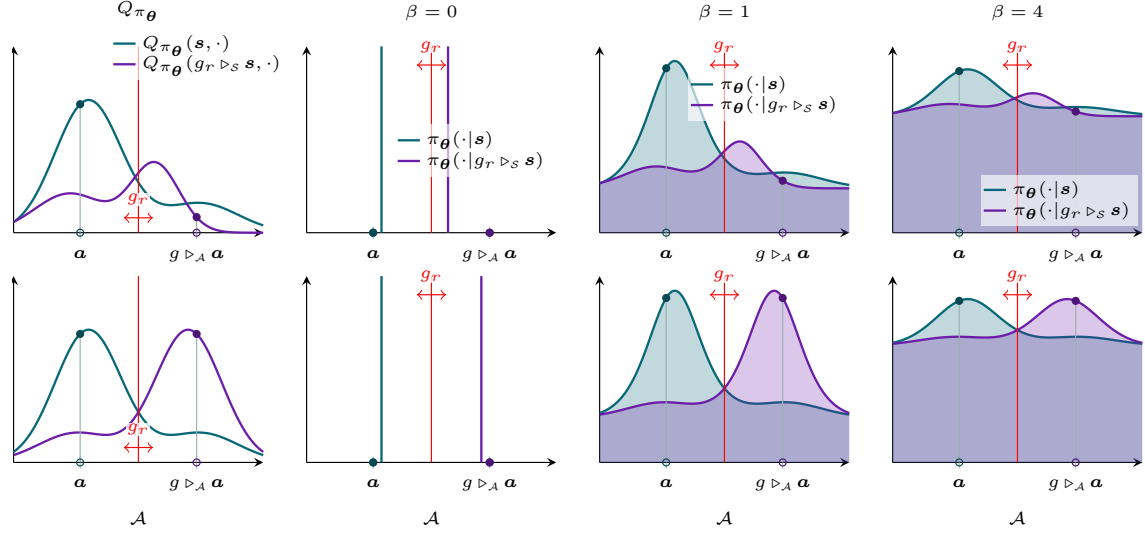


Figure 8.2 Control policies induced by an arbitrary quality function (top row) and a \mathbb{C}_2 -invariant quality function (bottom row) for a \mathbb{C}_2 -symmetric MDP with one-dimensional actions, where reflections act as $g_r \triangleright_{\mathcal{A}} \mathbf{a} = -\mathbf{a}$. Columns 2–4 show how the temperature parameter β reshapes probability mass in the Boltzmann parametrization (eq. 8.7).

Then following the results from [theorem 8.1](#), we have that V_k^* is \mathbb{G} -invariant for all $k \in \mathbb{N}$, and by the boundedness of the reward, $\{V_k^*\}_{k=0}^{\infty}$ converges to V^* , which is therefore \mathbb{G} -invariant. \square

Symmetry constraints on optimal control policies Having characterized the \mathbb{G} -invariance of the unique optimal value and quality functions, we now study the implications of our MDP symmetry priors for the family of stochastic optimal control policies. Recall that the optimal quality function $Q^* : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ quantifies the expected discounted cumulative reward obtained by following an optimal control policy π_* , starting from any state-action pair (\mathbf{s}, \mathbf{a}) . In this context, we can interpret each quality function $Q_{\pi_{\theta}} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ as defining a stochastic control policy via the Gibbs/Boltzmann distribution, such that:

$$\pi_{\theta}(\mathbb{A}|\mathbf{s}) := \frac{\int_{\mathbb{A}} \exp(Q_{\pi_{\theta}}(\mathbf{s}, \mathbf{a})/\beta) \mathbb{P}_{\mathbf{a}}(d\mathbf{a})}{p(\mathbf{s})} \quad \forall \mathbf{s} \in \mathcal{S}, \quad \text{with } p(\mathbf{s}) := \int_{\mathcal{A}} \exp(Q_{\pi_{\theta}}(\mathbf{s}, \mathbf{a}')/\beta) \mathbb{P}_{\mathbf{a}}(d\mathbf{a}'). \quad (8.7)$$

Here, $\mathbb{P}_{\mathbf{a}}$ is the marginal probability distribution of control actions, which is \mathbb{G} -invariant as a consequence of [eq. 8.4d](#); $p(\mathbf{s})$ is the normalizing partition function; and $\beta \in [0, \infty)$ is a temperature parameter that controls the stochasticity of the policy.

This parametrization offers strong theoretical flexibility, as it allows us to recover symmetry constraints on optimal deterministic and stochastic policies ($\beta = 0$), as well as optimal stochastic policies in the context of entropy regularized RL ($\beta > 0$) ([Sutton and Barto, 2018](#)). To see this, consider that the temperature parameter $\beta \in [0, \infty)$ controls the stochasticity of the policy in the following fashion (see [fig. 8.2](#)):

$\beta \rightarrow 0$: The policy assigns all probability mass to the set of optimal control actions, defined as $\mathbb{A}_{\pi}^*(\mathbf{s}) := \arg \max_{\mathbf{a} \in \mathbb{A}_{\pi}(\mathbf{s})} Q_{\pi_{\theta}}(\mathbf{s}, \mathbf{a})$ for all $\mathbf{s} \in \mathcal{S}$. If the set contains a single action for all \mathbf{s} , the policy becomes deterministic, otherwise its an optimal stochastic policy with uniform distribution over $\mathbb{A}_{\pi}^*(\mathbf{s})$.

- $\beta = 1$: The policy is a stochastic Gibbs/Boltzmann distribution over actions, with the optimal quality function as energy function.
- $\beta > 1$: The probability mass becomes increasingly smoothed across the action space, reducing the policy's exploitation of the quality function. In the limit as $\beta \rightarrow \infty$, the policy converges to the marginal distribution $\mathbb{P}_{\mathbf{a}}$.

Crucially, this parametrization also enables us to recover symmetry constraints on the Gibbs/Boltzmann policy induced by the optimal quality function Q^* (theorem 8.1).

Corollary 8.2 (Invariance of optimal stochastic policies). *For any \mathbb{G} -symmetric MDP satisfying definition 8.1, an optimal stochastic control policy is \mathbb{G} -invariant, that is:*

$$\pi_*(\mathbb{A}|\mathbf{s}) = \pi_*(g \triangleright_{\mathcal{A}} \mathbb{A} | g \triangleright_{\mathcal{S}} \mathbf{s}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbb{A} \subseteq \mathcal{A}.$$

Proof. By theorem 8.1, the optimal quality function Q^* is \mathbb{G} -invariant, which implies that any optimal control policy defined by Q^* is \mathbb{G} -invariant. Note that for $\beta \rightarrow \infty$, the distribution remains invariant as we assume the marginal distribution of actions $\mathbb{P}_{\mathbf{a}}$ is \mathbb{G} -invariant (definition 2.17). \square

This result extends directly to deterministic policies in the limit $\beta \rightarrow 0$, provided that ties within the optimal-action sets are resolved by a \mathbb{G} -equivariant selection rule. Recall that there may exist a family of deterministic optimal control policies associated with Q^* (Ravindran and Barto, 2001). A standard example is a driving task where obstacle avoidance can be achieved by turning either left or right; thus, two deterministic optimal control policies exist, each turning in a different direction. In the limit $\beta \rightarrow 0$, the optimal stochastic policy in corollary 8.2 converges, for every state $\mathbf{s} \in \mathcal{S}$, to a uniform distribution over optimal control actions $\mathbb{A}_{\pi}^*(\mathbf{s})$ —that is, a policy equally likely to turn left or right in this toy example. If a mechanism is introduced to select a single optimal action and this mechanism is \mathbb{G} -equivariant, the resulting deterministic policy $\hat{\pi}_* : \mathcal{S} \rightarrow \mathcal{A}$ is \mathbb{G} -equivariant by construction (definition 2.14). We summarize this result in the following corollary:

Corollary 8.3 (\mathbb{G} -equivariant optimal deterministic policy under equivariant tie-breaking). *For any \mathbb{G} -symmetric MDP satisfying definition 8.1, if one selects an action from each optimal-action set $\mathbb{A}_{\pi}^*(\mathbf{s})$ through a \mathbb{G} -equivariant selection rule, then the induced deterministic control policy $\hat{\pi}_* : \mathcal{S} \rightarrow \mathcal{A}$ is optimal and \mathbb{G} -equivariant, that is:*

$$g \triangleright_{\mathcal{A}} \hat{\pi}_*(\mathbf{s}) = \hat{\pi}_*(g \triangleright_{\mathcal{S}} \mathbf{s}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}.$$

Proof. By corollary 8.2, when $\beta \rightarrow 0$, the optimal stochastic policy converges to a uniform distribution over optimal control actions $\mathbb{A}_{\pi}^*(\mathbf{s})$ for all $\mathbf{s} \in \mathcal{S}$. If there is a unique optimal action, $|\mathbb{A}_{\pi}^*(\mathbf{s})| = 1$, for all $\mathbf{s} \in \mathcal{S}$, then the optimal deterministic policy is unique and \mathbb{G} -equivariant by construction.

If there exist states $\mathbf{s} \in \mathcal{S}$ with multiple optimal actions, $|\mathbb{A}_{\pi}^*(\mathbf{s})| > 1$, then symmetry-related states also have multiple optimal actions, i.e., $|\mathbb{A}_{\pi}^*(g \triangleright_{\mathcal{S}} \mathbf{s})| > 1$, for all $g \in \mathbb{G}$, by the \mathbb{G} -invariance of Q^* . Any selection rule picking an element from each set $\mathbb{A}_{\pi}^*(\mathbf{s})$ induces an optimal deterministic policy.

Define a \mathbb{G} -equivariant selection map of optimal actions per state, $\mathbb{A}_{\pi}^*(\mathbf{s}) \rightarrow \mathbf{a}^*(\mathbf{s})$, satisfying $\mathbf{a}^*(g \triangleright_{\mathcal{S}} \mathbf{s}) = g \triangleright_{\mathcal{A}} \mathbf{a}^*(\mathbf{s})$ for all $g \in \mathbb{G}$. The resulting optimal deterministic policy $\hat{\pi}_*(\mathbf{s}) := \mathbf{a}^*(\mathbf{s})$ is therefore \mathbb{G} -equivariant by construction (definition 2.14). \square

Symmetry constraints on marginal distributions Next, we characterize a fundamental implication for the distributions of state and action trajectories induced by *any* \mathbb{G} -invariant stochastic control policy in a \mathbb{G} -symmetric MDP. In a nutshell, the result states that for any \mathbb{G} -invariant policy, the marginal distributions of states and actions at any time $t \in \mathbb{R}$ induced by that policy are \mathbb{G} -invariant. Given that this result is fundamental for theoretically justifying the use of data augmentation in data-driven optimal control and the symmetry constraints of imitation learning algorithms (section 8.3), we state and prove it as a standalone proposition.

Proposition 8.1 (Marginal \mathbb{G} -invariance through time). *Let $(\mathcal{S}, \mathcal{A}, \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}, r, \mathbb{P}_{\mathbf{s}}^0, \mathbb{A}_{\pi})$ be a \mathbb{G} -symmetric MDP (definition 8.1), and let π be any \mathbb{G} -invariant stochastic control policy. Then, the marginal distributions of states and actions at any time $t \in \mathbb{R}$ induced by π are also \mathbb{G} -invariant, that is:*

$$\begin{aligned} \mathbb{P}_{\mathbf{s}}^0(\mathbb{S}) = \mathbb{P}_{\mathbf{s}}^0(g \triangleright_{\mathcal{S}} \mathbb{S}), \quad \text{and} \quad & \implies \quad \mathbb{P}_{\mathbf{s}}^t(\mathbb{S}) = \mathbb{P}_{\mathbf{s}}^t(g \triangleright_{\mathcal{S}} \mathbb{S}), \\ \pi(\mathbb{A}|\mathbf{s}) = \pi(g \triangleright_{\mathcal{A}} \mathbb{A}|g \triangleright_{\mathcal{S}} \mathbf{s}) & \implies \quad \mathbb{P}_{\mathbf{a}}^t(\mathbb{A}) = \mathbb{P}_{\mathbf{a}}^t(g \triangleright_{\mathcal{A}} \mathbb{A}) \quad \forall t \in \mathbb{R}, g \in \mathbb{G}, \mathbb{S} \subseteq \mathcal{S}, \mathbb{A} \subseteq \mathcal{A}. \end{aligned}$$

Proof. The proof follows by induction using prop. 2.3. For $t = 0$, the \mathbb{G} -invariance of the initial state distribution $\mathbb{P}_{\mathbf{s}}^0$ (eq. 8.4c) together with the \mathbb{G} -invariance of π implies, by prop. 2.3, that the marginal distribution of actions at time $t = 0$ is \mathbb{G} -invariant,

$$\mathbb{P}_{\mathbf{a}}^0(\mathbb{A}) = \int_{\mathcal{S}} \pi(\mathbb{A}|\mathbf{s}) \mathbb{P}_{\mathbf{s}}^0(d\mathbf{s}) = \int_{\mathcal{S}} \pi(g \triangleright_{\mathcal{A}} \mathbb{A}|g \triangleright_{\mathcal{S}} \mathbf{s}) \mathbb{P}_{\mathbf{s}}^0(g \triangleright_{\mathcal{S}} d\mathbf{s}) = \mathbb{P}_{\mathbf{a}}^0(g \triangleright_{\mathcal{A}} \mathbb{A}).$$

The \mathbb{G} -invariance of both the initial state and action distributions at time $t = 0$, together with the \mathbb{G} -invariance of the transition kernel (eq. 8.4a), implies the \mathbb{G} -invariance of the marginal distribution of states at time $t = \Delta t$ (see prop. 2.3):

$$\begin{aligned} \mathbb{P}_{\mathbf{s}}^{\Delta t}(\mathbb{S}) &= \int_{\mathcal{S}} \int_{\mathcal{A}} \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbb{S}|\mathbf{s}, \mathbf{a}) \pi(d\mathbf{a}|\mathbf{s}) \mathbb{P}_{\mathbf{s}}^0(d\mathbf{s}) \\ &= \int_{\mathcal{S}} \int_{\mathcal{A}} \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(g \triangleright_{\mathcal{S}} \mathbb{S}|g \triangleright_{\mathcal{S}} \mathbf{s}, g \triangleright_{\mathcal{A}} \mathbf{a}) \pi(g \triangleright_{\mathcal{A}} d\mathbf{a}|g \triangleright_{\mathcal{S}} \mathbf{s}) \mathbb{P}_{\mathbf{s}}^0(g \triangleright_{\mathcal{S}} d\mathbf{s}) = \mathbb{P}_{\mathbf{s}}^{\Delta t}(g \triangleright_{\mathcal{S}} \mathbb{S}). \end{aligned}$$

Then, the \mathbb{G} -invariance of the marginal state and action distributions at time $t = \Delta t$ together with the \mathbb{G} -invariance of the transition kernel implies the \mathbb{G} -invariance of the marginal distribution of states at time $t = k\Delta t$, for any $k \in \mathbb{N}$, by applying the same argument inductively. \square

Intuitively, this property of preserving the symmetries of state distributions through time can be interpreted as a characteristic property of symmetry-consistent optimal control policies in \mathbb{G} -symmetric MDPs, given that the Gibbs/Boltzmann optimal policy from corollary 8.2 is \mathbb{G} -invariant. In practical terms, the relevance of this result stems from characterizing that the distributions of optimal state-action trajectories are constrained to be \mathbb{G} -invariant, which is a crucial theoretical requirement for leveraging these symmetries in imitation and offline learning in section 8.3, leading to substantial improvements in sample efficiency and policy robustness.

Summary 8.1

The dynamics of contact-rich control tasks for robots with state/morphological symmetries can be modelled as a \mathbb{G} -symmetric MDP satisfying the symmetry priors in definition 8.1, including \mathbb{G} -invariance of the transition dynamics $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$, reward function r , and initial state distribution $\mathbb{P}_{\mathbf{s}}^0$, and \mathbb{G} -equivariance of the admissible action map \mathbb{A}_{π} .

Under these symmetry priors, the optimal quality function is constrained to be \mathbb{G} -invariant (theorem 8.1):

$$Q^*(g \triangleright_S \mathbf{s}, g \triangleright_{\mathcal{A}} \mathbf{a}) = Q^*(\mathbf{s}, \mathbf{a}), \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A},$$

Consequently, the optimal value function is also \mathbb{G} -invariant (corollary 8.1):

$$V^*(g \triangleright_S \mathbf{s}) = V^*(\mathbf{s}), \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S},$$

and symmetry-consistent optimal policy parametrizations are constrained to be \mathbb{G} -invariant (corollary 8.2):

$$\pi_*(\mathbb{A}|\mathbf{s}) = \pi_*(g \triangleright_{\mathcal{A}} \mathbb{A} | g \triangleright_S \mathbf{s}), \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbb{A} \subseteq \mathcal{A}.$$

8.2 Symmetry priors in optimal control under partial observability

As detailed in section 2.2, the world state in contact-rich robotic control tasks is only partially observable, since several quantities of interest cannot be directly measured by onboard proprioceptive and exteroceptive sensors. Among these unobservable or partially observable quantities is the robot contact mode—which determines the number of contacts at a given time t , the associated contact surfaces, friction regimes, and resultant contact forces.

Formally, the set of onboard proprioceptive and exteroceptive sensor measurements defines a state observation $\mathbf{o}_t \in \mathcal{O}$, assumed to be generated at each timestep by sampling $\mathbf{o}_t \sim \mathbb{P}_{\mathbf{o}|\mathbf{s}}(\cdot|\mathbf{s}_t)$ from a stochastic observation model $\mathbb{P}_{\mathbf{o}|\mathbf{s}} : \Sigma_{\mathcal{O}} \times \mathcal{S} \mapsto [0, 1]$. Such partially observable control problems are commonly modelled as a POMDP defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}, r, \mathbb{P}_{\mathbf{o}|\mathbf{s}}, \mathbb{P}_{\mathbf{s}}^0, \mathbb{A}_{\pi}, \gamma)$, composed of the underlying MDP $(\mathcal{S}, \mathcal{A}, \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}, r, \mathbb{P}_{\mathbf{s}}^0, \mathbb{A}_{\pi}, \gamma)$ (see section 8.1) together with the observation space \mathcal{O} and observation model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$.

To characterize how the symmetry priors of a \mathbb{G} -symmetric MDP (definition 8.1) translate to the associated POMDP, we study the *belief MDP*, which recasts optimal control under partial observability as an optimal-control problem over state-observation dynamics. We then characterize the relation between the belief MDP and a \mathbb{G} -symmetric MDP, yielding analogous \mathbb{G} -invariance constraints on the optimal belief value and quality functions, and on the family of optimal control policies.

Belief MDP We define a belief MDP as a POMDP with an associated optimal state-estimation model, $\mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \vec{\mathcal{O}} \mapsto [0, 1]$, that uses the full history of past observations $\vec{\mathbf{o}}_t := (\mathbf{o}_0, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}$ to define the probability distribution of the current world state $\mathbf{s}_t \in \mathcal{S}$ at time t . Thus, $\mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t}(\mathbb{S}|\vec{\mathbf{o}}_t)$ denotes the probability that $\mathbf{s}_t \in \mathbb{S}$ given the history of observations $\vec{\mathbf{o}}_t$. Here, $\vec{\mathcal{O}}$ denotes the space of finite past observation trajectories. See section 2.2 for details.

The optimal state-estimation model enables us to model state-observation dynamics using only directly measurable quantities—namely, the observation history $\vec{\mathbf{o}}_t$ and the control action \mathbf{a}_t at time t —by marginalizing out the unobservable world state \mathbf{s}_t (see section 2.2), that is:

$$\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathbf{o}}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \vec{\mathbf{o}}_t) = \int_{\mathcal{S}} \underbrace{\left(\int_{\mathcal{S}} \mathbb{P}_{\mathbf{o}|\mathbf{s}}(\mathbb{O}|\mathbf{s}_{t+\Delta t}) \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(d\mathbf{s}_{t+\Delta t}|\mathbf{a}_t, \hat{\mathbf{s}}_t) \right)}_{\mathbb{P}_{\mathbf{o}|\mathbf{a}\mathbf{s}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \hat{\mathbf{s}}_t)} \underbrace{\mathbb{P}_{\mathbf{s}|\vec{\mathbf{o}}}^{\Delta t}(d\hat{\mathbf{s}}_t|\vec{\mathbf{o}}_t)}_{\text{Optimal state estimation}}, \quad (8.8)$$

$$\forall \mathbb{O} \subseteq \mathcal{O}, \mathbf{a}_t \in \mathcal{A}, t \in \mathbb{R}, \vec{\mathbf{o}}_t := (\mathbf{o}_0, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}.$$

Here, $\mathbb{P}_{\mathbf{o}|\mathbf{as}}^{\Delta t}(\mathbb{O}|\mathbf{a}_t, \hat{\mathbf{s}}_t) := \mathbb{P}(\mathbf{o}_{t+\Delta t} \in \mathbb{O} | \mathbf{a}_t = \mathbf{a}_t, \mathbf{s}_t = \hat{\mathbf{s}}_t)$. We denote estimated/belief states by $\hat{\mathbf{s}}_t \in \mathcal{S}$ to distinguish them from the true partially observable world states $\mathbf{s}_t \in \mathcal{S}$.

In this context, we define the belief MDP by the tuple $(\vec{\mathcal{O}}, \mathcal{A}, \mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathcal{O}}}^{\Delta t}, \hat{r}, \mathbb{P}_{\mathbf{o}^*}^0, \mathbb{O}_\pi^*, \gamma)$. Its state space is the trajectory space $\vec{\mathcal{O}}$, and its action space is the same as that of the underlying MDP, \mathcal{A} . Its transition dynamics are given by $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathcal{O}}}^{\Delta t}$ in eq. 8.8. The initial observation-trajectory distribution $\mathbb{P}_{\mathbf{o}^*}^0$ is induced by the state initial distribution $\mathbb{P}_{\mathbf{s}}^0$, and the admissible control-action map $\mathbb{O}_\pi^* : \vec{\mathcal{O}} \rightarrow \Sigma_{\mathcal{A}}$ is defined from the underlying MDP admissible action map $\mathbb{A}_\pi : \mathcal{S} \rightarrow \Sigma_{\mathcal{A}}$. Finally, the belief reward function $\hat{r} : \vec{\mathcal{O}} \times \mathcal{A} \rightarrow \mathbb{R}$ is defined as the reward expectation over the belief/state-estimation distribution:

$$\hat{r}(\vec{\mathcal{O}}, \mathbf{a}) := \int_{\mathcal{S}} r(\hat{\mathbf{s}}, \mathbf{a}) \mathbb{P}_{\mathbf{s}|\vec{\mathcal{O}}}^{\Delta t}(d\hat{\mathbf{s}}|\vec{\mathcal{O}}) \quad \forall \vec{\mathcal{O}} \in \vec{\mathcal{O}}, \mathbf{a} \in \mathcal{A}, \quad (8.9)$$

In this context, a control policy depends only on the observable history of past observations, $\pi_\theta : \Sigma_{\mathcal{A}} \times \vec{\mathcal{O}} \mapsto [0, 1]$, and the policy quality function is defined as:

$$\hat{Q}_{\pi_\theta}(\vec{\mathcal{O}}_t, \mathbf{a}_t) = \hat{r}(\vec{\mathcal{O}}_t, \mathbf{a}_t) + \gamma \int_{\mathcal{O}} \hat{V}_{\pi_\theta}(\vec{\mathcal{O}}_{t+\Delta t}) \mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathcal{O}}}^{\Delta t}(d\mathbf{o}_{t+\Delta t}|\mathbf{a}_t, \vec{\mathcal{O}}_t) \quad \forall \vec{\mathcal{O}}_t \in \vec{\mathcal{O}}, \mathbf{a}_t \in \mathcal{A}, \quad (8.10)$$

and the belief value function is defined as the expected quality over the policy distribution of actions:

$$\hat{V}_{\pi_\theta}(\vec{\mathcal{O}}_t) = \int_{\mathcal{A}} \hat{Q}_{\pi_\theta}(\vec{\mathcal{O}}_t, \mathbf{a}_t) \pi_\theta(d\mathbf{a}_t|\vec{\mathcal{O}}_t) := \mathbb{E}_{\pi_\theta} \left(\sum_{k=0}^{\infty} \gamma^k \hat{r}(\vec{\mathcal{O}}_{t+k\Delta t}, \mathbf{a}_{t+k\Delta t}) \right) \quad \forall \vec{\mathcal{O}}_t \in \vec{\mathcal{O}}. \quad (8.11)$$

\mathbb{G} -symmetric belief MDP We now characterize the conditions under which the belief MDP modelling a partially observable, contact-rich control problem for a robot with environmental or morphological symmetries is itself a \mathbb{G} -symmetric MDP satisfying the symmetry priors in definition 8.1. This allows us to apply the results from section 8.1 directly to the belief MDP and obtain analogous \mathbb{G} -invariance constraints on the optimal belief value and quality functions, and on the family of optimal control policies.

As detailed in chapter 6, under the reasonable assumption that the observation space is endowed with a group action $\triangleright_{\mathcal{O}} : \mathbb{G} \times \mathcal{O} \mapsto \mathcal{O}$ and that the stochastic observation sensor model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ is \mathbb{G} -invariant (see discussion in assumption 6.1), the predictable state-observation dynamics defined by $\mathbb{P}_{\mathbf{o}|\mathbf{a}\vec{\mathcal{O}}}^{\Delta t}$ in eq. 8.8 become a \mathbb{G} -invariant conditional probability distribution (definition 2.17). The belief reward function $\hat{r} : \vec{\mathcal{O}} \times \mathcal{A} \rightarrow \mathbb{R}$ in eq. 8.9 is also \mathbb{G} -invariant by reward-function invariance (eq. 8.4b); the initial trajectory distribution $\mathbb{P}_{\mathbf{o}^*}^0$ is \mathbb{G} -invariant by initial-state-distribution invariance (eq. 8.4c); and the admissible action map \mathbb{O}_π^* is \mathbb{G} -equivariant by admissible-action-map equivariance (eq. 8.4d). Consequently, by translating theorem 8.1 and corollaries 8.1 and 8.2 to the belief-MDP setting, we guarantee that the optimal belief quality function \hat{Q}^* and optimal belief value function \hat{V}^* are \mathbb{G} -invariant, while the corresponding symmetry-constrained policy parametrizations inherit the same \mathbb{G} -invariance requirements.

Truncated history belief MDP In practice, when modelling or controlling dynamical/robotic systems using data-driven methods, memory and compute limitations prevent us from relying on the entire history of past observations $\vec{\mathcal{O}}_t = (\mathbf{o}_0, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}$. Instead, we rely on a truncated history of past observations $\vec{\mathcal{O}}_t := (\mathbf{o}_{t_0}, \dots, \mathbf{o}_t) \in \mathcal{O}^H$, with $t_0 = \max(0, t - H + 1)$ and $H \in \mathbb{N}$ denoting the

observation-history horizon. In this context, there exists an optimal state-estimation model that relies only on truncated observation histories, $\mathbb{P}_{\mathbf{s}|\bar{\mathbf{o}}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{O}^H \mapsto [0, 1]$, thereby defining a truncated-history belief MDP. The analysis of this section then applies directly, and the same \mathbb{G} -invariance constraints hold for the optimal belief value and quality functions, as well as for the family of optimal control policies.

Summary 8.2

In practice, contact-rich control tasks in robotics are partially observable. For robots with state/morphological symmetry group \mathbb{G} , under the assumption that the observation space is endowed with a group action $\triangleright_{\mathcal{O}} : \mathbb{G} \times \mathcal{O} \mapsto \mathcal{O}$ and that the sensor model $\mathbb{P}_{\mathbf{o}|\mathbf{s}}$ is \mathbb{G} -invariant (assumption 6.1), that is

$$\mathbb{P}_{\mathbf{o}|\mathbf{s}}(g \triangleright_{\mathcal{O}} \mathbb{O} | g \triangleright_{\mathcal{S}} \mathbf{s}) = \mathbb{P}_{\mathbf{o}|\mathbf{s}}(\mathbb{O} | \mathbf{s}) \quad \forall g \in \mathbb{G}, \mathbf{s} \in \mathcal{S}, \mathbb{O} \subseteq \mathcal{O},$$

the control problem can be modelled as a \mathbb{G} -symmetric belief MDP over the space of (potentially truncated) past observation trajectories $\vec{\mathcal{O}}$, with control policies that depend only on the history of past observations, $\pi_{\theta} : \Sigma_{\mathcal{A}} \times \vec{\mathcal{O}} \mapsto [0, 1]$. Consequently, even under partial observability, the optimal belief quality function \hat{Q}^* and optimal belief value function \hat{V}^* are \mathbb{G} -invariant, and the family of optimal stochastic control policies is subject to the same \mathbb{G} -invariance constraint, that is:

$$\pi_*(\mathbb{A} | \vec{\mathbf{o}}) = \pi_*(g \triangleright_{\mathcal{A}} \mathbb{A} | g \triangleright_{\vec{\mathcal{O}}} \vec{\mathbf{o}}) \quad \forall g \in \mathbb{G}, \vec{\mathbf{o}} := (\mathbf{o}_0, \dots, \mathbf{o}_t) \in \vec{\mathcal{O}}, \mathbb{A} \subseteq \mathcal{A}.$$

8.3 Leveraging symmetries in data-driven optimal control

This section shows how to leverage the symmetry priors from sections 8.1 and 8.2 in data-driven optimal control. In section 8.3.1, we formulate behaviour cloning as a symmetry-constrained problem and study symmetry-aware flow-matching policies for Push-T and bimanual mobile manipulation under bilateral reflection symmetry \mathbb{C}_2 . In section 8.3.2, we cast model-free RL objectives as symmetry-constrained optimization problems and study PPO-based methods for quadruped locomanipulation under \mathbb{C}_2 and dexterous manipulation under \mathbb{C}_2 and \mathbb{C}_4 . Across both settings, symmetry-aware methods improve sample efficiency, robustness, and generalization.

8.3.1 Imitation learning

Editorial Note: Section 8.3.1

The analysis and results from this section are based on the paper [Siebenborn et al. \(2026\)](#), done in collaboration with the PhD student Max Siebenborn, and other authors. Details on my personal contribution to the paper are provided in [section 1.5](#).

This section studies the implications of state and morphological symmetry priors in imitation learning—an offline learning paradigm in which a parametric policy π_{θ} is trained to approximate an expert policy π_{exp} from a dataset of expert demonstrations. In particular, for robotic systems with an environmental or morphological symmetry group \mathbb{G} , we show that the expert policy can be assumed \mathbb{G} -invariant (assumption 8.1), which frames imitation learning as a constrained problem.

Although the results apply to arbitrary symmetry groups, our experiments focus on bilateral reflection symmetry—the most common case in practice, since most robots (humanoids, quadrupeds, mobile manipulators, etc.) feature $\mathbb{G} = \mathbb{C}_2 = \{e, g_r | g_r^2 = e\}$ (see fig. 2.1). Leveraging these priors yields imitation-learning policies that are ambidextrous by construction: they can perform a manipulation or locomotion task and its mirrored version with identical control performance. A simple example is humanoid handwriting, or any dexterous bimanual manipulation task, where the robot can perform the task with either hand (see figs. 4.1 and 8.1).

As detailed in sections 2.2 and 8.2, contact-rich control tasks are modelled as a POMDP $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathbb{P}_{\mathcal{S}|\mathcal{A}\mathcal{S}}^{\Delta t}, r, \mathbb{P}_{\mathcal{O}|\mathcal{S}}, \mathbb{P}_{\mathcal{S}}^0, \mathbb{A}_\pi, \gamma)$, with $\mathbb{P}_{\mathcal{O}|\mathcal{S}} : \Sigma_{\mathcal{O}} \times \mathcal{S} \mapsto [0, 1]$ denoting the stochastic observation model. Since the world state is only partially observable, we assume that a history of H past observations, $\vec{o}_t := (o_{t_0}, \dots, o_t) \in \mathcal{O}^H$ with $t_0 = \max\{0, t - H + 1\}$, provides sufficient statistics for estimating the current state (see section 2.2). The expert control policy is then modelled as a conditional distribution of actions given observation history, $\pi_{\text{exp}} : \Sigma_{\mathcal{A}} \times \mathcal{O}^H \mapsto [0, 1]$. Under this assumption, we recast the POMDP as a truncated-history belief MDP $(\mathcal{O}^H, \mathcal{A}, \mathbb{P}_{\mathcal{O}|\mathcal{A}\vec{o}}^{\Delta t}, \hat{r}, \mathbb{P}_{\mathcal{O}^H}^0, \mathbb{O}_\pi^H, \gamma)$, where $\mathbb{P}_{\mathcal{O}|\mathcal{A}\vec{o}}^{\Delta t}$ denotes predictable observation dynamics (eq. 8.8), \hat{r} denotes the belief reward (eq. 8.9), $\mathbb{P}_{\mathcal{O}^H}^0$ denotes the initial trajectory distribution, and \mathbb{O}_π^H denotes the admissible action map (see section 8.2).

Behaviour cloning In its simplest form, imitation learning is a behaviour-cloning problem where the goal is to learn a parametric stochastic policy $\pi_\theta(\cdot|\vec{o})$ that approximates an expert policy $\pi_{\text{exp}}(\cdot|\vec{o})$ for each $\vec{o} \in \mathcal{O}^H$. This can be formalized as maximizing likelihood under the joint distribution of observation-history–action pairs induced by the expert policy (Ziebart et al., 2008), or equivalently as minimizing the Kullback–Leibler divergence between expert and learned policy distributions for each conditioning:

$$\begin{aligned} \theta^* &= \arg \max_{\theta} \mathcal{L}_{\text{bc}}(\pi_\theta) := \arg \max_{\theta} \prod_{\vec{o}, \mathbf{a} \sim \mathbb{P}_{\mathcal{O}^H \mathcal{A}}^{\text{exp}}} \pi_\theta(\mathbf{a}|\vec{o}) \\ &\equiv \arg \min_{\theta} \mathbb{E}_{\vec{o} \sim \mathbb{P}_{\mathcal{O}^H \mathcal{A}}^{\text{exp}}} D_{\text{KL}}(\pi_{\text{exp}}(\cdot|\vec{o}) \| \pi_\theta(\cdot|\vec{o})). \end{aligned} \quad (8.12)$$

Here $\mathbb{P}_{\mathcal{O}^H \mathcal{A}}^{\text{exp}} : \Sigma_{\mathcal{O}^H} \times \Sigma_{\mathcal{A}} \mapsto [0, 1]$ denotes the joint distribution over observation histories and actions induced by the expert policy π_{exp} . In practice, eq. 8.12 is approximated from a dataset of expert demonstrations composed of N i.i.d. samples from $\mathbb{P}_{\mathcal{O}^H \mathcal{A}}^{\text{exp}}$, denoted by $\mathbb{D}_H^{\text{exp}} := \{(\vec{o}_i, \mathbf{a}_i) \sim \mathbb{P}_{\mathcal{O}^H \mathcal{A}}^{\text{exp}}\}_{i=1}^N$.

Symmetry priors in imitation learning As detailed in sections 8.1 and 8.2, if a robotic system has a state/morphological symmetry group \mathbb{G} and the observation model $\mathbb{P}_{\mathcal{O}|\mathcal{S}}$ is \mathbb{G} -invariant (assumption 6.1), then the belief MDP for our contact-rich control problem is \mathbb{G} -symmetric and satisfies definition 8.1. Consequently, the family of optimal stochastic policies are constrained to be \mathbb{G} -invariant (corollary 8.2). In this context, it is reasonable to assume that the (potentially suboptimal) expert policy π_{exp} is a \mathbb{G} -invariant conditional distribution (definition 2.17). We state this assumption formally:

Assumption 8.1 (\mathbb{G} -invariance of expert control policies). *The expert control policy π_{exp} is a \mathbb{G} -invariant control policy satisfying:*

$$\pi_{exp}(\mathbb{A}|\vec{\sigma}) = \pi_{exp}(g \triangleright_{\mathcal{A}} \mathbb{A} | g \triangleright_{\mathcal{O}^H} \vec{\sigma}) \quad \forall g \in \mathbb{G}, \vec{\sigma} \in \mathcal{O}^H, \mathbb{A} \subseteq \mathcal{A}.$$

with $g \triangleright_{\mathcal{O}^H} \vec{\sigma}_t := (g \triangleright_{\mathcal{O}} \mathbf{o}_{t_0}, \dots, g \triangleright_{\mathcal{O}} \mathbf{o}_t) \in \mathcal{O}^H$, and $t_0 = \max(0, t - H + 1)$.

Justification: *In setups involving learning from human demonstrations and in most offline learning problems, one cannot expect the expert policy inducing the training dataset \mathbb{D}_H^{exp} to be an optimal control policy. Indeed, \mathbb{D}_H^{exp} is commonly composed of a collection of optimal and suboptimal trajectories from one or multiple expert demonstrators (Ziebart et al., 2008), and demonstrations can be biased—for instance, in the control task of writing with a humanoid robot, a dataset of expert demonstrations can be biased toward right-handed writing demonstrations. However, all expert demonstrations define trajectories solving the target control task, which, by the symmetries of the control problem (section 8.2), implies that symmetry-transformed demonstrations define trajectories solving the symmetry-transformed control task. For instance, any demonstration of writing with the right hand can be transformed into a demonstration of writing with the left hand, such that the robot can learn an ambidextrous writing policy from a dataset of right-handed writing demonstrations. Hence, if one is interested in solving all symmetry-related control tasks—i.e., ambidextrous locomanipulation—it is natural and beneficial to assume the \mathbb{G} -invariance of the expert policy π_{exp} .*

In practice, assumption 8.1 implies that for every expert demonstration trajectory in \mathbb{D}_H^{exp} , the corresponding symmetry-transformed trajectory is equally likely to occur under the joint distribution over observation histories and actions $\mathbb{P}_{\mathbf{o}^H \mathbf{a}}^{exp}$ (see the example in fig. 8.3). Formally, by prop. 8.1 and assumption 6.1, this implies that $\mathbb{P}_{\mathbf{o}^H \mathbf{a}}^{exp}$ is a \mathbb{G} -invariant joint distribution satisfying:

$$\mathbb{P}_{\mathbf{o}^H \mathbf{a}}^{exp}(\vec{\sigma}, \mathbf{a}) = \mathbb{P}_{\mathbf{o}^H \mathbf{a}}^{exp}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}, g \triangleright_{\mathcal{A}} \mathbf{a}) \quad \forall g \in \mathbb{G}, \vec{\sigma} \in \mathcal{O}^H, \mathbf{a} \in \mathcal{A}. \quad (8.13)$$

This has a direct practical interpretation: if one aims to solve a control task and its symmetry-related variants—as in locomotion (figs. 1.1 and 8.9), dexterous manipulation (fig. 8.12), and bimanual mobile manipulation (figs. 8.1 and 8.6)—then each expert trajectory in the training set also serves as a demonstration for the corresponding symmetry-transformed task. Hence, symmetry exploitation in imitation learning yields policies that solve multiple symmetry-related tasks, even when the dataset contains demonstrations of only one task variant.

For robots with bilateral reflection symmetry (e.g., humanoids, quadrupeds, and mobile manipulators), solving a task and its reflected version corresponds to an ambidextrous locomanipulation policy that solves both a task and its mirror variant—e.g., recovering from a fall and its mirrored counterpart, manipulating an object with the left/right hand, or turning left and right. For systems with larger symmetry groups, such as the Mini-Cheetah (see fig. 1.3), the same argument implies that, given expert locomotion demonstrations, the learned policy can be guaranteed to generalise to all symmetry-related locomotion tasks, including reflected gaits, forward and backward gaits, and gaits with the belly facing up and down.

As detailed next, there are two complementary approaches to leverage symmetry priors in data-driven optimal control: data augmentation and symmetry-constrained optimisation.

Data augmentation Eq. 8.13 motivates data augmentation: exploit known symmetries of the control problem to enlarge the expert dataset and obtain a symmetry-augmented dataset \mathbb{D}_H^{aug} with improved




Figure 8.3 Left: Push-T reflection symmetry with an expert rollout from s_0 and its reflected rollout from $g_r \triangleright_S s_0$. The first-action distribution $\mathbf{a}_0 \sim \pi_{\text{exp}}(\cdot | s_0)$ (orange, left) maps to $g_r \triangleright_A \mathbf{a}_0 \sim \pi_{\text{exp}}(\cdot | g_r \triangleright_S s_0)$ (orange, right) under [assumption 8.1](#); hence each demonstration in $\mathbb{D}_H^{\text{exp}}$ is also a demonstration for the reflected task. Right: \mathbb{G} -invariance of the joint distribution $\mathbb{P}_{\sigma^H \mathbf{a}}^{\text{exp}}$ induced by π_{exp} , motivating symmetry-based data augmentation to obtain $\mathbb{D}_H^{\text{aug}}$ (eq. 8.14) with improved empirical coverage of $\mathbb{P}_{\sigma^H \mathbf{a}}^{\text{exp}}$. Figure adapted from [Siebenborn et al. \(2026\)](#).

empirical coverage of $\mathbb{P}_{\sigma^H \mathbf{a}}^{\text{exp}}$ (see [fig. 8.3](#)), defined as:

$$\mathbb{D}_H^{\text{aug}} := \{(g \triangleright_{\mathcal{O}^H} \vec{\sigma}, g \triangleright_A \mathbf{a}) \mid \forall (\vec{\sigma}, \mathbf{a}) \in \mathbb{D}_H^{\text{exp}}, \forall g \in \mathbb{G}\}, \quad \text{with } |\mathbb{D}_H^{\text{aug}}| = |\mathbb{G}| \cdot |\mathbb{D}_H^{\text{exp}}|. \quad (8.14)$$

Symmetry-constrained behaviour cloning Both expert-policy \mathbb{G} -invariance ([assumption 8.1](#)) and optimal-policy \mathbb{G} -invariance ([corollary 8.2](#)) motivate casting [eq. 8.12](#) as a constrained optimisation problem in which \mathbb{G} -invariance of π_θ is explicitly enforced:

$$\theta^* = \arg \max_{\theta} \mathcal{L}_{\text{bc}}(\pi_\theta) = \arg \min_{\theta} \mathbb{E}_{\vec{\sigma} \sim \mathbb{P}_{\sigma^H \mathbf{a}}^{\text{exp}}} D_{\text{KL}}(\pi_{\text{exp}}(\cdot | \vec{\sigma}) \parallel \pi_\theta(\cdot | \vec{\sigma})), \quad (8.15a)$$

$$\text{s.t. } \pi_\theta(\mathbf{a} | \vec{\sigma}) = \pi_\theta(g \triangleright_A \mathbf{a} | g \triangleright_{\mathcal{O}^H} \vec{\sigma}), \quad \forall g \in \mathbb{G}, \mathbf{a} \in \mathcal{A}, \vec{\sigma} \in \mathcal{O}^H. \quad (8.15b)$$

Next, we address [eq. 8.15](#) using the flow-matching paradigm ([Lipman et al., 2022](#)) to parameterise π_θ .

Imitation Learning via Equivariant Flow Matching

To tackle [eq. 8.15](#), we adopt flow matching ([Lipman et al., 2022](#)), a continuous-time generative modelling framework for approximate sampling from a target distribution—here, the family of expert action distributions $\pi_{\text{exp}}(\cdot | \vec{\sigma})$ for all $\vec{\sigma} \in \mathcal{O}^H$.

For each conditioning $\vec{\sigma} \in \mathcal{O}^H$, flow matching learns a conditional probability path $\{\bar{\pi}_k : \Sigma_{\mathcal{A}} \times \mathcal{O}^H \mapsto [0, 1]\}_{k \in [0, 1]}$ that interpolates between a tractable base distribution, $\bar{\pi}_{k=0}(\cdot | \vec{\sigma})$ —e.g., a multivariate Gaussian—and the expert distribution, $\bar{\pi}_{k=1}(\cdot | \vec{\sigma}) := \pi_{\text{exp}}(\cdot | \vec{\sigma})$. Here $k \in [0, 1]$ is the *transport time*. The path is governed by a conditional optimal velocity field $\mathbf{v}_* : \mathcal{A} \times \mathcal{O}^H \times [0, 1] \rightarrow \mathcal{A}$, which defines directional changes of probability mass for each conditioning and transport time (see [fig. 8.4](#)). Sampling from π_{exp} is then modelled through solutions of the ODE induced by \mathbf{v}_* :

$$\mathbf{a}^{(1)} = \mathbf{a}^{(0)} + \int_0^1 \mathbf{v}_*(\mathbf{a}^{(k)}, \vec{\sigma}_t, k) dk, \quad (8.16)$$

with $\mathbf{a}^{(0)} \sim \bar{\pi}_0(\cdot | \vec{\sigma})$, and $\mathbf{a}_t := \mathbf{a}^{(1)} \sim \bar{\pi}_1(\cdot | \vec{\sigma}_t) := \pi_{\text{exp}}(\cdot | \vec{\sigma}_t)$.

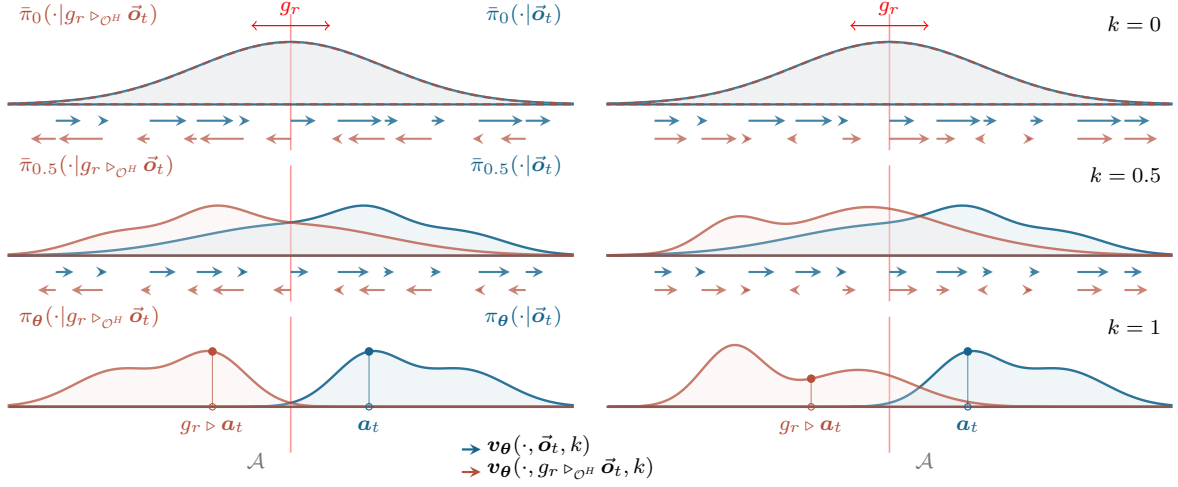


Figure 8.4 Flow-matching probability transport paths at $k \in [0, 0.5, 1]$. **Left:** Path induced by a \mathbb{C}_2 -equivariant velocity field \mathbf{v}_θ , satisfying $g_r \triangleright_{\mathcal{A}} \mathbf{v}_\theta(\mathbf{a}, \vec{\sigma}, k) = \mathbf{v}_\theta(g_r \triangleright_{\mathcal{A}} \mathbf{a}, g_r \triangleright_{\mathcal{O}^H} \vec{\sigma}, k)$ for all $\mathbf{a} \in \mathcal{A}$, $\vec{\sigma} \in \mathcal{O}^H$, and $k \in [0, 1]$. With $\mathbb{C}_2 = \{e, g_r \mid g_r^2 = e\}$ and $g_r \triangleright_{\mathcal{A}} \mathbf{a} := -\mathbf{a}$, the induced policy at $k = 1$ satisfies eq. 8.15b. **Right:** Path induced by a non-equivariant velocity field, yielding a policy π_θ that violates eq. 8.15b. Figure adapted from Ordoñez-Apraez et al. (2025).

Here, $\mathbf{a}^{(k)}$ denotes the *intermediate* action sample along the flow-matching ODE trajectory, and t denotes the *control-time* index of the underlying control problem. For each fixed conditioning $\vec{\sigma}_t$, the action \mathbf{a}_t is generated at endpoint $k = 1$, i.e., $\mathbf{a}_t := \mathbf{a}^{(1)}$ (see fig. 8.4). The objective is therefore to approximate \mathbf{v}_* with a parametric field \mathbf{v}_θ and, via eq. 8.16, parameterise a policy π_θ that approximates sampling from π_{exp} for every conditioning $\vec{\sigma} \in \mathcal{O}^H$.

Equivariant flow-matching Under assumption 8.1, π_{exp} is a \mathbb{G} -invariant conditional distribution (see fig. 8.3 and definition 2.17). This yields a key inductive bias: for the ODE in eq. 8.16 to produce $\bar{\pi}_{k=1} := \pi_{\text{exp}}$ with \mathbb{G} -invariance, \mathbf{v}_* must be equivariant (see corollary 2.2 and fig. 8.4), i.e.:

$$g \triangleright_{\mathcal{A}} \mathbf{v}_*(\mathbf{a}, \vec{\sigma}, k) = \mathbf{v}_*(g \triangleright_{\mathcal{A}} \mathbf{a}, g \triangleright_{\mathcal{O}^H} \vec{\sigma}, k), \quad \forall g \in \mathbb{G}, \mathbf{a} \in \mathcal{A}, \vec{\sigma} \in \mathcal{O}^H, k \in [0, 1]. \quad (8.17)$$

This motivates formulating the approximation of \mathbf{v}_* as a constrained regression problem of the form:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{FM}}(\theta) := \mathbb{E}_{\substack{\mathbf{a}^{(k)} \sim \bar{\pi}_k(\cdot | \vec{\sigma}) \\ \vec{\sigma} \sim \mathbb{D}_H^{\text{exp}}, k \sim U[0,1]}} \left\| \mathbf{v}_\theta(\mathbf{a}^{(k)}, \vec{\sigma}, k) - \mathbf{v}_*(\mathbf{a}^{(k)}, \vec{\sigma}, k) \right\|_2^2, \quad (8.18a)$$

$$\text{s.t.} \quad g \triangleright_{\mathcal{A}} \mathbf{v}_\theta(\mathbf{a}, \vec{\sigma}, k) = \mathbf{v}_\theta(g \triangleright_{\mathcal{A}} \mathbf{a}, g \triangleright_{\mathcal{O}^H} \vec{\sigma}, k), \quad \forall g \in \mathbb{G}, \mathbf{a} \in \mathcal{A}, \vec{\sigma} \in \mathcal{O}^H, k \in [0, 1]. \quad (8.18b)$$

Crucially, minimizing velocity-field error effectively minimizes the behaviour-cloning objective in eq. 8.12 for every conditioning $\vec{\sigma} \in \mathcal{O}^H$ (Su et al., 2025). Thus, flow matching provides a principled and scalable generative framework for imitation learning. Moreover, enforcing equivariance of \mathbf{v}_θ ensures that the resulting policy π_θ satisfies the invariance constraint in eq. 8.15b (see fig. 8.4 and section 2.4 and Klein et al. (2023)), as well as the optimal-policy \mathbb{G} -invariance constraint in corollary 8.2.

Note that, since \mathbf{v}_* is generally intractable, in practice, eq. 8.18a is minimized via an unbiased surrogate loss, termed the conditional flow-matching loss \mathcal{L}_{CFM} (see details in Lipman et al. (2022)).

Within this formulation, symmetry priors enter through two complementary mechanisms: (i) data augmentation (eq. 8.14 and fig. 8.3), and (ii) constrained optimisation that enforces \mathbb{G} -invariance of the learned policy (eqs. 8.15b and 8.18b), either via explicit constraints or regularisation:

Equivariance constraint regularisation One approach to approximately enforce the equivariance constraint in eq. 8.15b is to penalise constraint violations through regularisation, yielding:

$$\begin{aligned} & \arg \min_{\theta} \mathcal{L}_{\text{FM}}(\theta) + \lambda \mathcal{R}_{\text{eq}}(\theta) \\ & \text{with } \mathcal{R}_{\text{eq}}(\theta) := \sum_{g \in \mathbb{G}} \|g \triangleright_{\mathcal{A}} \mathbf{v}_{\theta}(\mathbf{a}, \vec{\sigma}, k) - \mathbf{v}_{\theta}(g \triangleright_{\mathcal{A}} \mathbf{a}, g \triangleright_{\mathcal{O}^H} \vec{\sigma}, k)\|^2, \forall \mathbf{a} \in \mathcal{A}, \vec{\sigma} \in \mathcal{O}^H, k \in [0, 1]. \end{aligned} \quad (8.19)$$

Here, \mathcal{L}_{FM} denotes the flow-matching loss in eq. 8.18a, and $\lambda > 0$ is a Lagrange multiplier that controls the strength of the quadratic penalty on constraint violations, $\mathcal{R}_{\text{eq}}(\theta)$.

The main advantage of this approach is simplicity: it requires no modification to the NN architecture used to parameterize \mathbf{v}_{θ} . Its main limitation is that equivariance is enforced only approximately on the support of the augmented training distribution (see fig. 8.3), with no guarantees for out-of-distribution inputs. From a control-theoretic perspective, violations of corollary 8.2 render the learned policy suboptimal by construction, since performance may differ across symmetry-related initial states (e.g., the two initial states in fig. 8.3-left). In addition, approximating eq. 8.19 from a batch of size B requires augmentation by a factor $|\mathbb{G}|$, with memory cost $\mathcal{O}(B \cdot |\mathbb{G}|)$. This can be prohibitive for large symmetry groups and impossible for continuous compact groups¹, thus requiring approximation techniques for the regularisation term, such as random sampling of group elements for each batch.

Explicit equivariance constraint Instead of penalizing violations, we can explicitly enforce eq. 8.18b by adopting an equivariant NN architecture whose trainable parameters satisfy equivariance by construction, including out-of-distribution inputs Bronstein et al. (2021); Ordoñez-Apaez et al. (2025); Weiler et al. (2021). This enables direct minimization of eq. 8.18 with guaranteed constraint satisfaction.

Imposing explicit equivariance requires modifying the NN used for \mathbf{v}_{θ} so that each submodule is equivariant to \mathbb{G} . Although several libraries support equivariant NNs (Cesa et al., 2021; Geiger et al., 2022; Kondor et al., 2018; Thomas et al., 2018; Weiler et al., 2021), most target point-cloud or image data and are not directly applicable to trajectories of symmetric vector-valued signals. As a result, adapting state-of-the-art (s.o.t.a.) imitation-learning architectures—such as the conditional encoder-decoder transformer or Unet1D models in Chi et al. (2023, 2024)—remains non-trivial and often requires substantial implementation effort and representation-theory expertise (Ordoñez-Apaez et al., 2023; Weiler et al., 2021).

To address these challenges and lower the barrier to exploiting state/morphological symmetries in s.o.t.a. imitation learning and data-driven control, this thesis introduces `symm_learning`, a symmetry-aware ML library that provides \mathbb{G} -equivariant baselines of the encoder-decoder transformer (see fig. 8.5) and Unet1D architectures used in (Chi et al., 2023, 2024), for any finite group \mathbb{G} (see section 1.5.1).

¹For continuous compact groups, the regularisation in eq. 8.19 is defined by an integral over the group weighted by the Haar measure.

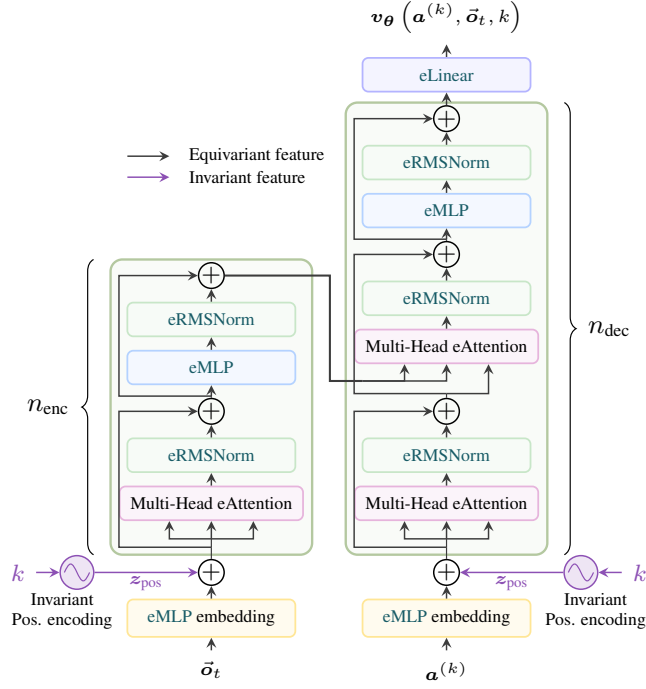


Figure 8.5 \mathbb{G} -equivariant transformer (encoder-decoder) architecture for parameterizing the flow-matching velocity field v_θ with built-in \mathbb{G} -equivariance (eq. 8.17). The encoder applies n_{enc} equivariant transformer layers to $\bar{o}_t + z_{\text{pos}}(k)$. The decoder applies n_{dec} equivariant transformer layers with cross-attention to the encoder latent, taking $a^{(k)} + z_{\text{pos}}$ as input and outputting the velocity used in eq. 8.16. Source code and documentation: `symm_learning.models.eCondTransformerRegressor`. Figure adapted from Siebenborn et al. (2026).

Experimental results

Our experimental setup compares three strategies for leveraging symmetry priors in imitation learning when optimizing eq. 8.15: (i) data augmentation (*SymAug*, eq. 8.14), (ii) soft policy-invariance enforcement via loss regularization (*EquivReg*, eq. 8.19), and (iii) explicit invariance enforcement through equivariant NNs (*EquivNet*). Tasks include synthetic Push-T (fig. 8.3) and simulated/real-world mobile-manipulation benchmarks (fig. 8.6). Results report 3 training seeds and 50 evaluation environments per seed. Full experimental details are in Siebenborn et al. (2026).

All symmetry-aware methods are compared against a symmetry-agnostic flow-matching baseline policy (*BL*) with velocity field v_θ parameterized by an encoder-decoder transformer adapted from (Chi et al., 2023, 2024). The same backbone is used for *SymAug* and *EquivReg*, while *EquivNet* uses our \mathbb{C}_2 -equivariant transformer (*eTransformer*) with an equivalent architectural footprint (see fig. 8.5). Thus, experiments control for backbone size. Equivariance constraints in the *eTransformer* reduce trainable parameters by roughly half, trading expressivity for global satisfaction of the policy-symmetry constraint motivated by corollary 8.2.

Push-T Task The first experiment examines reflection-symmetry exploitation on the state-based planar Push-T task from Chi et al. (2023), a standard imitation-learning benchmark. As shown in fig. 8.3, the task has state reflection symmetry group \mathbb{C}_2 , also present as a morphological symmetry in bimanual mobile manipulation. Hence, it is modelled as a \mathbb{C}_2 -symmetric POMDP satisfying the priors in section 8.2, making it a suitable setting to evaluate symmetry exploitation.




Figure 8.6 Benchmark environments used in the mobile-manipulation experiments (section 8.3.1). The four subplots (three simulation tasks and one real-world task) show the original scene and its reflection-transformed counterpart under $g_r \in \mathbb{C}_2$, illustrating the symmetry-related task pairs solved by a single policy. Figure adapted from Siebenborn et al. (2026).

Figure 8.7-left compares sample efficiency of all policy variants, reporting target T-pose reaching performance versus the number of expert rollouts. Across all dataset-size regimes, symmetry-informed policies consistently outperform the baseline, indicating improved generalization and optimality in low-, mid-, and high-data regimes; *EquivNet* performs best in low data. On average, symmetry exploitation yields a $2\times$ (or larger) sample-efficiency gain: the baseline needs 200 rollouts to reach a 0.9 score, whereas symmetry-informed policies reach 0.95 with 100 rollouts and near-perfect performance with 200. These results highlight the benefits of reflection-symmetry exploitation even in low-dimensional control tasks.

Simulated bimanual mobile manipulation This experiment studies symmetry exploitation in three bimanual mobile-manipulation tasks: *bimanual reaching*, *box lifting*, and *cabinet opening* (see fig. 8.6). The bimanual manipulator *Tiago++* has bilateral reflection symmetry group \mathbb{C}_2 , which induces reflection symmetry in these tasks (see figs. 8.1 and 8.6). Because the goal is to solve each task and its reflection-transformed version with a single policy, we assume expert-policy \mathbb{G} -invariance (assumption 8.1), yielding the symmetry-constrained formulation in eq. 8.15.




Figure 8.7 Control performance as a function of the number of expert demonstrations. Left: Push-T. Middle and right: simulated bimanual mobile manipulation (box lifting and cabinet opening). Across tasks and data regimes, symmetry-informed policies (*SymAug*, *EquivReg*, *EquivNet*) outperform the symmetry-agnostic baseline (*BL*), with *EquivReg* and *EquivNet* generally strongest. On average, symmetry-informed policies yield about a $2\times$ sample-efficiency gain over *BL*.

| Method | Box lifting (Zero-shot, success rate \uparrow) | | | Cabinet opening (Zero-shot, success rate \uparrow) | | |
|-------------|---|-----------------------------------|-----------------|---|-----------------------------------|-----------------|
| | e | g_r | total | e | g_r | total |
| FM-baseline | 0.59 \pm 0.14 | 0.00 \pm 0.00 | 0.29 \pm 0.07 | 0.89 \pm 0.03 | 0.00 \pm 0.00 | 0.44 \pm 0.02 |
| FM-SymAug | 0.63 \pm 0.07 | 0.67 \pm 0.06 | 0.65 \pm 0.04 | 0.90 \pm 0.07 | 0.88 \pm 0.02 | 0.89 \pm 0.04 |
| FM-EquivReg | 0.70 \pm 0.12 | 0.72 \pm 0.03 | 0.71 \pm 0.06 | 0.85 \pm 0.08 | 0.87 \pm 0.02 | 0.86 \pm 0.04 |
| FM-EquivNet | 0.67 \pm 0.11 | 0.66 \pm 0.16 | 0.67 \pm 0.13 | 0.87 \pm 0.02 | 0.83 \pm 0.03 | 0.85 \pm 0.02 |

Table 8.1 *Zero-shot* evaluation in simulated mobile manipulation. Policies are trained only on demonstrations from the original configuration (e) and evaluated on both the original (e) and reflected (g_r) configurations. Results report success rate (mean \pm standard deviation; higher is better) for box lifting and cabinet opening.

Expert demonstration trajectories are collected in a `mujoco` simulation (Todorov et al., 2012) using scripted end-effector policies and a whole-body tracking controller based on `mink` (Zakka, 2025). The learned policy must capture the distribution of approaching, grasping, and manipulation motions from these demonstrations. Control actions are $\mathbf{a} = [\mathbf{X}_{\text{tgt},L}, \mathbf{X}_{\text{tgt},R}, \mathbf{c}_{\text{tgt}}]$, where $\mathbf{X}_{\text{tgt},L}, \mathbf{X}_{\text{tgt},R} \in \mathbb{SE}_3$ are target left and right end-effector configurations tracked by the whole-body controller, and $\mathbf{c}_{\text{tgt}} = [b_{\text{tgt},L}, b_{\text{tgt},R}] \in \mathbb{R}^2$ is the target gripper-closing state. Likewise, each observation is $\mathbf{o} = [\mathbf{X}_{\text{eef},L}, \mathbf{X}_{\text{eef},R}, \mathbf{c}, \mathbf{e}_{\text{goal},L}, \mathbf{e}_{\text{goal},R}]$, comprising left and right end-effector configurations $\mathbf{X}_{\text{eef},L}, \mathbf{X}_{\text{eef},R} \in \mathbb{SE}_3$, binary gripper-closing state $\mathbf{c} = [b_L, b_R] \in \mathbb{R}^2$, and target object poses plus pose errors for the two arms, i.e., $\mathbf{e}_A = [\mathbf{X}_{\text{goal},A}, \mathbf{X}_{\text{goal},A} \ominus \mathbf{X}_{\text{eef},A}]$ for $A \in \{L, R\}$. Here, $\mathbf{X}_{\text{goal},\cdot} \ominus \mathbf{X}_{\text{eef},\cdot}$ denotes the end-effector configuration error expressed as a vector in the Lie algebra of \mathbb{SE}_3 (Selig, 2004).

The experiment quantifies three advantages of symmetry-informed policies over the symmetry-agnostic baseline: (i) zero-shot generalization to reflected task states unseen during training, (ii) sample efficiency versus number of expert demonstrations, and (iii) spatial generalization under broader initial/target-state distributions.

1. **Symmetric zero-shot generalization.** Policies are trained only on the original configuration $e \in \mathcal{C}_2$ and evaluated on both e and reflected $g_r \in \mathcal{C}_2$ (fig. 8.6). In *box lifting*, this means training on right-table lifting and evaluating also on left-table lifting. Results in section 8.3.1 show that all symmetry-informed methods preserve baseline performance on e and generalize to unseen g_r , with broadly comparable zero-shot performance. This supports ambidextrous behavior without degrading in-distribution performance.
2. **Sample efficiency.** Policies are trained with demonstrations from both e and g_r variants of each task, and evaluated on both. As shown in fig. 8.7-middle for *box lifting* and *cabinet opening*, all symmetry-informed methods outperform the baseline across low-, middle-, and high-data regimes, with at least $2\times$ sample-efficiency gains. Methods that enforce the policy-invariance optimality constraint (*EquivReg* and *EquivNet*) perform best across regimes, motivating constrained optimization beyond data augmentation alone.
3. **Spatial generalization.** Policies are trained on *bimanual reaching* and *cabinet opening* under two difficulty levels (*narrow*, *wide*), corresponding to broader target-pose distributions and larger cabinet-pose variation. Figure 8.8 reports relative performance gains over the baseline, with all methods trained on 25 demonstrations. Symmetry-informed methods outperform the baseline at both difficulty levels. Gains are comparable across methods in the *narrow* setting, while *EquivNet* is most robust in the *wide* setting, where out-of-distribution evaluations are more likely. We attribute this to global satisfaction of the policy-invariance optimality constraint (corollary 8.2),

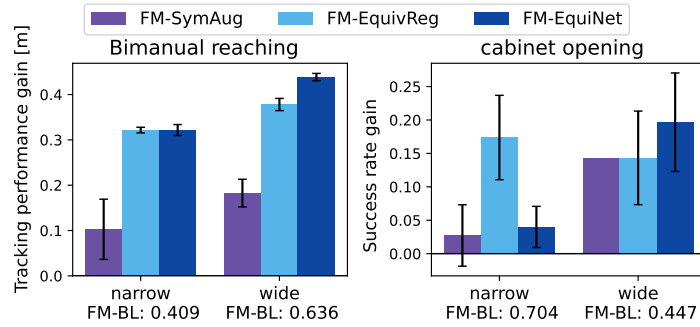


Figure 8.8 Relative gain over the symmetry-agnostic baseline in spatial-generalization experiments for *narrow* and *wide* settings of bimanual reaching and cabinet opening (all methods trained with 25 demonstrations). Symmetry-informed methods improve both tracking error and success rate in both settings. Methods enforcing the policy-invariance optimality constraint (corollary 8.2) provide the largest gains in the *wide* regime, where out-of-distribution evaluations are more common. Figure adapted from Siebenborn et al. (2026).

which *EquivNet* guarantees by construction through equivariance constraints in the NN architecture parameterizing the flow-matching velocity field.

Summary 8.3.1: Imitation learning

In imitation learning for locomanipulation control tasks involving robotic systems with a state or morphological symmetry group \mathbb{G} , if the policy is expected to generalize across symmetry-related tasks, the expert policy can be assumed \mathbb{G} -invariant (assumption 8.1):

$$\pi_{\text{exp}}(\mathbf{a}|\vec{\sigma}) = \pi_{\text{exp}}(g \triangleright_{\mathcal{A}} \mathbf{a} | g \triangleright_{\mathcal{O}^H} \vec{\sigma}), \quad \forall g \in \mathbb{G}, \mathbf{a} \in \mathcal{A}, \vec{\sigma} \in \mathcal{O}^H.$$

In imitation learning via flow-matching, expert-policy invariance makes the learning problem equivariant (eq. 8.18a). This motivates both data augmentation (eq. 8.14) to improve empirical coverage of expert demonstrations, and constrained optimization to enforce \mathbb{G} -invariance of the learned policy—either through regularization (eq. 8.19) or explicit equivariance constraints in the NN architecture parameterizing the flow-matching velocity field (eq. 8.17).

For bimanual mobile manipulation with bilateral reflection symmetry, leveraging these symmetry priors consistently yields ambidextrous control policies with at least $2\times$ improvements in sample efficiency and enhanced robustness in out-of-distribution settings.

8.3.2 Model-free reinforcement learning

In the last decade, model-free RL has become a primary approach for contact-rich robotic control problems, such as locomotion and locomanipulation, especially when no extensive expert-demonstration dataset is available (Grigorescu et al., 2020; Kumar et al., 2016; Rudin et al., 2022; Team et al., 2025). These methods are effective in regimes where model-based trajectory optimization struggles with combinatorial contact-mode planning (Mastalli et al., 2020a,b) and hard-to-model contact dynamics Todorov et al. (2012).

A key limitation is poor sample efficiency, which increases data and compute requirements and often yields brittle or suboptimal locomanipulation policies. This section shows how the symmetry

priors from sections 8.1 and 8.2 provide inductive biases that improve sample efficiency and control performance.

We focus on PPO algorithms (Schulman et al., 2017), but the same foundational priors apply to any model-free on-/off-policy RL method, in the sense that the \mathbb{G} -invariance of the MDP optimal value and quality functions (theorem 8.1 and corollary 8.1) together with the induced symmetry constraints on optimal policy parametrizations (corollary 8.2).

Proximal policy optimization algorithms PPO is a family of policy-gradient actor-critic algorithms that learn a parametric stochastic control policy $\pi_\theta : \Sigma_{\mathcal{A}} \times \mathcal{O}^H \rightarrow [0, 1]$ and a parametric value function $V_\theta : \mathcal{O}^H \rightarrow \mathbb{R}$, approximating the optimal policy and value function of the truncated-history belief MDP modelling contact-rich control tasks. In this setup, the belief-MDP state space \mathcal{O}^H is the space of past observation-history trajectories of length $H \in \mathbb{N}_+$ (see section 8.2).

These algorithms alternate between sampling experience interactions—with a fixed policy $\pi_{\theta_{\text{old}}}$ —and optimizing π_θ using the following advantage-weighted importance-sampling estimator:

$$\mathcal{P}_{\text{PG}}(\vec{o}_t, \mathbf{a}_t; \theta) := \frac{\pi_\theta(d\mathbf{a}_t | \vec{o}_t)}{\pi_{\theta_{\text{old}}}(d\mathbf{a}_t | \vec{o}_t)} \widehat{A}_\theta(\vec{o}_t, \mathbf{a}_t). \quad (8.20)$$

Here, $A_{\pi_{\theta_{\text{old}}}} : \mathcal{O}^H \times \mathcal{A} \rightarrow \mathbb{R}$ denotes the intractable *advantage* function, defined as $A_{\pi_{\theta_{\text{old}}}}(\vec{o}_t, \mathbf{a}_t) := \widehat{Q}_{\pi_{\theta_{\text{old}}}}(\vec{o}_t, \mathbf{a}_t) - \widehat{V}_{\pi_{\theta_{\text{old}}}}(\vec{o}_t)$ (see eqs. 8.10 and 8.11). In practice, it is estimated from the last collected batch and the parametric value function as:

$$\begin{aligned} \widehat{A}_\theta(\vec{o}_t, \mathbf{a}_t) &:= \widehat{Q}_{\pi_{\theta_{\text{old}}}}(\vec{o}_t, \mathbf{a}_t, P) - V_\theta(\vec{o}_t) \\ \text{with } \widehat{Q}_{\pi_{\theta_{\text{old}}}}(\vec{o}_t, \mathbf{a}_t, P) &:= \gamma^P V_\theta(\vec{o}_{t+P\Delta t}) + \sum_{k=0}^{P-1} \gamma^k \widehat{r}(\vec{o}_{t+k\Delta t}, \mathbf{a}_{t+k\Delta t}) \end{aligned} \quad (8.21)$$

$\widehat{Q}_{\pi_{\theta_{\text{old}}}}$ denotes an approximate quality function based on a fixed look-ahead horizon P and the bootstrap estimate V_θ . Using the same batch, we define the pointwise value-regression error as

$$\mathcal{E}_{\text{VR}}(\vec{o}_t; \theta) := (V_{\theta_{\text{old}}}(\vec{o}_t) - V_\theta(\vec{o}_t))^2, \quad (8.22)$$

In this context, the training objective for any PPO algorithm can be written with a single expectation over the sampling distribution as:

$$\mathcal{L}_{\text{PPO}}(\theta) = \mathbb{E}_{\pi_{\theta_{\text{old}}}} [\mathcal{P}_{\text{RPG}}(\vec{o}_t, \mathbf{a}_t; \theta) - \lambda_V \mathcal{E}_{\text{VR}}(\vec{o}_t; \theta) + \lambda_E \mathcal{R}(\pi_\theta(\cdot | \vec{o}_t))], \quad (8.23)$$

The expectation $\mathbb{E}_{\pi_{\theta_{\text{old}}}}$ is taken over the distribution of observation histories and actions induced by $\pi_{\theta_{\text{old}}}$, which is fixed while updating π_θ in each optimization phase and updated at the start of the next sampling iteration. Moreover, \mathcal{P}_{RPG} denotes a regularized pointwise version of \mathcal{P}_{PG} in eq. 8.20 that penalizes large policy updates (the most common variant is the CLIP surrogate; see Schulman et al. (2017)). The scalars $\lambda_V, \lambda_E > 0$ control the value-regression and entropy-regularization terms, respectively.

Symmetry priors in PPO algorithms For contact-rich control tasks of robots featuring a state/morphological symmetry group \mathbb{G} , the \mathbb{G} -invariance of the optimal value and quality functions

(theorem 8.1 and corollary 8.1) together with the corresponding symmetry constraints on optimal policy parametrizations (corollary 8.2) render the PPO objective a symmetry-constrained optimization problem of the form:

$$\begin{aligned} & \arg \min_{\boldsymbol{\theta}} \mathcal{L}_{\text{PPO}}(\boldsymbol{\theta}) \\ \text{s.t.} \quad & \pi_{\boldsymbol{\theta}}(\mathbf{a}|\vec{\boldsymbol{o}}) = \pi_{\boldsymbol{\theta}}(g \triangleright_{\mathcal{A}} \mathbf{a} | g \triangleright_{\mathcal{O}^H} \vec{\boldsymbol{o}}), \quad \forall g \in \mathbb{G}, \mathbf{a} \in \mathcal{A}, \vec{\boldsymbol{o}} \in \mathcal{O}^H. \\ & V_{\boldsymbol{\theta}}(\vec{\boldsymbol{o}}) = V_{\boldsymbol{\theta}}(g \triangleright_{\mathcal{O}^H} \vec{\boldsymbol{o}}), \end{aligned} \quad (8.24)$$

Similarly to imitation learning (section 8.3.1), the symmetry priors of this constrained optimization objective can be exploited via (i) data augmentation, and (ii) enforcing policy and value-function invariance constraints through regularization or explicit equivariance constraints in the NN architectures parameterizing $\pi_{\boldsymbol{\theta}}$ and $V_{\boldsymbol{\theta}}$.

Data augmentation For any trajectory sampled during one iteration, $\{\vec{\boldsymbol{o}}_t, \mathbf{a}_t \sim \pi_{\boldsymbol{\theta}_{\text{old}}}(\cdot|\vec{\boldsymbol{o}}_t), \hat{r}(\vec{\boldsymbol{o}}_t, \mathbf{a}_t)\}_{t=0}^T$, the symmetry-transformed trajectory $\{g \triangleright_{\mathcal{O}^H} \vec{\boldsymbol{o}}_t, g \triangleright_{\mathcal{A}} \mathbf{a}_t, \hat{r}(g \triangleright_{\mathcal{O}^H} \vec{\boldsymbol{o}}_t, g \triangleright_{\mathcal{A}} \mathbf{a}_t) := \hat{r}(\vec{\boldsymbol{o}}_t, \mathbf{a}_t)\}_{t=0}^T$ is also valid experience for the symmetry-related task. Augmenting batches with these transformed samples improves empirical coverage and, therefore, estimation of the policy-gradient \mathcal{P}_{PG} , value-regression \mathcal{E}_{VR} , and entropy-regularization \mathcal{R} terms in eq. 8.23, all of which are \mathbb{G} -invariant by construction (Mittal et al., 2024; Mondal et al., 2022; Su et al., 2024; Wang et al., 2022b).

Equivariance/invariance constraint regularization To enforce policy and value-function invariance constraints via regularization, define the following pointwise equivariance/invariance violation penalties:

$$\begin{aligned} \mathcal{R}_{\text{eq}}(\vec{\boldsymbol{o}}_t, \mathbf{a}_t; \boldsymbol{\theta}) &:= \sum_{g \in \mathbb{G}} \|\pi_{\boldsymbol{\theta}}(\mathbf{a}_t|\vec{\boldsymbol{o}}_t) - \pi_{\boldsymbol{\theta}}(g \triangleright_{\mathcal{A}} \mathbf{a}_t | g \triangleright_{\mathcal{O}^H} \vec{\boldsymbol{o}}_t)\|^2, \\ \mathcal{R}_{\text{inv}}(\vec{\boldsymbol{o}}_t; \boldsymbol{\theta}) &:= \sum_{g \in \mathbb{G}} \|V_{\boldsymbol{\theta}}(\vec{\boldsymbol{o}}_t) - V_{\boldsymbol{\theta}}(g \triangleright_{\mathcal{O}^H} \vec{\boldsymbol{o}}_t)\|^2. \end{aligned} \quad (8.25)$$

Then, the regularized PPO objective can be defined with a Lagrange multiplier $\lambda_{\text{equiv}} > 0$ as:

$$\arg \min_{\boldsymbol{\theta}} \mathbb{E}_{\pi_{\boldsymbol{\theta}_{\text{old}}}} [\mathcal{P}_{\text{RPG}}(\vec{\boldsymbol{o}}_t, \mathbf{a}_t; \boldsymbol{\theta}) - \lambda_{\text{V}} \mathcal{E}_{\text{VR}}(\vec{\boldsymbol{o}}_t; \boldsymbol{\theta}) + \lambda_{\text{E}} \mathcal{R}(\pi_{\boldsymbol{\theta}}(\cdot|\vec{\boldsymbol{o}}_t)) + \lambda_{\text{equiv}} (\mathcal{R}_{\text{eq}}(\vec{\boldsymbol{o}}_t, \mathbf{a}_t; \boldsymbol{\theta}) + \mathcal{R}_{\text{inv}}(\vec{\boldsymbol{o}}_t; \boldsymbol{\theta}))] \quad (8.26)$$

Since PPO is already hyperparameter-rich and evaluating eq. 8.26 can be prohibitively expensive for large symmetry groups \mathbb{G} , a practical alternative is to enforce invariance constraints explicitly.

Explicit equivariance/invariance constraints Value-function invariance can be readily enforced by using a \mathbb{G} -invariant NN architecture to parameterize $V_{\boldsymbol{\theta}}$ (Cesa et al., 2021; Ordoñez-Apraes et al., 2025). In contrast, enforcing policy invariance depends on the chosen parametrization of the conditional distribution $\pi_{\boldsymbol{\theta}}(\cdot|\vec{\boldsymbol{o}})$. In most practical uses of PPO, $\pi_{\boldsymbol{\theta}}$ is parameterized as a Gaussian distribution with mean $\mu_{\boldsymbol{\theta}} : \mathcal{O}^H \rightarrow \mathcal{A}$ and covariance map $\Sigma_{\boldsymbol{\theta}} : \mathcal{O}^H \rightarrow \mathbb{R}^{|\mathcal{A}| \times |\mathcal{A}|}$. In this case, satisfying the optimality invariance constraint (corollary 8.2) is achieved by enforcing \mathbb{G} -equivariance

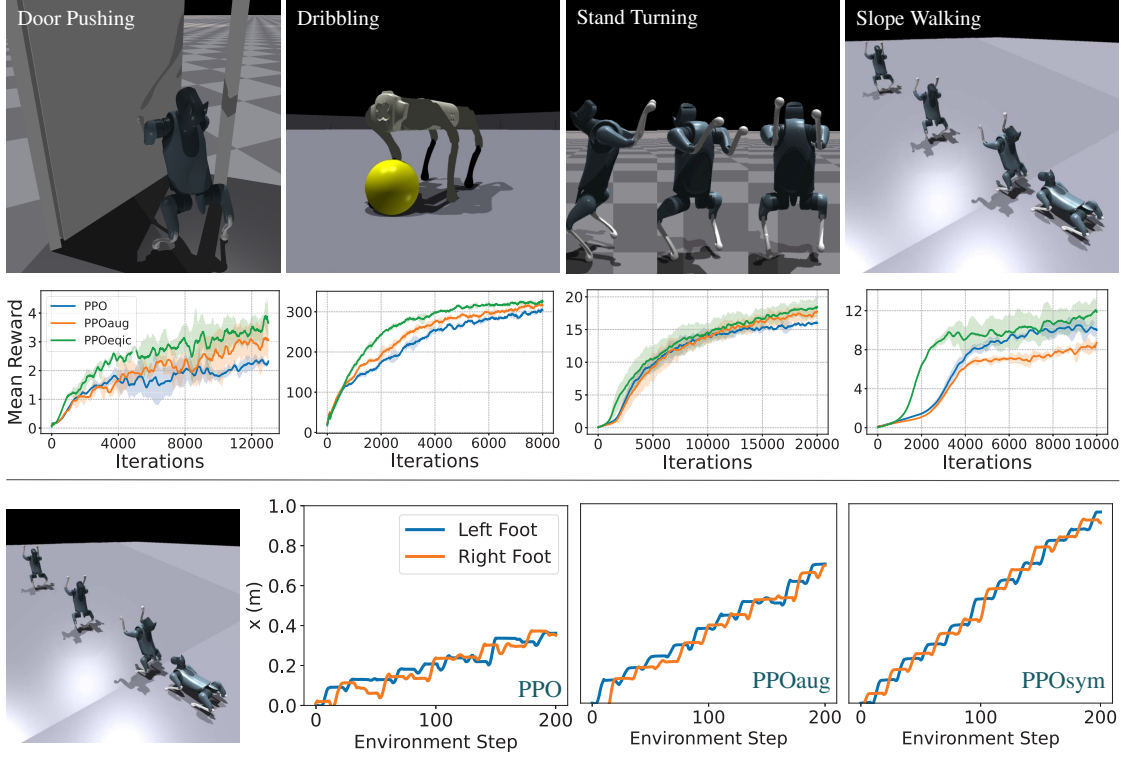


Figure 8.9 Top: snapshots and training curves (average expected discounted reward, i.e., return, vs. experience samples/iterations) for four quadruped locomanipulation tasks under PPO, PPOaug, and PPOsym. Symmetry-aware methods improve sample efficiency and return, with PPOsym strongest overall. Bottom: qualitative slope-walking comparison. The left panel shows the task; the remaining panels show each foot x coordinate over steps, with the x axis aligned to commanded motion. Vanilla PPO shows slower speed tracking, less regular stepping, and transient right-foot dragging; PPOaug improves these behaviors; PPOsym is most regular with best tracking. This highlights that return alone is not a direct proxy for gait quality. Figure adapted from Su et al. (2024).

of these mean and covariance maps, defining the Gaussian distribution for each $\vec{\sigma} \in \mathcal{O}^H$, i.e.:

$$\begin{aligned} \pi_{\theta}(\cdot | \vec{\sigma}) &:= \mathcal{N}(\cdot; \mu_{\theta}(\vec{\sigma}), \Sigma_{\theta}(\vec{\sigma})), & \forall \vec{\sigma} \in \mathcal{O}^H, \\ \text{s.t. } \mu_{\theta}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}) &= g \triangleright_A \mu_{\theta}(\vec{\sigma}), & \Sigma_{\theta}(g \triangleright_{\mathcal{O}^H} \vec{\sigma}) = g \triangleright_A \Sigma_{\theta}(\vec{\sigma}) \triangleleft g^{-1}, & \forall g \in \mathbb{G}, \vec{\sigma} \in \mathcal{O}^H. \end{aligned} \quad (8.27)$$

This can be achieved by using \mathbb{G} -equivariant NN architectures to parameterize μ_{θ} and Σ_{θ} (see `symm_learning.nn.MultivariateNormal` for details and examples).

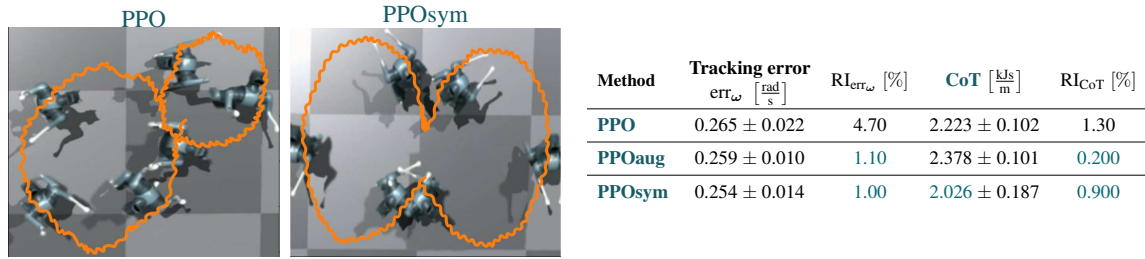


Figure 8.10 Stand-and-turn qualitative/quantitative comparison. Left: turn-left and turn-right trajectories under PPO and PPOsym, with the robot CoM trace in orange. Right: turning-velocity tracking error and cost-of-transport (CoT) (energy proxy). Tracking errors are comparable across methods, while PPOsym attains the lowest CoT and tracking error. We also report reflection-invariance errors that measure left-right asymmetry. For CoT, $\text{RI}_{\text{CoT}} = \frac{|\text{CoT}_{\text{left}} - \text{CoT}_{\text{right}}|}{\text{CoT}_{\text{left}} + \text{CoT}_{\text{right}}} \times 100[\%]$, which measures relative left-right energy asymmetry and thus policy suboptimality under corollary 8.2. Symmetry-agnostic policies show about one-order-of-magnitude larger CoT asymmetry than symmetry-aware policies. Figure adapted from Su et al. (2024).

Quadruped legged locomanipulation

Editorial Note: Section Quadruped legged locomanipulation

The analysis and results from this section are based on the paper Su et al. (2024), done in collaboration with the undergraduate student Zhi Su. Details on my personal contribution to the paper are provided in section 1.5.

This section analyzes symmetry exploitation in model-free RL for contact-rich quadruped locomanipulation. We compare vanilla PPO against two symmetry-aware variants, PPO with data-augmentation (PPOaug) and PPO with hard equivariance / invariance symmetry constraints (PPOsym) (eq. 8.24), using equal or equivalent model footprints and the same number of on-policy interactions. Full experimental details are given in Su et al. (2024).

The experimental suite covers four quadruped locomanipulation tasks: *door pushing*, *ball dribbling*, *stand turning left/right*, and *slope walking* (top row of fig. 8.9). Experiments use Xiaomi CyberDog2 (Xiaomi, 2024) and/or Unitree Go1, both with bilateral reflection symmetry $\mathcal{C}_2 = \{e, g_r | g_r^2 = e\}$. Modelling these tasks as truncated-history belief MDPs (section 8.2) transfers this symmetry to the control problem, so optimality constraints on value, quality, and policy functions apply (theorem 8.1 and corollaries 8.1 and 8.2). In this setting, policy invariance corresponds to ambidextrous behaviour: if a policy rejects a disturbance and solves a task from \vec{o}_t , it should also solve the symmetry-transformed task from $g_r \triangleright_{\mathcal{O}^H} \vec{o}_t$ under the transformed disturbance (section 8.1 and chapter 4).

As shown in fig. 8.9, PPOsym consistently outperforms the other PPO variants in sample efficiency and average expected discounted reward (hereafter, return) across all four tasks. PPOaug improves over vanilla PPO on most tasks but underperforms on *slope walking* in average return; still, fig. 8.9 shows clear gait-quality gains (more regular stepping and better velocity tracking). This highlights that, in RL, return alone is not always an informative proxy for locomotion gait quality, since rewards are often aggregates of multiple terms (e.g., velocity tracking, foot clearance, contact penalizations) tailored to encode a curriculum in early exploration.

To complement return, table 8.2 reports turning-velocity tracking error, CoT (energy proxy), and left-right reflection-invariance errors. Figure 8.11 further shows zero-shot real-world transfer for

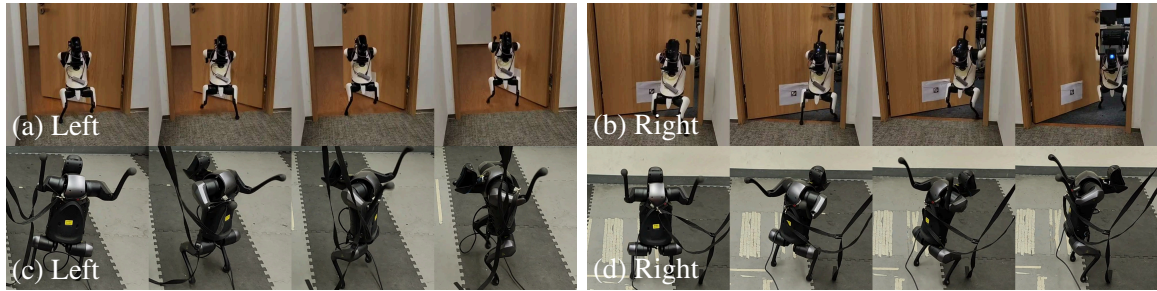


Figure 8.11 Zero-shot real-world deployment snapshots for the locomanipulation tasks of door pushing and stand turning, together with their reflection-transformed counterparts. Symmetry-aware policies [PPOaug](#) and [PPOsym](#), trained on one task variant (left), transfer directly to the reflected variant (right). In real-world deployments, [PPOaug](#) is more robust to disturbances than [PPOsym](#). Figure adapted from [Su et al. \(2024\)](#).

door pushing and stand turning, where symmetry-aware policies transfer better to hardware. Overall, symmetry exploitation improves sample efficiency and control performance, with the strongest results obtained by explicitly enforcing policy and value-function invariance constraints ([corollaries 8.1](#) and [8.2](#)).

Dexterous bimanual manipulation

Editorial Note: Section Dexterous bimanual manipulation

The analysis and results from this section are based on the paper [Li et al. \(2025\)](#), done in collaboration with the PhD student Zechu Li. Details on my personal contribution to the paper are provided in [section 1.5](#). For further details, videos, and animations, visit the project [website](#).

This section studies symmetry exploitation in model-free [RL](#) for dexterous bimanual manipulation. We evaluate six benchmark tasks—*box-lift*, *table-clean*, *drawer-insert*, *threading*, *bowl-stir*, and *handover* ([fig. 8.12](#))—covering cooperative lifting, timed pick-and-place, open–insert–close manipulation, dual-arm insertion with reorientation, stirring, and coordinated handover. Each task is decomposed into two role-specific subtasks, and success is based on sustained goal completion. Policies are trained and evaluated in simulation for all six tasks and deployed in the real world on *box-lift* and *table-clean*. Full task definitions, rewards, and implementation details are provided in [Li et al. \(2025\)](#) and the project website.

Because left- and right-arm in-hand manipulation dynamics are equivalent up to reflection, we model these dexterous bimanual tasks as a truncated-history belief [MDP](#) (see [section 8.2](#)) with bilateral reflection symmetry $\mathbb{C}_2 = \{e, g_r | g_r^2 = e\}$, where g_r reflects the workspace and permutes left/right arm configurations (see [fig. 8.9](#) and [chapter 4](#)). Actions $\mathbf{a}_t \in \mathcal{A}$ are target arm-and-hand joint positions tracked by low-level controllers, and observations $\mathbf{o}_t \in \mathcal{O}$ include the robot state and relative end-effector configuration errors to task-specific goals.

Several tasks involve sequential, role-dependent subtasks (e.g., one hand stabilizes while the other manipulates), so success depends on stage-wise completion. Because vanilla [PPO](#) explores this long-horizon structure inefficiently, we use a [Multi-Task Multi-Agent POMDP \(MTMA-POMDP\)](#) formulation with intermediate rewards and hand-specific policies/tasks ([Li et al., 2025](#)), then distill them into a unified policy for the full task. Since [MTMA-POMDP](#) details and distillation are outside

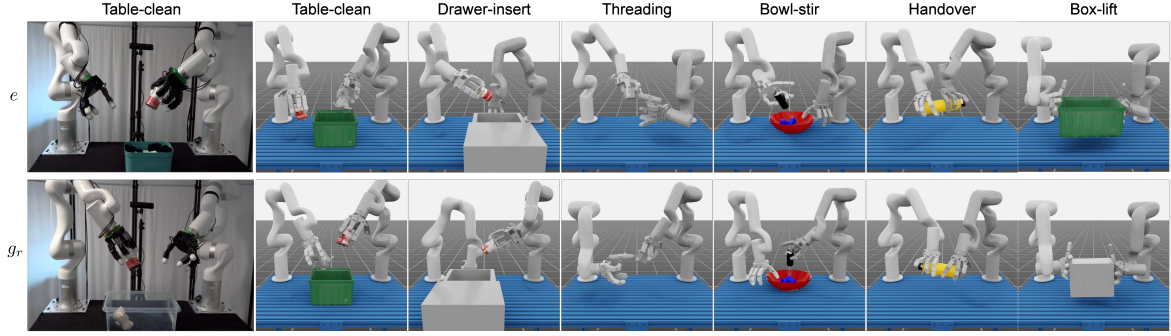


Figure 8.12 Dexterous bimanual manipulation tasks (e row) and their reflected counterparts (g_r row). Since left/right in-hand dynamics are equivalent, each task is symmetric under bilateral reflection $\mathbb{C}_2 = \{e, g_r | g_r^2 = e\}$, where g_r reflects the workspace and permutes left/right hand configurations (see chapter 4). Figure adapted from Li et al. (2025).

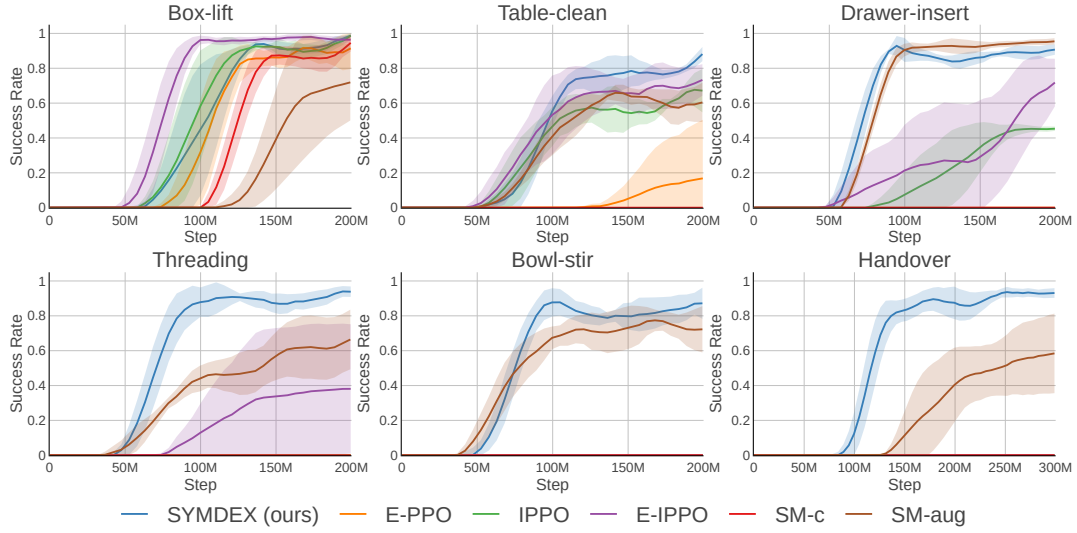


Figure 8.13 Performance of **Symmetric Dexterous Manipulation (SymDex)** and baseline methods on six benchmark tasks. **SymDex** consistently learns all six tasks and achieves success rates exceeding 80%, outperforming all baselines. Figure adapted from Li et al. (2025).

this thesis’s scope, we refer to Li et al. (2025) and focus on how symmetry exploitation—via data augmentation and value/policy invariance constraints (corollaries 8.1 and 8.2)—affects intermediate-policy learning and final-policy distillation.

We evaluate six PPO-based policies to ablate the role of three priors in the **MTMA-POMDP** control problem: subtask and actor/critic decomposition, state/action-space dimensionality, and symmetry exploitation. The policies are:

- * **Equivariant PPO (ePPO)**: a single \mathbb{C}_2 -invariant policy (eq. 8.27) for the full bimanual task (i.e., the 44-DoF action space), trained with policy equivariance constraints (corollary 8.2), leveraging *explicit symmetry constraints of actor and critic functions*;
- * **Independent IPPO (IPPO)**: two independent symmetry-agnostic arm-specific policies and critics for each subtask, leveraging *subtask and actor/critic decomposition*;
- * **Equivariant IPPO (eIPPO)**: a shared \mathbb{C}_2 -invariant single-arm policy (eq. 8.27) and critic, with task-conditioned observations, leveraging *symmetry of the control subtasks*;

- * **SymDex**: an IPPO variant with per-task \mathbb{C}_2 -invariant policy (eq. 8.27) and critic, leveraging *subtask and actor/critic decomposition* and *explicit symmetry constraints of actor and critic functions*;
- * **SYMDEX-aug (SM-aug)**: a SymDex variant that exploits symmetry via data augmentation, as described earlier in this section, instead of explicit invariant constraint enforcement;
- * **SYMDEX-c (SM-c)**: a SymDex variant with a centralized \mathbb{C}_2 -invariant value function, leveraging *subtask and actor decomposition* and *explicit symmetry constraints of actor/critic functions*.

The results shown in fig. 8.13 and table 8.3 indicate that methods exploiting symmetry and stronger task decomposition consistently outperform the other variants, with SymDex achieving the most robust performance. We briefly summarize these findings below and refer the reader to Li et al. (2025) for a detailed analysis.

- * **Advantage of Task Decomposition**: Decomposing tasks into subtasks and sub-action/critic modules, with lower-dimensional observation/action spaces, is critical when subtasks are distinct and mutually dependent. Although most methods achieve some success on the simpler tasks (*box-lift* and *table-clean*), ePPO, SM-c, and eIPPO fail on the remaining tasks, reflecting the difficulty of learning in larger action/observation spaces.
- * **Advantage of Symmetry Exploitation**: Exploiting symmetry through invariant constraints or data augmentation improves learning efficiency and cross-task generalization. Combined with task and actor/critic decomposition (SymDex and SM-aug), it yields the strongest overall performance. The performance gap between ePPO and IPPO further indicates that symmetry exploitation is beneficial even without full task decomposition.
- * **Symmetry Exploitation in Distillation**: In distilling the unified global policy via teacher-student behaviour cloning, we compare a symmetry-agnostic Gaussian student policy with a \mathbb{C}_2 -invariant student policy (eq. 8.27). The \mathbb{C}_2 -invariant student performs better on all six tasks, confirming the benefits of leveraging policy optimality constraints (corollary 8.2).

The primary metric is task success rate, averaged over five random seeds with 4096 rollouts per seed in simulation (fig. 8.13). Real-world evaluation reports both overall task success and per-hand subtask success over 30 independent trials.

| Gaussian student policy | box-lift | table-clean | drawer-insert | threading | bowl-stir | handover |
|---------------------------|-------------|-------------|---------------|-------------|-------------|-------------|
| Symmetry-agnostic | 0.83 ± 0.03 | 0.74 ± 0.05 | 0.69 ± 0.09 | 0.62 ± 0.13 | 0.75 ± 0.12 | 0.54 ± 0.23 |
| \mathbb{C}_2 -invariant | 0.89 ± 0.01 | 0.83 ± 0.01 | 0.87 ± 0.07 | 0.63 ± 0.17 | 0.87 ± 0.08 | 0.86 ± 0.12 |

Table 8.3 Performance of the distilled student policy on six tasks, comparing symmetry-agnostic Gaussian and \mathbb{C}_2 -invariant Gaussian parameterizations (eq. 8.27).

Dexterous multi-arm manipulation The symmetry-aware bimanual framework (SymDex) extends to systems with larger symmetry groups, such as the four-manipulator setup in fig. 8.14. This system has the morphological symmetry group $\mathbb{G} = \mathbb{C}_4 = \{e, g_c, g_c^2, g_c^3 \mid g_c^4 = e\}$, corresponding to 90° rotations about the vertical axis, which permute manipulator roles and enable a single \mathbb{C}_4 -invariant policy. Figure 8.14 shows approximate \mathbb{C}_4 -equivariant evolution from symmetry-related initial states $\mathbb{G}s_0 = \{s_0, g_c \triangleright_S s_0, g_c^2 \triangleright_S s_0, g_c^3 \triangleright_S s_0\}$, where s_0 is the upper-left state. Because the optimal policy is \mathbb{G} -invariant, trajectories from these states evolve approximately \mathbb{G} -equivariantly, and the policy solves the task from any of the four symmetry-related initial states.




Figure 8.14 Environment-policy rollout for the multi-arm task from s_0 and symmetry-related states $g_c \triangleright_S s_0$, $g_c^2 \triangleright_S s_0$, and $g_c^3 \triangleright_S s_0$. The task inherits $\mathbb{G} = \mathbb{C}_4 = \{e, g_c, g_c^2, g_c^3 \mid g_c^4 = e\}$, where g_c is a 90° rotation about the vertical axis. Figure adapted from Li et al. (2025).

Summary 8.3.2: Model-free reinforcement learning

For model-free RL in contact-rich control tasks of robots with state/morphological symmetry group \mathbb{G} , the \mathbb{G} -invariance of the optimal value and quality functions (theorem 8.1 and corollary 8.1) together with the induced symmetry constraints on policy parametrizations (corollary 8.2) induce symmetry constraints on the learning objective.

In the PPO setting, this yields a constrained optimization problem with policy and value-function invariance constraints (eq. 8.24). For Gaussian policy parametrizations, these constraints can be enforced by requiring \mathbb{G} -equivariance of the mean and covariance maps (eq. 8.27); in practice, symmetry priors can be exploited through data augmentation, regularization (eq. 8.26), or explicit equivariant/invariant architectures.

In quadruped and dexterous manipulation benchmarks, symmetry-aware PPO variants consistently improve sample efficiency and control quality over symmetry-agnostic baselines, with the strongest results generally obtained when invariance constraints are enforced explicitly (figs. 8.9 and 8.13 and table 8.3).

Chapter 9

\mathbb{G} -invariant conditional probability estimation

As discussed in [section 2.2](#) and [chapters 6 to 8](#), the stochastic forms of data-driven dynamics modelling, optimal control, and state estimation can all be reduced to the foundational learning task of approximating a conditional probability distribution from data.

This chapter presents a novel approach to conditional probability estimation by approximating the associated conditional expectation operator ([Kostic et al., 2024](#); [Ryu et al., 2024](#)). In state estimation, this operator-theoretic perspective enables parametric uncertainty quantification for unobserved quantities of interest in robotics (see [section 7.3.1](#)). In dynamics modelling, the conditional expectation operator associated with the Δt -transition dynamics $\mathbb{P}_{s|as}^{\Delta t}$ (see [eq. 2.13](#)) is known as the *controlled transfer* or *Koopman operator* ([Kostic et al., 2024](#); [Mauroy et al., 2020](#); [Turri et al., 2025](#)), and is used in [section 9.3.2](#) to learn a linear dynamics model for quadruped legged locomotion.

The aim of this chapter is to provide a self-contained introduction to operator-theoretic modelling of conditional probabilities and to discuss the implications of symmetry priors in this context ([section 9.1](#)). To approximate these operators using deep contrastive representation learning, [section 9.2](#) discusses the [Neural Conditional Probability \(NCP\)](#) framework ([Kostic et al., 2024](#)), and its symmetry-aware extension, [Equivariant Neural Conditional Probability \(eNCP\)](#) ([Ordoñez-Apaez et al., 2026](#)).

9.1 Operator-theoretic modelling of conditional distributions

Let (\mathbf{x}, \mathbf{y}) be two random variables taking values in the probability spaces $(\mathcal{X}, \Sigma_{\mathcal{X}}, \mathbb{P}_{\mathbf{x}})$ and $(\mathcal{Y}, \Sigma_{\mathcal{Y}}, \mathbb{P}_{\mathbf{y}})$, respectively. The conditional probability of an event $\mathbb{Y} \subseteq \mathcal{Y}$ given the realization $\mathbf{x} = \mathbf{x}$ is

$$\mathbb{P}(\mathbf{y} \in \mathbb{Y} | \mathbf{x} = \mathbf{x}) := \int_{\mathbb{Y}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) = \int_{\mathcal{Y}} \mathbb{1}_{\mathbb{Y}}(\mathbf{y}) \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) := \mathbb{E}[\mathbb{1}_{\mathbb{Y}}(\mathbf{y}) | \mathbf{x} = \mathbf{x}], \quad (9.1)$$

where $\mathbb{P}_{\mathbf{y}|\mathbf{x}} : \Sigma_{\mathcal{Y}} \times \mathcal{X} \mapsto [0, 1]$ is the conditional probability measure, and $\mathbb{1}_{\mathbb{Y}} : \mathcal{Y} \rightarrow \{0, 1\}$ is the characteristic function of the set \mathbb{Y} . The joint probability measure $\mathbb{P}_{\mathbf{xy}}$ is then defined as

$$\mathbb{P}_{\mathbf{xy}}(\mathbb{X}, \mathbb{Y}) := \int_{\mathbb{X}} \mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}). \quad (9.2)$$

Whenever the joint probability measure $\mathbb{P}_{\mathbf{xy}}$ is absolutely continuous with respect to the product of the marginals $\mathbb{P}_{\mathbf{x}} \times \mathbb{P}_{\mathbf{y}}$, i.e., $\mathbb{P}_{\mathbf{xy}} \ll \mathbb{P}_{\mathbf{x}} \times \mathbb{P}_{\mathbf{y}}$, there exists a unique Radon-Nikodym derivative $\kappa : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$, known as the **Pointwise Mutual Dependency (PMD)** (Tsai et al., 2020), defined as $\kappa(\mathbf{x}, \mathbf{y}) := \frac{d\mathbb{P}_{\mathbf{xy}}}{d(\mathbb{P}_{\mathbf{x}} \times \mathbb{P}_{\mathbf{y}})}(\mathbf{x}, \mathbf{y})$. Then, by construction, we have:

$$\mathbb{P}_{\mathbf{xy}}(\mathbb{X}, \mathbb{Y}) := \int_{\mathbb{X}} \int_{\mathbb{Y}} \kappa(\mathbf{x}, \mathbf{y}) \mathbb{P}_{\mathbf{y}}(d\mathbf{y}) \mathbb{P}_{\mathbf{x}}(d\mathbf{x}). \quad (9.3)$$

Combining eqs. 9.2 and 9.3 shows that the conditional probability measure can be expressed in terms of the **PMD** as $\mathbb{P}_{\mathbf{y}|\mathbf{x}}(\mathbb{Y}|\mathbf{x}) = \int_{\mathbb{Y}} \kappa(\mathbf{x}, \mathbf{y}) \mathbb{P}_{\mathbf{y}}(d\mathbf{y})$ for all $\mathbb{Y} \subseteq \mathcal{Y}$ and $\mathbf{x} \in \mathcal{X}$. This lets us compute conditional probabilities via a *marginal* expectation of an indicator function weighted by the **PMD**:

$$\mathbb{P}(\mathbf{y} \in \mathbb{Y} | \mathbf{x} = \mathbf{x}) := \mathbb{E}[\mathbb{1}_{\mathbb{Y}}(\mathbf{y}) | \mathbf{x} = \mathbf{x}] = \int_{\mathbb{Y}} \mathbb{1}_{\mathbb{Y}}(\mathbf{y}) \kappa(\mathbf{x}, \mathbf{y}) \mathbb{P}_{\mathbf{y}}(d\mathbf{y}) = [\mathbb{E}_{\mathbf{y}|\mathbf{x}} \mathbb{1}_{\mathbb{Y}}](\mathbf{x}). \quad (9.4)$$

Here, $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$ denotes the *conditional expectation operator* associated with $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$: a linear integral operator that acts on any bounded function $h : \mathcal{Y} \rightarrow \mathbb{R}$ and returns its conditional expectation, i.e., $(h, \mathbf{x}) \mapsto \mathbb{E}[h(\mathbf{y}) | \mathbf{x} = \mathbf{x}]$. In this context, the **PMD** κ is also interpreted as the kernel function defining the integral operator $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$.

This relationship between conditional probabilities and the conditional expectation operator is the key theoretical insight we leverage in this chapter: *approximating the conditional distribution $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ is equivalent to approximating the conditional expectation operator $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$* . Specifically, approximating conditional probabilities for a sufficiently rich set of relevant events $\widehat{\Sigma}_{\mathcal{Y}} \subset \Sigma_{\mathcal{Y}}$ according to $\mathbb{P}_{\mathbf{y}}$ is equivalent to approximating the action of $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$ on functions in the span of the corresponding indicator functions $\{\mathbb{1}_{\mathbb{Y}} : \mathbb{Y} \in \widehat{\Sigma}_{\mathcal{Y}}\}$. That is, approximating the operator for any $\mathbb{P}_{\mathbf{y}}$ -measurable function at conditioning points $\mathbf{x} \in \mathcal{X}$ with $\mathbb{P}_{\mathbf{x}}(d\mathbf{x}) > 0$.

Next, we provide a self-contained definition of the conditional expectation operator, constrained to act on the Hilbert spaces of square-integrable functions of the random variables \mathbf{x} and \mathbf{y} , namely $\mathcal{L}_{\mathbb{P}_{\mathbf{x}}}^2 := \mathcal{L}_{\mathbb{P}_{\mathbf{x}}}^2(\mathcal{X}, \mathbb{R})$ and $\mathcal{L}_{\mathbb{P}_{\mathbf{y}}}^2 := \mathcal{L}_{\mathbb{P}_{\mathbf{y}}}^2(\mathcal{Y}, \mathbb{R})$, respectively. We then study the implications of our assumed symmetry priors (prop. 2.3) for the structure of the conditional expectation operator.

Definition 9.1 (Conditional expectation operator). *Let (\mathbf{x}, \mathbf{y}) be two random variables defined on the measure spaces $(\mathcal{X}, \Sigma_{\mathcal{X}}, \mathbb{P}_{\mathbf{x}})$ and $(\mathcal{Y}, \Sigma_{\mathcal{Y}}, \mathbb{P}_{\mathbf{y}})$, respectively, and let $\mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{L}_{\mathbf{y}}^2$ denote the corresponding spaces of square-integrable functions. The conditional expectation operator $\mathbb{E}_{\mathbf{y}|\mathbf{x}} : \mathcal{L}_{\mathbf{y}}^2 \rightarrow \mathcal{L}_{\mathbf{x}}^2$ is a linear integral operator—defined via the **PMD** Radon–Nikodym derivative $\kappa(\mathbf{x}, \mathbf{y}) = \mathbb{P}_{\mathbf{xy}}(d\mathbf{x}, d\mathbf{y}) / \mathbb{P}_{\mathbf{x}}(d\mathbf{x}) \mathbb{P}_{\mathbf{y}}(d\mathbf{y})$ —which acts on any function $h \in \mathcal{L}_{\mathbf{y}}^2$ by computing its conditional expectation:*

$$[\mathbb{E}_{\mathbf{y}|\mathbf{x}} h](\mathbf{x}) = \mathbb{E}[h(\mathbf{y}) | \mathbf{x} = \mathbf{x}] := \int_{\mathcal{Y}} h(\mathbf{y}) \mathbb{P}_{\mathbf{y}|\mathbf{x}}(d\mathbf{y} | \mathbf{x}) = \int_{\mathcal{Y}} h(\mathbf{y}) \kappa(\mathbf{x}, \mathbf{y}) \mathbb{P}_{\mathbf{y}}(d\mathbf{y}).$$

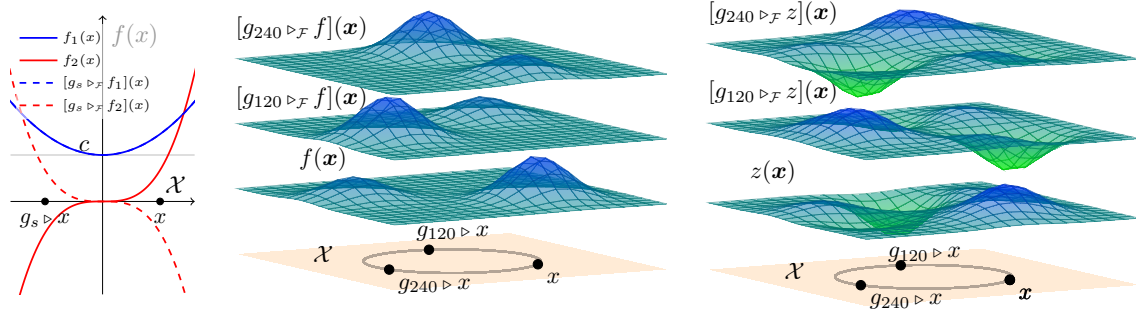


Figure 9.1 **Left:** Diagram of the group action $\triangleright_{\mathcal{F}}$ on functions $f_1(x) = x^2 + c$ and $f_2(x) = x^3$ defined on the domain $\mathcal{X} := \mathbb{R}$ endowed with the reflectional symmetry group $\mathbb{G} := \mathbb{C}_2 = \{e, g_s\}$, with the reflection action acting on the domain by $g_s \triangleright_{\mathcal{X}} x = -x$ and on the function space $\mathcal{F} := \{f \mid f : \mathcal{X} \mapsto \mathbb{R}\}$ by $[g \triangleright_{\mathcal{F}} f](x) = f(g^{-1} \triangleright_{\mathcal{X}} x) = f(-x)$. Hence, f_1 is a \mathbb{G} -invariant function, $g_s \triangleright_{\mathcal{F}} f_1(x) = f_1(x)$, and f_2 is a \mathbb{G} -equivariant function, $g_s \triangleright_{\mathcal{F}} f_2(x) = -x^3$. **Center:** Diagram representing the action $\triangleright_{\mathcal{F}}$ on the (arbitrarily chosen) function $f(x) = \mathcal{N}(x; c_1, 2) + \mathcal{N}(x; c_2, 1)$ defined over the symmetric domain $\mathcal{X} = \mathbb{R}^2$ with the cyclic symmetry group $\mathbb{G} = \mathbb{C}_3 = \{e, g_{120}, g_{240}\}$ and group action $g \triangleright_{\mathcal{X}} x = \rho_{\mathcal{X}}(g)x = R_g x$, where R_g is a rotation matrix in 2D. Here, $g_{120} \triangleright_{\mathcal{F}} f$ is equivalent to evaluating f on a domain rotated by -120° . The same holds for $g_{240} \triangleright_{\mathcal{F}} f$. Note that the z -offsets are added for visualization purposes. **Right:** Diagram representing the action $\triangleright_{\mathcal{F}}$ on the function $z \in \widehat{\mathcal{F}}$, defined to be a member of the finite-dimensional symmetric function space $\widehat{\mathcal{F}} := \text{span}(\mathbb{I}_{\widehat{\mathcal{F}}})$, constructed from a basis set composed of the group orbit of the (arbitrarily chosen) function $f \in \mathcal{F}$, that is $\mathbb{I}_{\widehat{\mathcal{F}}} := \mathbb{G}f = \{f, g_{120} \triangleright_{\mathcal{F}} f, g_{240} \triangleright_{\mathcal{F}} f\}$. This function space is \mathbb{G} -stable by construction, since $\mathbb{G}\mathbb{I}_{\widehat{\mathcal{F}}} = \mathbb{I}_{\widehat{\mathcal{F}}}$. Note that the z -offsets are added for visualization purposes.

It is important to clarify when the conditional expectation operator is not properly defined as a linear integral operator defined by **PMD** kernel, i.e., when the κ does not exist. In practice, this occurs when the joint distribution $\mathbb{P}_{\mathbf{x}\mathbf{y}}$ is not absolutely continuous with respect to the product of the marginals $\mathbb{P}_{\mathbf{x}} \times \mathbb{P}_{\mathbf{y}}$, i.e., $\mathbb{P}_{\mathbf{x}\mathbf{y}} \not\ll \mathbb{P}_{\mathbf{x}} \times \mathbb{P}_{\mathbf{y}}$. Equivalently, $\mathbb{P}_{\mathbf{x}\mathbf{y}}$ assigns non-zero probability to sets of measure zero under $\mathbb{P}_{\mathbf{x}} \times \mathbb{P}_{\mathbf{y}}$. Crucially, this includes deterministic relationships between \mathbf{x} and \mathbf{y} , i.e., when $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ is a Dirac delta distribution for every $\mathbf{x} \in \mathcal{X}$, induced by a deterministic mapping $\mathbf{y} = f(\mathbf{x})$. Note, however, that introducing mild noise in the data-generating process, e.g., $\mathbf{y} = f(\mathbf{x} + \epsilon)$ with ϵ a small noise term, is sufficient to ensure existence of the **PMD** and well-definedness of the conditional expectation operator. Since in this thesis \mathbf{x} is commonly defined by noisy state observations of a robot's dynamics, this operator is well-defined in the settings we study.

9.1.1 Symmetry priors of the conditional expectation operator

Under the assumption that $\mathbb{P}_{\mathbf{y}|\mathbf{x}}$ and $\mathbb{P}_{\mathbf{x}}$ are \mathbb{G} -invariant, by [prop. 2.3](#), the target marginal $\mathbb{P}_{\mathbf{y}}$, the joint $\mathbb{P}_{\mathbf{x}\mathbf{y}}$, and the **PMD** κ are also \mathbb{G} -invariant (see [fig. 2.3](#)). As discussed in [chapters 6 to 8](#), these symmetry priors hold for the tasks of dynamics modelling, control, and estimation in robotic systems with environmental or morphological symmetry groups \mathbb{G} . As we show next, these priors will make the conditional expectation operator $E_{\mathbf{y}|\mathbf{x}}$ a \mathbb{G} -equivariant linear operator, a crucial inductive bias that we will leverage for its approximation. To achieve this, we first need to define the symmetry actions on the function spaces $\mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{L}_{\mathbf{y}}^2$.

The \mathbb{G} -invariance of the marginal probabilities implies that the function spaces $\mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{L}_{\mathbf{y}}^2$ become symmetric Hilbert spaces, as these inherit unitary group actions $\triangleright_{\mathcal{L}_{\mathbf{x}}^2} : \mathbb{G} \times \mathcal{L}_{\mathbf{x}}^2 \rightarrow \mathcal{L}_{\mathbf{x}}^2$ and $\triangleright_{\mathcal{L}_{\mathbf{y}}^2} : \mathbb{G} \times$

$\mathcal{L}_y^2 \rightarrow \mathcal{L}_y^2$ defined via the push-forward of symmetry transformations of the data spaces. This action defines the symmetry transformed function as the original function evaluated on the transformed domain (see fig. 9.1):

$$[g \triangleright_{\mathcal{L}_x^2} f](\cdot) := f(g^{-1} \triangleright_x \cdot) \in \mathcal{L}_x^2, \quad \text{and} \quad [g \triangleright_{\mathcal{L}_y^2} h](\cdot) := h(g^{-1} \triangleright_y \cdot) \in \mathcal{L}_y^2, \quad \forall g \in \mathbb{G}. \quad (9.5)$$

A fundamental property of \mathbb{G} -symmetric Hilbert spaces is their orthogonal decomposition into $n_{\text{iso}} \leq |\mathbb{G}|$ subspaces referred to as *isotypic subspaces*: $\mathcal{L}_x^2 = \bigoplus_{k \in [1, n_{\text{iso}}]}^{\perp} \mathcal{L}_x^{2(k)}$, and $\mathcal{L}_y^2 = \bigoplus_{k \in [1, n_{\text{iso}}]}^{\perp} \mathcal{L}_y^{2(k)}$ (see theorem 2.1). where $\mathcal{L}_x^{2(k)}$ and $\mathcal{L}_y^{2(k)}$ denote the spaces of $\mathbb{G}^{(k)}$ -equivariant functions of \mathbf{x} and \mathbf{y} , with $\mathbb{G}^{(k)} := \mathbb{G}/\mathbb{N}_k$ denoting the **quotient group** induced by the normal subgroup \mathbb{N}_k , defined by the kernel of the group's k^{th} irreducible representation (see fig. 9.2 and definition 2.11). This standard result from harmonic analysis (Mackey, 1980) enables us to express any \mathbb{G} -equivariant function as a sum of its projections onto the isotypic subspaces:

$$f(\cdot) = f^{\text{inv}}(\cdot) + \sum_{k=2}^{n_{\text{iso}}} f^{(k)}(\cdot), \quad h(\cdot) = h^{\text{inv}}(\cdot) + \sum_{k=2}^{n_{\text{iso}}} h^{(k)}(\cdot), \quad \text{with } f^{(k)} \in \mathcal{L}_x^{2(k)}, h^{(k)} \in \mathcal{L}_y^{2(k)}, \forall k \in [n_{\text{iso}}], \quad (9.6)$$

where $f^{(k)}$ and $h^{(k)}$ denote the $\mathbb{G}^{(k)}$ -equivariant components of f and h . Moreover, by convention, we associate the first subspace ($k = 1$) with the space of \mathbb{G} -invariant functions, i.e., $\mathbb{G}^{(1)} = \mathbb{G}^{\text{inv}} = \{e\}$ (see fig. 9.2 and Ordoñez-Apaez et al. (2026)).

Equivariant conditional expectation operator The \mathbb{G} -invariance of the PMD kernel (definition 9.1) ensures that $E_{y|x}$ commutes with $\triangleright_{\mathcal{L}_x^2}$ and $\triangleright_{\mathcal{L}_y^2}$, and is therefore a \mathbb{G} -equivariant linear operator:

$$g \triangleright_{\mathcal{L}_x^2} [E_{y|x} h](\cdot) = [E_{y|x} [g \triangleright_{\mathcal{L}_y^2} h]](\cdot) \quad \iff \quad \begin{array}{ccc} \mathcal{L}_y^2 & \xrightarrow{\triangleright_{\mathcal{L}_y^2}} & \mathcal{L}_y^2 \\ E_{y|x} \downarrow & & \downarrow E_{y|x} \\ \mathcal{L}_x^2 & \xrightarrow{\triangleright_{\mathcal{L}_x^2}} & \mathcal{L}_x^2 \end{array} \quad (9.7)$$

This can be directly verified using the \mathbb{G} -invariance of κ and \mathbb{P}_y , and eq. 9.5:

$$\begin{aligned} [g \triangleright_{\mathcal{L}_x^2} [E_{y|x} h]](\cdot) &= [E_{y|x} h](g^{-1} \triangleright_x \cdot) = \int_{\mathcal{Y}} h(\mathbf{y}) \kappa(g^{-1} \triangleright_x \cdot, \mathbf{y}) \mathbb{P}_y(d\mathbf{y}) \\ &= \int_{\mathcal{Y}} h(\mathbf{y}) \kappa(\cdot, g \triangleright_y \mathbf{y}) \mathbb{P}_y(d\mathbf{y}) \quad (\text{by } \kappa \text{ } \mathbb{G}\text{-invariance}) \\ &= \int_{\mathcal{Y}} h(g^{-1} \triangleright_y \mathbf{y}) \kappa(\cdot, \mathbf{y}) \mathbb{P}_y(d\mathbf{y}) \quad (\text{by corollary 2.1}) \\ &= [E_{y|x} [g \triangleright_{\mathcal{L}_y^2} h]](\cdot) \end{aligned} \quad (9.8)$$

Furthermore, as a consequence of the \mathbb{G} -equivariance of $E_{y|x}$ and the isotypic decomposition of \mathcal{L}_x^2 and \mathcal{L}_y^2 , the conditional expectation operator decomposes into operators acting on the corresponding isotypic subspaces. This implies that for every $h \in \mathcal{L}_y^2$, with decomposition $h = \sum_{k=1}^{n_{\text{iso}}} h^{(k)}$ and $h^{(k)} \in \mathcal{L}_y^{2(k)}$, the action of $E_{y|x}$ on h can be expressed as the sum of the actions of $E_{y|x}$ on each isotypic component of h :

$$[E_{y|x} h](\cdot) = \sum_{k=1}^{n_{\text{iso}}} [E_{y|x}^{(k)} h^{(k)}](\cdot) \quad \text{with } E_{y|x}^{(k)} : \mathcal{L}_y^{2(k)} \rightarrow \mathcal{L}_x^{2(k)} \quad \forall k \in [n_{\text{iso}}], \quad (9.9)$$

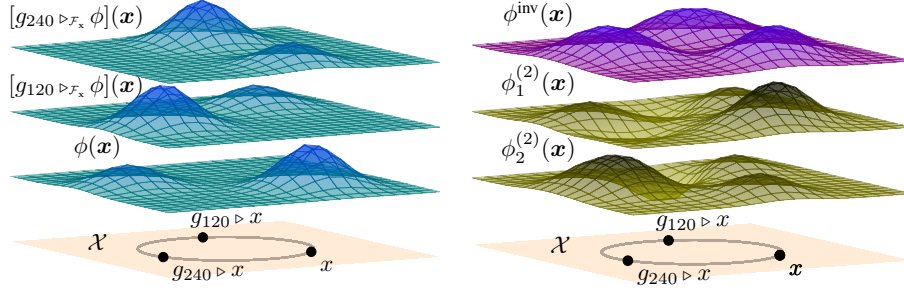


Figure 9.2 Visualization of the basis functions of a finite-dimensional symmetric function space \mathcal{F}_x . **Left:** Depiction of the basis functions in the regular basis $\mathbb{I}_{\mathcal{F}_x} = \{\phi, g_{120} \triangleright \phi, g_{240} \triangleright \phi\}$, generated by the action of the cyclic group \mathbb{C}_3 on an arbitrary function $\phi \in \mathcal{F}_x$. **Right:** Depiction of the basis functions in the isotypic basis $\mathbb{I}_{\mathcal{F}_x}^{\text{iso}} = \{u^{\text{inv}}, u_1^{(2)}, u_2^{(2)}\}$, obtained via the change of basis matrix Q . The first basis function u^{inv} corresponds to the \mathbb{G} -invariant subspace $\mathcal{F}_x^{\text{inv}}$ and is visually invariant under the action of \mathbb{C}_3 on \mathcal{X} . The other two basis functions $u_1^{(2)}, u_2^{(2)}$ are constrained to span a \mathbb{G} -stable subspace of the finite-dimensional space \mathcal{F}_x , denoted by $\mathcal{F}_x^{(2)}$, that transforms according to the irreducible representation $\bar{\rho}_{2\pi/3}$. For any function $f \in \mathcal{F}_x^{(2)}$, the group action $g \triangleright_{\mathcal{F}_x} f$ can be computed by a linear transformation of its basis expansion coefficients.

where each operator $E_{y|x}^{(k)} : \mathcal{L}_y^{2(k)} \rightarrow \mathcal{L}_x^{2(k)}$ models the conditional expectation for $\mathbb{G}^{(k)}$ -equivariant functions. Note that this decomposition is a consequence of Schur's lemma (lemma 2.1), which states that there are no non-trivial equivariant linear maps between isotypic subspaces of different types (see details in prop. 2.2).

Overall, the \mathbb{G} -equivariance of $E_{y|x}$ and its isotypic decomposition are crucial inductive biases for the approximation of $E_{y|x}$ from data, as discussed in the next sections.

Summary 9.1

Modelling a non-degenerate conditional probability distribution $\mathbb{P}_{y|x}$ can be cast as approximating its conditional expectation operator $E_{y|x} : \mathcal{L}_{\mathbb{R}^d}^2(\mathcal{Y}) \rightarrow \mathcal{L}_{\mathbb{R}^d}^2(\mathcal{X})$, a linear integral operator that maps observables of y to conditional expectations given $x = x$:

$$[E_{y|x}h](x) := \int_{\mathcal{Y}} h(y) \mathbb{P}_{y|x}(dy|x)$$

Whenever $\mathbb{P}_{y|x}$ and \mathbb{P}_x are \mathbb{G} -invariant (prop. 2.3), $E_{y|x}$ is a \mathbb{G} -equivariant linear operator that commutes with the group actions on function spaces (eq. 9.5):

$$g \triangleright_{\mathcal{L}_y^2} [E_{y|x}h](\cdot) = [E_{y|x}[g \triangleright_{\mathcal{L}_y^2} h]](\cdot) \quad \forall g \in \mathbb{G}, h \in \mathcal{L}_y^2.$$

For compact symmetry groups \mathbb{G} , the isotypic decompositions (theorem 2.1) of $\mathcal{L}_x^2 := \mathcal{L}_x^{2\text{inv}} \oplus \dots \oplus \mathcal{L}_x^{2(k)}$ and $\mathcal{L}_y^2 := \mathcal{L}_y^{2\text{inv}} \oplus \dots \oplus \mathcal{L}_y^{2(k)}$ induce an analogous decomposition of $E_{y|x}$ into independent operators between isotypic subspaces of the same type: $E_{y|x} = \sum_{k=1}^{n_{\text{iso}}} E_{y|x}^{(k)}$, with $E_{y|x}^{(k)} : \mathcal{L}_y^{2(k)} \rightarrow \mathcal{L}_x^{2(k)}$ modelling conditional expectations between $\mathbb{G}^{(k)}$ -symmetric function spaces, where $\mathbb{G}^{(k)} < \mathbb{G}$ is a quotient group.

9.2 Contrastive spectral representation learning

This section studies how the conditional expectation operator $E_{y|x}$ can be approximated from data via contrastive deep representation learning. We first present NCP (Kostic et al., 2024), and then

| | | |
|-------------------|--|--|
| Task | $f(\mathbf{x}) := \mathbb{E}_{\mathbf{y}}[\mathbf{y} \mathbf{x}=\mathbf{x}] \approx \hat{f}_{\theta}(\mathbf{x})$ | $\mathbb{P}[\mathbf{y} \in \mathbb{B} \mathbf{x} \in \mathbb{A}] \approx \hat{\mathbb{P}}_{\theta}[\mathbf{y} \in \mathbb{B} \mathbf{x} \in \mathbb{A}]$ |
| Estimate | $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbf{y}] + \left(\phi_{\theta}(\mathbf{x})^{\top} \mathbf{E}_{\theta} \hat{\mathbb{E}}_{\mathbf{y}}[\psi_{\theta}(\mathbf{y})\mathbf{y}^{\top}] \right)^{\top}$ | $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{B}}] + \frac{\hat{\mathbb{E}}_{\mathbf{x}}[\mathbb{1}_{\mathbb{A}}(\mathbf{x})\phi_{\theta}(\mathbf{x})]^{\top} \mathbf{E}_{\theta} \hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{B}}(\mathbf{y})\psi_{\theta}(\mathbf{y})]}{\hat{\mathbb{E}}_{\mathbf{x}}[\mathbb{1}_{\mathbb{A}}(\mathbf{x})]}$ |
| Guarantees | $\ f - \hat{f}_{\theta}\ _{\mathcal{L}_{\mathbf{x}}^2} \lesssim \sqrt{\text{Var}[\ \mathbf{y}\]} \left(\mathcal{E}_{\theta}^r + \frac{\ln(1/\delta)}{N \Gamma^2 2\alpha} \right)$ | $ \mathbb{P} - \hat{\mathbb{P}}_{\theta} \lesssim \sqrt{\frac{\mathbb{P}[\mathbf{y} \in \mathbb{B}]}{\mathbb{P}[\mathbf{x} \in \mathbb{A}]}} \left(\mathcal{E}_{\theta}^r + \frac{\ln(1/\delta)}{N \Gamma^2 2\alpha} \right)$ |

Table 9.1 Statistical learning guarantees of NCP (Kostic et al., 2024) for regression and conditional probability estimation, with bounds that hold with probability at least $1 - \delta$ with respect to an i.i.d. draw of $\mathbb{D}_N = \{(\mathbf{x}_n, \mathbf{y}_n) \sim \mathbb{P}_{\mathbf{x}\mathbf{y}}\}_{n=1}^N$. In the estimators, $\hat{\mathbb{E}}_{\mathbf{y}}[\psi_{\theta}(\mathbf{y})\mathbf{y}^{\top}]$, $\hat{\mathbb{E}}_{\mathbf{x}}[\mathbb{1}_{\mathbb{A}}(\mathbf{x})\phi_{\theta}(\mathbf{x})]$, and $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{B}}(\mathbf{y})\psi_{\theta}(\mathbf{y})]$ are basis-expansion coefficient moments in the learned bases $(\phi_{\theta}, \psi_{\theta})$ for vector-valued and scalar targets. The bounds are shaped by the quality of the learned representations $\mathcal{E}_{\theta}^r = \|\mathbb{E}_{\mathbf{y}|\mathbf{x}} - \mathbf{E}_{\theta}\|_{\text{op}} \leq \sqrt{\mathcal{L}_{\gamma}(\theta) - \mathcal{L}_{\gamma}(\star)}$ (see eq. 9.13), the sample size N , and the decay rate of $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$ singular-values $\alpha > 0$, which quantifies the task difficulty via the regularity of the PMD κ .

introduce its symmetry-aware extension, eNCP (Ordoñez-Apaez et al., 2026), which leverages the symmetry priors introduced in this chapter.

The learning problem of finding the best finite-dimensional truncation of the conditional expectation operator $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$ as a rank- r operator \mathbf{E}_{θ} with matrix representation $\mathbf{E}_{\theta} \in \mathbb{R}^{r \times r}$ is defined as:

$$\begin{aligned} \arg \min_{\theta} \|\mathbb{E}_{\mathbf{y}|\mathbf{x}} - \mathbf{E}_{\theta}\|_{\text{HS}}^2 &= \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} (\kappa(\mathbf{x}, \mathbf{y}) - \kappa_{\theta}(\mathbf{x}, \mathbf{y}))^2, \\ \text{s.t. } \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \kappa_{\theta}(\mathbf{x}, \mathbf{y}) &= 1, \quad \kappa_{\theta}(\mathbf{x}, \mathbf{y}) \geq 0, \forall \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}, \quad \text{and} \quad \text{rank}(\mathbf{E}_{\theta}) \leq r + 1. \end{aligned} \quad (9.10)$$

where κ_{θ} denotes the approximated PMD kernel defining the truncated operator \mathbf{E}_{θ} , constrained to ensure that κ_{θ} is a valid PMD kernel.

Given that eq. 9.10 is a low-rank approximation problem, the optimal solution, denoted \mathbf{E}_{\star} , is given by the r -truncated SVD of $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$ (Baker, 1973; Eckart and Young, 1936; Weidmann, 2012), namely

$$[\mathbf{E}_{\star} f](\mathbf{x}) = \sum_{i=0}^r \sigma_i \langle f, v_i \rangle_{\mathbb{P}_{\mathbf{y}}} u_i(\mathbf{x}), \quad \text{with} \quad \sigma_i u_i(\mathbf{x}) = [\mathbb{E}_{\mathbf{y}|\mathbf{x}} v_i](\mathbf{x}), \quad \forall i \in [r], \quad (9.11)$$

where (σ_i, u_i, v_i) denotes the i^{th} singular value and left/right singular functions of $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$, with $(\sigma_0=1, u_0=\mathbb{1}_{\mathbb{P}_{\mathbf{x}}}, v_0=\mathbb{1}_{\mathbb{P}_{\mathbf{y}}})$ being the constant functions supported on $\mathbb{P}_{\mathbf{x}}$ and $\mathbb{P}_{\mathbf{y}}$ (Kostic et al., 2024).

Consequently, NCP parameterizes the truncated operator \mathbf{E}_{θ} using a bilinear JEPAN architecture: $\kappa_{\theta}(\mathbf{x}, \mathbf{y}) = \mathbb{1}_{\mathbb{P}_{\mathbf{x}}}(\mathbf{x})\mathbb{1}_{\mathbb{P}_{\mathbf{y}}}(\mathbf{y}) + \phi_{\theta}(\mathbf{x})^{\top} \mathbf{E}_{\theta} \psi_{\theta}(\mathbf{y})$ (see fig. 9.3-left). This model comprises two encoder NNs, $\phi_{\theta}: \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_{\theta}: \mathcal{Y} \rightarrow \mathbb{R}^r$, which define the centered basis functions for the learned finite-dimensional representation spaces:

$$\mathcal{F}_{\mathbf{x}}^{\theta} := \text{span}(\mathbb{1}_{\mathbb{P}_{\mathbf{x}}}, \phi_{\theta,1}, \dots, \phi_{\theta,r}) \subset \mathcal{L}_{\mathbf{x}}^2 \quad \text{and} \quad \mathcal{F}_{\mathbf{y}}^{\theta} := \text{span}(\mathbb{1}_{\mathbb{P}_{\mathbf{y}}}, \psi_{\theta,1}, \dots, \psi_{\theta,r}) \subset \mathcal{L}_{\mathbf{y}}^2. \quad (9.12)$$

Given knowledge of the first singular triplet $(\sigma_0=1, u_0=\mathbb{1}_{\mathbb{P}_{\mathbf{x}}}, v_0=\mathbb{1}_{\mathbb{P}_{\mathbf{y}}})$, the representation functions are centered by construction; that is, $\mathbb{E}_{\mathbf{x}}[\phi_{\theta}(\mathbf{x})] = \mathbb{E}_{\mathbf{y}}[\psi_{\theta}(\mathbf{y})] = \mathbf{0}$. By eq. 9.11, the representations are trained to approximate the *span* of the top r (non-constant) left and right singular functions of $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$, namely $\mathcal{F}_{\mathbf{x}}^{\star} = \text{span}(u_0, \dots, u_r)$ and $\mathcal{F}_{\mathbf{y}}^{\star} = \text{span}(v_0, \dots, v_r)$, respectively.

Because the true operator $\mathbb{E}_{y|x}$ and its associated PMD κ are generally unknown, eq. 9.10 is optimized via the regularized contrastive low-rank (cLoRa) loss (Kostic et al., 2024; Ryu et al., 2024):

$$\begin{aligned} \mathcal{L}_{\text{cLoRa}}(\boldsymbol{\theta}, \gamma) = & -2\mathbb{E}_{\mathbf{x}, \mathbf{y}} \kappa_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y}) + \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \kappa_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y})^2 + 2\gamma (\|\mathbb{E}_{\mathbf{x}} \boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x})\|_F^2 + \|\mathbb{E}_{\mathbf{y}} \boldsymbol{\psi}_{\boldsymbol{\theta}}(\mathbf{y})\|_F^2 \\ & + \|\text{Cov}(\boldsymbol{\phi}_{\boldsymbol{\theta}}) - \text{Id}_r\|_F^2 + \|\text{Cov}(\boldsymbol{\psi}_{\boldsymbol{\theta}}) - \text{Id}_r\|_F^2), \end{aligned} \quad (9.13)$$

where the first two regularization terms center the learned representations, ensuring that $\mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \kappa_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y}) \approx 1$ (Kostic et al., 2024), while the last two promote rank- r solutions by encouraging orthonormality of the learned basis functions so that $\mathcal{F}_{\mathbf{x}}^{\boldsymbol{\theta}}$ and $\mathcal{F}_{\mathbf{y}}^{\boldsymbol{\theta}}$ (eq. 9.12) converge to r -dimensional subspaces of $\mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{L}_{\mathbf{y}}^2$, respectively.

After learning $\boldsymbol{\phi}_{\boldsymbol{\theta}}$, $\boldsymbol{\psi}_{\boldsymbol{\theta}}$, and $\mathbf{E}_{\boldsymbol{\theta}}$ by minimizing (eq. 9.13), the architecture provides estimators for regression and conditional probability estimation tasks, as detailed in table 9.1. This procedure uses the pre-trained representations $\boldsymbol{\phi}_{\boldsymbol{\theta}}$ and $\boldsymbol{\psi}_{\boldsymbol{\theta}}$ together with an additional linear regression step that expands the target function in the learned representation spaces $\mathcal{F}_{\mathbf{x}}^{\boldsymbol{\theta}}$ and $\mathcal{F}_{\mathbf{y}}^{\boldsymbol{\theta}}$. These estimators are endowed with non-asymptotic statistical learning guarantees that directly relate the quality of the learned representation spaces to the quality of conditional expectation (i.e., regression) and conditional probability estimates.

9.2.1 Equivariant contrastive spectral representation learning

The \mathbb{G} -equivariant structure of $\mathbb{E}_{y|x}$ and its decomposition into isotypic components imply that computing the conditional expectation of a \mathbb{G} -equivariant function amounts to summing the conditional expectations of its $\mathbb{G}^{(k)}$ -equivariant components for all $k \in [n_{\text{iso}}]$ (see eq. 9.9). Consequently, the low-rank approximation problem in eq. 9.10, where $\mathbb{E}_{y|x}$ is approximated by a finite-rank operator acting on the finite-dimensional representation spaces $\mathcal{F}_{\mathbf{x}}^{\boldsymbol{\theta}} \subset \mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{F}_{\mathbf{y}}^{\boldsymbol{\theta}} \subset \mathcal{L}_{\mathbf{y}}^2$, decouples into n_{iso} independent (disentangled) low-rank subproblems. In each subproblem, the conditional expectation operator for $\mathbb{G}^{(k)}$ -equivariant functions, $\mathbb{E}_{y|x}^{(k)} : \mathcal{L}_{\mathbf{y}}^{2(k)} \rightarrow \mathcal{L}_{\mathbf{x}}^{2(k)}$, is approximated by a rank- r_k operator $\mathbb{E}_{\boldsymbol{\theta}}^{(k)} : \mathcal{F}_{\mathbf{y}}^{\boldsymbol{\theta}^{(k)}} \rightarrow \mathcal{F}_{\mathbf{x}}^{\boldsymbol{\theta}^{(k)}}$, acting on finite-dimensional representation spaces $\mathcal{F}_{\mathbf{x}}^{\boldsymbol{\theta}^{(k)}} \subset \mathcal{L}_{\mathbf{x}}^{2(k)}$ and $\mathcal{F}_{\mathbf{y}}^{\boldsymbol{\theta}^{(k)}} \subset \mathcal{L}_{\mathbf{y}}^{2(k)}$. Namely:

$$\arg \min_{\boldsymbol{\theta}} \|\mathbb{E}_{y|x} - \mathbf{E}_{\boldsymbol{\theta}}\|_{\text{HS}}^2 = \sum_{k=1}^{n_{\text{iso}}} \|\mathbb{E}_{y|x}^{(k)} - \mathbf{E}_{\boldsymbol{\theta}}^{(k)}\|_{\text{HS}}^2 = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \sum_{k=1}^{n_{\text{iso}}} (\kappa^{(k)}(\mathbf{x}, \mathbf{y}) - \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}))^2, \quad (9.14a)$$

$$\text{s.t. } \kappa_{\boldsymbol{\theta}}^{(k)}(g \triangleright_{\mathcal{X}} \mathbf{x}, g \triangleright_{\mathcal{Y}} \mathbf{y}) = \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}), \quad \forall g \in \mathbb{G}^{(k)}, k \in [n_{\text{iso}}], \quad (9.14b)$$

$$\mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}) = \delta_{k, \text{inv}}, \quad \text{rank}(\mathbf{E}_{\boldsymbol{\theta}}^{(k)}) \leq r_k + \delta_{k, \text{inv}}, \quad (9.14c)$$

$$\kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}) \geq 0, \quad \forall \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}, \quad (9.14d)$$

Here $\kappa^{(k)}$ and $\kappa_{\boldsymbol{\theta}}^{(k)}$ denote the PMD kernels associated with $\mathbb{E}_{y|x}^{(k)}$ and $\mathbf{E}_{\boldsymbol{\theta}}^{(k)}$. These components aggregate to define the PMDs of $\mathbb{E}_{y|x}$ and $\mathbf{E}_{\boldsymbol{\theta}}$, respectively:

$$\kappa(\mathbf{x}, \mathbf{y}) = \sum_{k=1}^{n_{\text{iso}}} \kappa^{(k)}(\mathbf{x}, \mathbf{y}) \quad \text{and} \quad \kappa_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y}) = \sum_{k=1}^{n_{\text{iso}}} \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}) \quad \forall \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}. \quad (9.15)$$

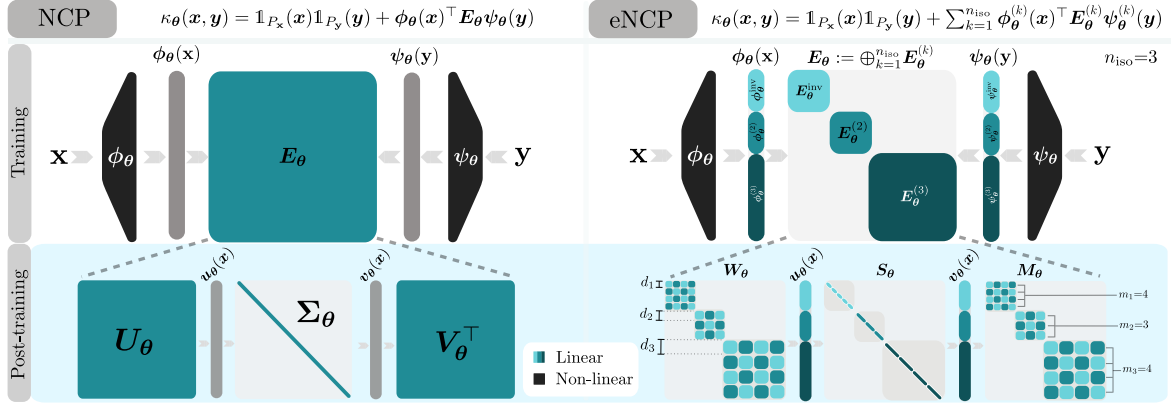


Figure 9.3 **Left:** NCP’s bilinear JEPA NN architecture. **Right:** eNCP’s \mathbb{G} -equivariant bilinear NN architecture, featuring ϕ_θ and ψ_θ as \mathbb{G} -equivariant NNs and E_θ as a \mathbb{G} -equivariant block-diagonal matrix. Each block is equivariant to a quotient group $\mathbb{G}^{(k)} = \mathbb{G}/\mathbb{N}_k$ and is constrained to have singular spaces of dimension at least d_k —the minimal dimension for a representation of the action of $\mathbb{G}^{(k)}$.

The constraints in eq. 9.14b ensure that the learned low-rank operators $E_\theta^{(k)}$ satisfy the $\mathbb{G}^{(k)}$ -equivariance of the true operators $E_{y|x}^{(k)}$, and therefore the \mathbb{G} -equivariance of $E_{y|x}$ (eq. 9.7). Furthermore, the PMDs constraints in eqs. 9.14c and 9.14d ensure that κ_θ is a valid Radon-Nikodym derivative, directly mirroring the constraints in eq. 9.10. Since the first singular triplet of $E_{y|x}$, $(\sigma_0 = 1, u_0 = \mathbb{1}_{\mathbb{P}_x}, v_0 = \mathbb{1}_{\mathbb{P}_y})$, is by construction a singular triplet of $E_{y|x}^{\text{inv}}$ —because constant functions are \mathbb{G} -invariant and therefore belong to $\mathcal{L}_x^{2\text{inv}}$ and $\mathcal{L}_y^{2\text{inv}}$ —the constraints differ for the $k = \text{inv}$ and $k \neq \text{inv}$ cases ($\delta_{k,\text{inv}} = 1$ if $k = \text{inv}$, and 0 otherwise).

Importantly, to satisfy the \mathbb{G} -invariance constraint on κ_θ , the truncated operator $E_\theta: \mathcal{F}_y^\theta \rightarrow \mathcal{F}_x^\theta$ must act on symmetric finite-dimensional Hilbert spaces that are stable under \mathbb{G} , i.e., $g \triangleright_{\mathcal{L}_x^2} f \in \mathcal{F}_x^\theta$ and $g \triangleright_{\mathcal{L}_y^2} h \in \mathcal{F}_y^\theta$ for all $f \in \mathcal{F}_x^\theta$ and $h \in \mathcal{F}_y^\theta$, and have a \mathbb{G} -equivariant matrix representation E_θ . Thus, as in the infinite-dimensional case, these finite-dimensional spaces decompose into n_{iso} isotypic subspaces $\mathcal{F}_x^\theta = \bigoplus_{k=1}^{n_{\text{iso}}} \mathcal{F}_x^{\theta(k)}$ and $\mathcal{F}_y^\theta = \bigoplus_{k=1}^{n_{\text{iso}}} \mathcal{F}_y^{\theta(k)}$. Accordingly, the truncated operator matrix decomposes block-diagonally into n_{iso} blocks, i.e., $E_\theta = \bigoplus_{k=1}^{n_{\text{iso}}} E_\theta^{(k)}$, where each k -th block needs to be a $\mathbb{G}^{(k)}$ -equivariant matrix approximating the restriction of the conditional expectation operator on its corresponding isotypic subspace (Salova et al., 2019; Ordoñez-Apraéz et al., 2024), see fig. 9.3-right.

Structure of the spectral decomposition of \mathbb{G} -equivariant operators To further motivate the constraints in eq. 9.14b and derive constraints on the truncation dimensions r_k for each isotypic component, we analyze the structure of the spectral decomposition of \mathbb{G} -equivariant operators. Analogous to eq. 9.11, the optimal truncation of a \mathbb{G} -equivariant $E_{y|x}$ is given by its truncated SVD (Weidmann, 2012). However, the spectral decomposition of the operator is constrained by the symmetry group \mathbb{G} .

Consider that if (σ_i, u_i, v_i) is a singular triplet of $E_{y|x}$, then for any $g \in \mathbb{G}$, $(\sigma_i, g \triangleright_{\mathcal{L}_x^2} u_i, g \triangleright_{\mathcal{L}_y^2} v_i)$ is also a singular triplet of $E_{y|x}$ with the same singular value (Ordoñez-Apraéz et al., 2024), given that

$$\text{if } [E_{y|x} v_i](\cdot) = \sigma_i u_i(\cdot) \quad \text{then} \quad \begin{aligned} [E_{y|x}[g \triangleright_{\mathcal{L}_y^2} v_i]](\cdot) &= g \triangleright_{\mathcal{L}_x^2} [E_{y|x} v_i](\cdot) \quad \forall g \in \mathbb{G} \\ &= \sigma_i [g \triangleright_{\mathcal{L}_x^2} u_i](\cdot). \end{aligned} \quad (9.16)$$

This constrains the dimensionality of individual singular spaces, depending on the number of distinct functions in $\mathbb{G}v_i := \{g \triangleright_{\mathcal{L}_y^{2(k)}} v_i \mid g \in \mathbb{G}\}$. To better understand these constraints, we leverage the known decomposition of the operator into suboperators per isotypic subspace, $E_{y|x} = \sum_{k=1}^{n_{\text{iso}}} E_{y|x}^{(k)}$ (eq. 9.9), which decouples the spectral decomposition of each suboperator $E_{y|x}^{(k)}$ from the others. Moreover, each isotypic subspace $\mathcal{L}_y^{2(k)}$ features only the quotient symmetry subgroup $\mathbb{G}^{(k)}$, so its spectral decomposition inherits constraints induced by $\mathbb{G}^{(k)}$. In particular, if $(\sigma_i^{(k)}, u_i^{(k)}, v_i^{(k)})$ is a singular triplet of $E_{y|x}^{(k)}$, all $\mathbb{G}^{(k)}$ -related triplets are also singular triplets, namely:

$$\text{if } \left[E_{y|x}^{(k)} v_i^{(k)} \right] (\cdot) = \sigma_i^{(k)} u_i^{(k)} (\cdot) \quad \text{then} \quad \begin{aligned} \left[E_{y|x}^{(k)} \left[g \triangleright_{\mathcal{L}_y^{2(k)}} v_i^{(k)} \right] \right] (\cdot) &= g \triangleright_{\mathcal{L}_y^{2(k)}} \left[E_{y|x}^{(k)} v_i^{(k)} \right] (\cdot) \quad \forall g \in \mathbb{G}^{(k)} \\ &= \sigma_i^{(k)} \left[g \triangleright_{\mathcal{L}_y^{2(k)}} u_i^{(k)} \right] (\cdot). \end{aligned} \quad (9.17)$$

This is a crucial inductive bias for the **NCP** architecture, because it changes the learning objective from finding the top r left/right singular functions of $E_{y|x}$ to finding the top symmetry-structured singular spaces composed of group-related left/right singular functions. Accordingly, the spectral decomposition of the true operator is constrained to be of the form:

$$\left[E_{y|x} h \right] (\cdot) = \sum_{k=1}^{n_{\text{iso}}} \left[E_{y|x}^{(k)} h^{(k)} \right] (\cdot) = \sum_{k=1}^{n_{\text{iso}}} \sum_{i=1}^{\infty} \sum_{g \in \mathbb{G}^{(k)}} \sigma_i^{(k)} \langle h^{(k)}, g \triangleright_{\mathcal{L}_y^{2(k)}} v_i^{(k)} \rangle_{\mathbb{P}_y} \left[g \triangleright_{\mathcal{L}_y^{2(k)}} u_i^{(k)} \right] (\cdot), \quad (9.18)$$

where $\sigma_i^{(k)}$ denotes the singular value of the i^{th} singular space of $E_{y|x}^{(k)}$, spanned by the symmetry-related singular functions $\mathbb{G}^{(k)} v_i^{(k)} := \{g \triangleright_{\mathcal{L}_y^{2(k)}} v_i^{(k)} \mid g \in \mathbb{G}^{(k)}\}$ and $\mathbb{G}^{(k)} u_i^{(k)} := \{g \triangleright_{\mathcal{L}_y^{2(k)}} u_i^{(k)} \mid g \in \mathbb{G}^{(k)}\}$, respectively. Crucially, by [prop. 2.1](#), these spaces have a minimum dimension of d_k , defined by the dimensionality of the irreducible representation of type k , i.e., $d_k = |\bar{\rho}_k|$ (see [fig. 9.3-right](#)). Therefore, when approximating these operators in finite dimensions, the truncation dimension r_k must be a multiple of d_k to ensure both $\mathbb{G}^{(k)}$ -equivariance and the full approximation of singular spaces.

\mathbb{G} -equivariant bilinear NN JEP A architecture To tackle the symmetry-constrained low-rank truncation problem in eq. 9.14, in [Ordoñez-Apaez et al. \(2026\)](#), we proposed an equivariant contrastive representation method termed **eNCP** (see [fig. 9.3-right](#)). The approach combines a \mathbb{G} -equivariant bilinear JEP A NN architecture with a disentangled training loss that directly leverages the isotypic decomposition of $E_{y|x}$ and the associated constraints on the spectral decomposition of \mathbb{G} -equivariant operators.

To parameterize the \mathbb{G} -equivariant truncated operator E_{θ} , we use two \mathbb{G} -equivariant NN encoders, $\phi_{\theta} : \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_{\theta} : \mathcal{Y} \rightarrow \mathbb{R}^r$, designed to expose the isotypic subspace¹, i.e.,

$$\phi_{\theta}(\cdot) = \left[\phi_{\theta}^{\text{inv}}(\cdot)^{\top}, \dots, \phi_{\theta}^{(n_{\text{iso}})}(\cdot)^{\top} \right]^{\top} \quad \text{and} \quad \psi_{\theta}(\cdot) = \left[\psi_{\theta}^{\text{inv}}(\cdot)^{\top}, \dots, \psi_{\theta}^{(n_{\text{iso}})}(\cdot)^{\top} \right]^{\top}. \quad (9.19)$$

where each component $\phi_{\theta}^{(k)} : \mathcal{X} \rightarrow \mathbb{R}^{r_k}$ and $\psi_{\theta}^{(k)} : \mathcal{Y} \rightarrow \mathbb{R}^{r_k}$ spans the approximated isotypic subspace of $\mathbb{G}^{(k)}$ -equivariant functions, i.e., $\mathcal{F}_x^{\theta(k)} := \text{span}(\phi_{\theta}^{(k)}) \subset \mathcal{L}_x^{2(k)}$ and $\mathcal{F}_y^{\theta(k)} := \text{span}(\psi_{\theta}^{(k)}) \subset \mathcal{L}_y^{2(k)}$, with $r_k = m_k d_k$ being the dimensionality of the space, governed by the chosen multiplicity m_k of the irreducible representation of type k in the learned representation spaces. Furthermore, the truncated

¹The parameterization of a NN in the isotypic basis can be achieved in `symm_learning` ([section 1.5.1](#)) using `symm_learning.nn.Change2DisentangledBasis`.

| | | |
|-------------------|---|---|
| Task | $\mathbf{f}(\mathbf{x}) := \mathbb{E}_{\mathbf{y}}[\mathbf{y} \mathbf{x}=\mathbf{x}] \approx \hat{\mathbf{f}}_{\boldsymbol{\theta}}(\mathbf{x})$ | $\mathbb{P}[\mathbf{y} \in \mathbb{B} \mathbf{x} \in \mathbb{A}] \approx \hat{\mathbb{P}}_{\boldsymbol{\theta}}[\mathbf{y} \in \mathbb{B} \mathbf{x} \in \mathbb{A}]$ |
| Estimate | $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbf{y}] + \left(\boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x})^{\top} \mathbf{E}_{\boldsymbol{\theta}} \hat{\mathbb{E}}_{\mathbf{y}}[\boldsymbol{\psi}_{\boldsymbol{\theta}}(\mathbf{y})\mathbf{y}^{\top}] \right)^{\top}$ | $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{B}}] + \frac{\hat{\mathbb{E}}_{\mathbf{x}}[\mathbb{1}_{\mathbb{A}}(\mathbf{x})\boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x})]^{\top} \mathbf{E}_{\boldsymbol{\theta}} \hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{B}}(\mathbf{y})\boldsymbol{\psi}_{\boldsymbol{\theta}}(\mathbf{y})]}{\hat{\mathbb{E}}_{\mathbf{x}}[\mathbb{1}_{\mathbb{A}}(\mathbf{x})]}$ |
| Guarantees | $\ \mathbf{f} - \hat{\mathbf{f}}_{\boldsymbol{\theta}}\ _{\mathcal{L}_{\mathbf{x}}^2} \lesssim \sqrt{\text{Var}[\ \mathbf{y}\]} \left(\mathcal{E}_{\boldsymbol{\theta}}^r + \frac{\ln(n_{\text{iso}}/\delta)}{(d_{\text{iso}}N)^{\frac{1}{1+2\alpha}}} \right)$ | $ \mathbb{P} - \hat{\mathbb{P}}_{\boldsymbol{\theta}} \lesssim \sqrt{\frac{\mathbb{P}[\mathbf{y} \in \mathbb{B}]}{\mathbb{P}[\mathbf{x} \in \mathbb{G}_{\triangleright \mathbf{x}} \mathbb{A}]}} \left(\mathcal{E}_{\boldsymbol{\theta}}^r + \frac{\ln(n_{\text{iso}}/\delta)}{(d_{\text{iso}}N)^{\frac{1}{1+2\alpha}}} \right)$ |

Table 9.2 Statistical learning guarantees for eNCP with bounds that hold with probability at least $1 - \delta$ with respect to an i.i.d. draw of $\mathbb{D}_N = \{(\mathbf{x}_n, \mathbf{y}_n) \sim \mathbb{P}_{\mathbf{x}\mathbf{y}}\}_{n=1}^N$. In the estimators, $\hat{\mathbb{E}}_{\mathbf{y}}[\boldsymbol{\psi}_{\boldsymbol{\theta}}(\mathbf{y})\mathbf{y}^{\top}]$, $\hat{\mathbb{E}}_{\mathbf{x}}[\mathbb{1}_{\mathbb{A}}(\mathbf{x})\boldsymbol{\phi}_{\boldsymbol{\theta}}(\mathbf{x})]$, and $\hat{\mathbb{E}}_{\mathbf{y}}[\mathbb{1}_{\mathbb{B}}(\mathbf{y})\boldsymbol{\psi}_{\boldsymbol{\theta}}(\mathbf{y})]$ are basis-expansion coefficient moments in the learned equivariant bases $(\boldsymbol{\phi}_{\boldsymbol{\theta}}, \boldsymbol{\psi}_{\boldsymbol{\theta}})$ for vector-valued and scalar targets. The error bounds are shaped by (i) the structure of the symmetry group \mathbb{G} —the number of isotypic subspaces $n_{\text{iso}} \leq d_{\text{iso}}$ and their minimum singular space dimensions $d_{\text{iso}} = \sum_{k=1}^{n_{\text{iso}}} d_k$ (see fig. 9.3), which enlarge the effective sample size—, (ii) the quality of the learned representations $\mathcal{E}_{\boldsymbol{\theta}}^r = \|\mathbb{E}_{\mathbf{y}|\mathbf{x}} - \mathbf{E}_{\boldsymbol{\theta}}\|_{\text{op}} \leq \sqrt{\mathcal{L}_{\gamma}(\boldsymbol{\theta}) - \mathcal{L}_{\gamma}(\star)}$, and (iii) the operator’s singular-value decay rate $\alpha > 0$. Note that $\mathbb{G}_{\triangleright \mathbf{x}} \mathbb{A} := \cup_{g \in \mathbb{G}} g \triangleright_{\mathbf{x}} \mathbb{A}$ denotes the group orbit of \mathbb{A} . For details and derivation refer to Ordoñez-Præz et al. (2026).

operator matrix is parameterized in block-diagonal form $\mathbf{E}_{\boldsymbol{\theta}} = \bigoplus_{k=1}^{n_{\text{iso}}} \mathbf{E}_{\boldsymbol{\theta}}^{(k)}$, with each block being a $r_k \times r_k$ $\mathbb{G}^{(k)}$ -equivariant matrix,² such that eq. 9.14b is satisfied by construction³.

The corresponding approximated PMD kernel is then given by:

$$\kappa_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y}) = \mathbb{1}_{\mathbb{P}_{\mathbf{x}}}(\mathbf{x})\mathbb{1}_{\mathbb{P}_{\mathbf{y}}}(\mathbf{y}) + \sum_{k=1}^{n_{\text{iso}}} \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}), \quad \text{with} \quad \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}) := \boldsymbol{\phi}_{\boldsymbol{\theta}}^{(k)}(\mathbf{x})^{\top} \mathbf{E}_{\boldsymbol{\theta}}^{(k)} \boldsymbol{\psi}_{\boldsymbol{\theta}}^{(k)}(\mathbf{y}), \quad (9.20)$$

where $\mathbb{1}_{\mathbb{P}_{\mathbf{x}}}(\mathbf{x})\mathbb{1}_{\mathbb{P}_{\mathbf{y}}}(\mathbf{y})$ are the first left and right singular constant functions of $\mathbb{E}_{\mathbf{y}|\mathbf{x}}$ (see eq. 9.11).

Crucially, this parameterization inherently satisfies the symmetry constraints of the spectral decomposition of the operator (see eq. 9.18), since any \mathbb{G} -equivariant truncated operator $\mathbf{E}_{\boldsymbol{\theta}}$ must satisfy:

$$[\mathbf{E}_{\boldsymbol{\theta}} h](\cdot) = \sum_{k=1}^{n_{\text{iso}}} [\mathbf{E}_{\boldsymbol{\theta}}^{(k)} h^{(k)}](\cdot) = \sum_{k=1}^{n_{\text{iso}}} \sum_{i=1}^{m_k} \sum_{g \in \mathbb{G}^{(k)}} \sigma_i^{\boldsymbol{\theta}^{(k)}} \langle h^{(k)}, g \triangleright_{\mathcal{L}_{\mathbf{y}}^2} v_i^{\boldsymbol{\theta}^{(k)}} \rangle_{\mathbb{P}_{\mathbf{y}}} [g \triangleright_{\mathcal{L}_{\mathbf{x}}^2} u_i^{\boldsymbol{\theta}^{(k)}}](\cdot), \quad (9.21)$$

with $\sigma_i^{\boldsymbol{\theta}^{(k)}}$ denoting the i^{th} singular value of the singular space spanned by the symmetry-related singular functions $\mathbb{G}^{(k)} v_i^{\boldsymbol{\theta}^{(k)}} = \{g \triangleright_{\mathcal{L}_{\mathbf{y}}^2} v_i^{\boldsymbol{\theta}^{(k)}} \mid g \in \mathbb{G}^{(k)}\}$ and $\mathbb{G}^{(k)} u_i^{\boldsymbol{\theta}^{(k)}} = \{g \triangleright_{\mathcal{L}_{\mathbf{x}}^2} u_i^{\boldsymbol{\theta}^{(k)}} \mid g \in \mathbb{G}^{(k)}\}$.

Disentangled training loss To train this architecture, we decompose the contrastive loss (eq. 9.13) to reflect the separability of the optimization induced by the operator’s isotypic decomposition (eq. 9.14):

$$\begin{aligned} \mathcal{L}_{\text{cLoRa}}(\boldsymbol{\theta}, \gamma) &= \sum_{k=1}^{n_{\text{iso}}} \left(-2\mathbb{E}_{\mathbf{x}\mathbf{y}} \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y}) + \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \kappa_{\boldsymbol{\theta}}^{(k)}(\mathbf{x}, \mathbf{y})^2 \right) + \\ &\gamma \left(\|\mathbb{E}_{\mathbf{x}} \boldsymbol{\phi}_{\boldsymbol{\theta}}^{\text{inv}}(\mathbf{x})\|_F^2 + \|\mathbb{E}_{\mathbf{y}} \boldsymbol{\psi}_{\boldsymbol{\theta}}^{\text{inv}}(\mathbf{y})\|_F^2 + \sum_{k=1}^{n_{\text{iso}}} \|\text{Cov}(\boldsymbol{\phi}^{(k)}) - \text{Id}_{r_k}\|_F^2 + \|\text{Cov}(\boldsymbol{\psi}^{(k)}) - \text{Id}_{r_k}\|_F^2 \right). \end{aligned} \quad (9.22)$$

This decomposes learning \mathbb{G} -equivariant representations of \mathbf{x} and \mathbf{y} into learning n_{iso} less constrained $\mathbb{G}^{(k)}$ -equivariant representations for distinct quotient groups of \mathbb{G} . Such representations are known in the literature as *disentangled* representations (Higgins et al., 2018) (see definition 2.13).

²We use square matrices for notational convenience. Dimensions for $\mathcal{F}_{\mathbf{x}}^{\boldsymbol{\theta}}$ and $\mathcal{F}_{\mathbf{y}}^{\boldsymbol{\theta}}$ need not match.

³A linear layer can be instantiated in `symm_learning` in the isotypic basis via `symm_learning.nn.eLinear` and `symm_learning.isotypic_decomp_rep`

Moreover, we improve estimates of the regularization terms in (eq. 9.13) by leveraging our symmetry priors to: (i) tighten the centering regularization (eq. 9.22), since functions in $\mathcal{F}_x^{(k)}$ and $\mathcal{F}_y^{(k)}$ are centered by construction for $k \neq \text{inv}$ (see Ordoñez-Apaez et al. (2026)); and (ii) exploit orthogonality between isotypic subspaces (eq. 9.6) to regularize orthonormality independently within each isotypic subspace (see fig. 9.2), leading to better covariance estimates (Shah and Chandrasekaran, 2012):

$$\Omega^{(k)}(\theta) := \sum_{k=1}^{n_{\text{iso}}} \|\text{Cov}(\phi^{(k)}) - \text{Id}_{r_k}\|_F^2 + \|\text{Cov}(\psi^{(k)}) - \text{Id}_{r_k}\|_F^2. \quad (9.23)$$

After learning the \mathbb{G} -equivariant representations ϕ_θ , ψ_θ , and the operator E_θ by minimizing (eq. 9.22), the architecture provides estimators for \mathbb{G} -equivariant regression and \mathbb{G} -invariant conditional probability estimation tasks, as detailed in table 9.2. While these estimators retain the same structural form as NCP, their non-asymptotic statistical learning guarantees are tighter: symmetry priors enlarge the effective sample size by a factor $d_{\text{iso}} = \sum_{k=1}^{n_{\text{iso}}} d_k$ and mitigate rare-event estimation bottlenecks (see Ordoñez-Apaez et al. (2026) for details and derivations).

Experimental validation

We present two experiments evaluating the impact of symmetry exploitation in operator-theoretic approximation of conditional probabilities. In the first experiment, presented next, we evaluate eNCP on approximation of a known conditional expectation operator for which approximation error can be computed directly. The second experiment, presented in section 7.3.1, evaluates eNCP on the downstream task of symmetry-aware uncertainty quantification of relevant unobservable quantities in legged locomotion, where reliable conditional probability estimation is critical for well-calibrated uncertainty estimates.

Conditional expectation operator learning This synthetic experiment directly compares the performance of eNCP when approximating a diverse range of conditional expectation operators with different finite symmetry groups \mathbb{G} . To the best of our knowledge, this is the first synthetic experiment that directly estimates truncation error of the conditional expectation operator in an inference-task-agnostic setting, serving as a benchmark for future work in symmetry-aware and symmetry-agnostic operator-theoretic modelling of conditional probabilities.

Specifically, we extend the conditional Gaussian Mixture Model (cGMM) of Gilardi et al. (2002) to parametrically construct \mathbb{G} -invariant vector-valued random variables $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ that satisfy the symmetry priors assumed in section 9.1.1, for arbitrary finite symmetry groups \mathbb{G} (see a 2D example in fig. 2.3). Each cGMM is defined by

$$\mathbf{z} := \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \sim \sum_{g \in \mathbb{G}} \sum_{c=1}^{n_g} \mathcal{N}(g \triangleright_{\mathcal{Z}} \mu_{\mathbf{z},c}, g \triangleright_{\mathcal{Z}} \Sigma_{\mathbf{z},c} \triangleleft g^{-1}),$$

where \mathbf{z} is a random variable of arbitrary dimension that stacks target \mathbf{y} and conditioning \mathbf{x} , \mathbb{G} is a finite symmetry group, and $\triangleright_{\mathcal{Z}} : \mathbb{G} \times \mathcal{Z} \mapsto \mathcal{Z}$ is an arbitrary group action defined by symmetries of the target and conditioning variables. Furthermore, n_g is the number of unique Gaussians with randomly sampled means $\mu_{\mathbf{z}} := \mu_{\mathbf{x}} \oplus \mu_{\mathbf{y}}$ and block-diagonal covariances $\Sigma_{\mathbf{z}} := \Sigma_{\mathbf{x}} \oplus \Sigma_{\mathbf{y}}$. It is straightforward to verify that the marginal distributions of \mathbf{x} and \mathbf{y} are \mathbb{G} -invariant, and so is their joint distribution $\mathbb{P}_{\mathbf{xy}}$ (see fig. 2.3). Moreover, since both the joint distribution and the product of marginals are GMMs, each cGMM has an analytically well-defined PMD ratio κ , enabling direct

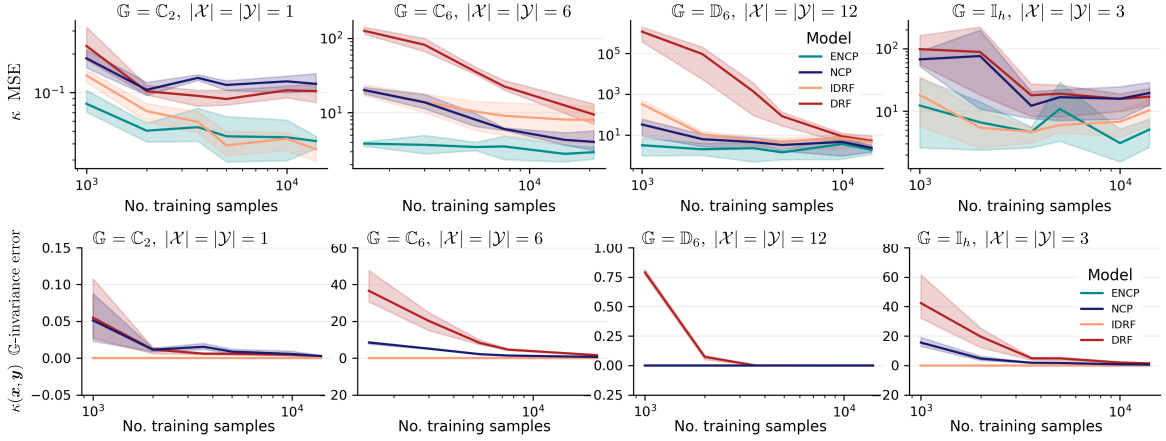


Figure 9.4 Top row: Sample-efficiency plots comparing the test-set PMD MSE $\kappa_{\text{mse}} := \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} (\kappa(\mathbf{x}, \mathbf{y}) - \kappa_{\theta}(\mathbf{x}, \mathbf{y}))^2$ versus the number of training samples, in log scales. Each column corresponds to a symmetric cGMM with distinct symmetry groups and (\mathbf{x}, \mathbf{y}) dimensionality. The tested groups are the cyclic groups \mathbb{C}_2 and \mathbb{C}_6 , the Dihedral group \mathbb{D}_6 , and the Icosahedral group \mathbb{I}_h (order 60). Bottom row: Sample-efficiency plots comparing test-set PMD G-invariance-constraint violation (eq. 9.14) $\kappa_{\text{inv-err}} := \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \sum_{g \in \mathbb{G}} (\kappa_{\theta}(\mathbf{x}, \mathbf{y}) - \kappa_{\theta}(g_{\triangleright_{\mathcal{X}} \mathbf{x}}, g_{\triangleright_{\mathcal{Y}} \mathbf{y}}))^2$ versus the number of training samples, in log scales. Each column corresponds to a symmetric cGMM with distinct symmetry groups and (\mathbf{x}, \mathbf{y}) dimensionality. The tested groups are the cyclic groups \mathbb{C}_2 and \mathbb{C}_6 , the Dihedral group \mathbb{D}_6 (order 12), and the Icosahedral group \mathbb{I}_h (order 60).

estimation of conditional expectation operator approximation error (see eqs. 9.10 and 9.14)—usually impossible for real-world datasets—via

$$\kappa_{\text{mse}} := \mathcal{L}_{\text{cLoRa}}(\boldsymbol{\theta}, \gamma) = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y}} \|\kappa(\mathbf{x}, \mathbf{y}) - \kappa_{\theta}(\mathbf{x}, \mathbf{y})\|^2.$$

The results, depicted in fig. 9.4, compare approximation quality of our eNCP model against NCP (Kostic et al., 2024), the baseline Density Ratio Fitting (DRF) (Tsai et al., 2020), and our Invariant Density Ratio Fitting (iDRF) adaptation. The baseline models approximate κ as a single NN, $\kappa_{\theta}^{\text{drf}}: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$, trained via the contrastive loss in eq. 9.13. Since they do not enforce the separable structure in fig. 9.3, neither DRF nor iDRF can be used for downstream regression or conditional probability estimation tasks. Therefore, comparison to these methods isolates the impact of the bilinear NN architecture used in eNCP and NCP.

We evaluate performance across cGMMs with diverse symmetry groups with order ranging from $|\mathbb{G}| = 2$ to $|\mathbb{G}| = 60$ and varying (\mathbf{x}, \mathbf{y}) dimensionalities. Across all settings, eNCP consistently demonstrates superior performance and sample efficiency compared to iDRF and the symmetry-agnostic models NCP and DRF (see fig. 9.4-top). Furthermore, fig. 9.4-bottom shows that the symmetry-agnostic models (NCP and DRF) struggle to satisfy the \mathbb{G} -invariance constraint of the true PMD from data alone, yielding final \mathbb{G} -invariance violations that are an order of magnitude larger than those of eNCP and iDRF, which enforce the invariance constraint in eq. 9.14 globally.

These results highlight two main conclusions: (i) the separable bilinear NN architecture of the eNCP and NCP models—namely $\phi_{\theta}: \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_{\theta}: \mathcal{Y} \rightarrow \mathbb{R}^r$ —is more capable of approximating the conditional expectation operator than the single—and more expressive—NN architecture of the DRF and iDRF models, $\kappa_{\theta}^{\text{drf}}: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$, when both architectures share a similar number of parameters; and (ii) leveraging the symmetry priors of the \mathbb{G} -equivariance of the conditional expectation operator results in significantly improved approximation quality and sample efficiency.

Summary 9.2

Approximating a conditional expectation operator $E_{y|x} : \mathcal{L}_y^2 \rightarrow \mathcal{L}_x^2$ with a finite-rank operator is a low-rank approximation learning problem (eq. 9.10). We tackle it via contrastive spectral representation learning: a bilinear JEPA NN architecture (fig. 9.3-left) learns representations $\phi_\theta : \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_\theta : \mathcal{Y} \rightarrow \mathbb{R}^r$ that approximate the leading singular subspaces of $E_{y|x}$ through the cLoRa loss (eq. 9.13).

The approximated operator can estimate both conditional expectations, $\mathbb{E}[h(y)|x]$, and conditional event probabilities, $\mathbb{P}(y \in \mathbb{Y}|x = x)$, for any $h \in \mathcal{L}_y^2$ and $\mathbb{Y} \subset \mathcal{Y}$ (table 9.1).

Whenever \mathbb{G} is compact and $E_{y|x}$ is \mathbb{G} -equivariant, approximating $E_{y|x}$ with a finite-rank \mathbb{G} -equivariant operator E_θ becomes a symmetry-constrained low-rank approximation problem (eq. 9.14). We address it via equivariant contrastive spectral representation learning, where a \mathbb{G} -equivariant bilinear JEPA NN architecture (fig. 9.3-right) learns \mathbb{G} -equivariant representations $\phi_\theta : \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_\theta : \mathcal{Y} \rightarrow \mathbb{R}^r$ that approximate the leading symmetry-constrained singular spaces of $E_{y|x}$ through a disentangled contrastive loss (eq. 9.22).

Extensive synthetic experiments demonstrate that leveraging symmetry priors is crucial for approximating \mathbb{G} -equivariant conditional expectation operators, and for improving the quality of the resulting \mathbb{G} -equivariant regression estimates and \mathbb{G} -invariant conditional probability estimates (fig. 9.4 and section 7.3.1).

9.3 Koopman/Transfer dynamics operators

In dynamical and robotic systems, a conditional expectation operator of central interest is the one associated with stochastic transition dynamics. This operator is referred to as the *transfer operator* in stochastic Markovian dynamics (Kostic et al., 2023; Novelli et al., 2024) and as the *Koopman operator* in deterministic autonomous dynamics (Mauroy et al., 2020). In this section, we study how the formalism of section 9.1 applies to operator-theoretic modelling of robot dynamics, and how the symmetry priors of the dynamics, discussed in detail in chapter 6, relate to the symmetry priors of conditional expectation operators, discussed in section 9.1.1.

Recall that the world dynamics of a robot control task in a contact-rich environment are modelled by the conditional probability of the next world state $s_{t+\Delta t}$ given the current state s_t and action a_t , which, for a sufficiently small discretization time step Δt , can be approximated by a Gaussian distribution:

$$s_{t+\Delta t} \sim \mathbb{P}_{s|as}^{\Delta t}(\cdot | a_t, s_t) \approx \mathcal{N}(F_{\Delta t}(s_t, a_t), D(s_t)D(s_t)^\top \Delta t), \quad \forall s_t \in \mathcal{S}, a_t \in \mathcal{A}, t \in \mathbb{R}.$$

Here, $\mathbb{P}_{s|as}^{\Delta t}$ denotes the Δt -transition dynamics, defined as a controlled diffusion process with deterministic drift term $F_{\Delta t}$ and diffusion term D (see eq. 2.13 and section 2.2).

This conditional distribution $\mathbb{P}_{s|as}^{\Delta t}$ can be modelled via a conditional expectation operator (definition 9.1), referred to as the controlled transfer operator $E_{s|sa}^{\Delta t} : \mathcal{L}_s^2 \rightarrow \mathcal{L}_{sa}^2$ (Novelli et al., 2024). This infinite-dimensional linear world-dynamics model evolves any state observable $h \in \mathcal{L}_s^2$ in time, given $(a, s) \in \mathcal{A} \times \mathcal{S}$, via conditional expectation:

$$[E_{s|sa}^{\Delta t} h](a, s) := \widehat{h}(s_{t+\Delta t}) = \mathbb{E}[h(s_{t+\Delta t}) | a = a, s = s] = \int_{\mathcal{S}} h(s') \kappa_{\Delta t}(s', (a, s)) \mathbb{P}_s(ds'), \quad (9.24)$$

Here, $\mathcal{L}_s^2 := \mathcal{L}_{\mathbb{P}_s}^2(\mathcal{S})$ and $\mathcal{L}_{sa}^2 := \mathcal{L}_{\mathbb{P}_{sa}}^2(\mathcal{A} \times \mathcal{S})$ are the Hilbert spaces of square-integrable functions over the state space and the action-state space, respectively, with \mathbb{P}_s and \mathbb{P}_{sa} denoting the marginal distributions of states and state-action pairs. Furthermore, $\kappa_{\Delta t} : \mathcal{S} \times (\mathcal{A} \times \mathcal{S}) \rightarrow \mathbb{R}^+$ is the PMD between the joint distribution of next-states-present-state-action triplets $\mathbb{P}_{s'as}$ and the product of marginals $\mathbb{P}_s \times \mathbb{P}_{sa}$, i.e., $\kappa_{\Delta t}(s', (a, s)) = \frac{d\mathbb{P}_{s'as}}{d(\mathbb{P}_s \times \mathbb{P}_{sa})}(s', (a, s))$, which is well defined whenever the diffusion term D is non-degenerate (see section 9.1). As we discuss next, the controlled transfer operator $E_{s|sa}^{\Delta t}$ is an *infinite-dimensional bilinear model of the dynamics*, which we approximate with a finite-dimensional bilinear model.

To approximate this operator from data using the NCP formalism (see section 9.2), we must find suitable representation functions $\psi_\theta : \mathcal{S} \rightarrow \mathbb{R}^{r_s}$ and $\phi_\theta : \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^{r_{sa}}$ spanning the finite-dimensional representation spaces $\mathcal{F}_{sa} \subset \mathcal{L}_{sa}^2$ and $\mathcal{F}_s \subset \mathcal{L}_s^2$ that approximate the span of the leading singular subspaces of $E_{s|sa}^{\Delta t}$, with $\mathcal{F}_s = \text{span}(\psi_\theta \cup \{\mathbb{1}_{\mathbb{P}_s}\}) \subset \mathcal{L}_s^2$ and $\mathcal{F}_{sa} = \text{span}(\phi_\theta \cup \{\mathbb{1}_{\mathbb{P}_{sa}}\}) \subset \mathcal{L}_{sa}^2$. Under this construction, the dynamics PMD $\kappa_{\Delta t}$ can be approximated by the bilinear model:

$$\kappa_{\theta, \Delta t}(s', (a, s)) := \mathbb{1}_{\mathbb{P}_s}(s') \mathbb{1}_{\mathbb{P}_{sa}}(a, s) + \phi_\theta(a, s)^\top \mathbf{E}_\theta \psi_\theta(s')$$

Here, $\mathbf{E}_\theta \in \mathbb{R}^{r_{sa} \times r_s}$ is the matrix representation of the truncated operator E_θ in the chosen basis sets of \mathcal{F}_{sa} and \mathcal{F}_s (see eq. 9.10). Then, the conditional expectation of any state observable $h \in \mathcal{L}_s^2$ is estimated via table 9.1 as:

$$[E_{s|sa}^{\Delta t} h](a_t, s_t) := \widehat{h}(s_{t+\Delta t}) \approx \mathbb{E}_s[h(s)] + (\phi_\theta(a_t, s_t))^\top \mathbf{E}_\theta \mathbb{E}_s[\psi_\theta(s)h(s)]. \quad (9.25)$$

The resulting approximate dynamics can be expressed in a control-theoretic familiar form as a latent bilinear dynamics model (Colonius et al., 2000, Sec. 2.7):

Proposition 9.1 (Bilinear dynamics of controlled transfer operator). *Assume that $\psi_\theta(s)$ denotes a latent state representation of s and impose the structural factorization*

$$\phi_\theta(a, s) = \mu_\theta(a) \otimes \psi_\theta(s) \in \mathbb{R}^{r_a r_s},$$

where $\mu_\theta : \mathcal{A} \rightarrow \mathbb{R}^{r_a}$ is a non-linear representation of the action space, so that $\mathbf{E}_\theta \in \mathbb{R}^{r_a r_s \times r_s}$. Under these assumptions, applying eq. 9.25 to the latent state observable, $\psi_\theta : \mathcal{S} \rightarrow \mathbb{R}^{r_s}$, yields the following approximate bilinear latent dynamics model:

$$\begin{aligned} [E_{s|sa}^{\Delta t} \psi_\theta](a_t, s_t) &:= \widehat{\psi_\theta}(s_{t+\Delta t}) \approx \mathbb{E}_s[\psi_\theta(s)] + \left((\mu_\theta(a_t) \otimes \psi_\theta(s_t))^\top \mathbf{E}_\theta \mathbb{E}_s[\psi_\theta(s) \psi_\theta(s)^\top] \right)^\top. \\ &\approx \mathbf{b}_\theta + \sum_{i=1}^{r_a} \mu_{\theta,i}(a_t) \mathbf{A}_{\theta,i} \psi_\theta(s_t) \in \mathbb{R}^{r_s}, \end{aligned}$$

with $\mathbf{b}_\theta := \mathbb{E}_s[\psi_\theta(s)] \in \mathbb{R}^{r_s}$, $\mathbf{A}_\theta = \mathbb{E}_s[\psi_\theta(s) \psi_\theta(s)^\top]^\top \mathbf{E}_\theta^\top = \begin{bmatrix} \mathbf{A}_{\theta,1} \\ \vdots \\ \mathbf{A}_{\theta,r_a} \end{bmatrix} \in \mathbb{R}^{r_a r_s \times r_s}$, and $\mathbf{A}_{\theta,i} \in \mathbb{R}^{r_s \times r_a}$.

Remark: Section 9.3

Several works assume a linear control-affine approximation of the controlled transfer operator:

$$\psi_\theta(s_{t+\Delta t}) \approx \mathbf{A}_\theta \psi_\theta(s_t) + \mathbf{B} \mu_\theta(a_t), \quad \mathbf{A}_\theta \in \mathbb{R}^{r_s \times r_s}, \mathbf{B} \in \mathbb{R}^{r_s \times r_a}.$$

While this greatly facilitates the use of linear control methods (Bevanda et al., 2022), it introduces a strong and unjustified bias: it assumes the state-action observable space can be approximated as $\mathcal{L}_{\mathbb{P}_{sa}}^2(\mathcal{S} \times \mathcal{A}) \approx \mathcal{L}_{\mathbb{P}_s}^2(\mathcal{S}) \oplus \mathcal{L}_{\mathbb{P}_a}^2(\mathcal{A})$. This structural constraint is unlikely to approximate the true right singular functions of the transfer operator and therefore leads to a biased estimate in eq. 9.10, which should be avoided when the goal is to learn an accurate dynamics model for downstream control tasks. Note that the tensor-product structure in prop. 9.1 does not introduce this bias: given orthonormal bases of $\mathcal{L}_{\mathbb{P}_s}^2(\mathcal{S})$ and $\mathcal{L}_{\mathbb{P}_a}^2(\mathcal{A})$, their tensor product is an orthonormal basis of $\mathcal{L}_{\mathbb{P}_{sa}}^2(\mathcal{S} \times \mathcal{A})$ (Reed, 2012).

\mathbb{G} -equivariant controlled transfer operators

The symmetry priors discussed in section 9.1.1 apply directly to the controlled transfer operator $E_{s|sa}^{\Delta t}$ associated with the world dynamics of contact-rich control tasks for robots featuring an environmental or morphological symmetry group \mathbb{G} (see chapters 3, 4 and 6). In such settings, the Δt -transition dynamics $\mathbb{P}_{s|as}^{\Delta t}$ is a \mathbb{G} -invariant conditional distribution (see eq. 6.4). Consequently, the controlled transfer operator is \mathbb{G} -equivariant (section 9.1.1), satisfying:

$$g \triangleright_{\mathcal{L}_{sa}^2} [E_{s|sa}^{\Delta t} h](\cdot, \cdot) = [E_{s|sa}^{\Delta t} [g \triangleright_{\mathcal{L}_s^2} h]](\cdot, \cdot) \iff \begin{array}{ccc} \mathcal{L}_s^2 & \xrightarrow{\triangleright_{\mathcal{L}_s^2}} & \mathcal{L}_s^2 \\ E_{s|sa}^{\Delta t} \downarrow & & \downarrow E_{s|sa}^{\Delta t} \\ \mathcal{L}_{sa}^2 & \xrightarrow{\triangleright_{\mathcal{L}_{sa}^2}} & \mathcal{L}_{sa}^2 \end{array} \quad (9.26)$$

Here, $\triangleright_{\mathcal{L}_{sa}^2}$ and $\triangleright_{\mathcal{L}_s^2}$ are the group actions of \mathbb{G} on the function spaces \mathcal{L}_{sa}^2 and \mathcal{L}_s^2 , respectively, induced by the group actions of \mathbb{G} on the state and action spaces (see eq. 9.5).

Then, under the assumption that the joint distribution of state-action pairs is \mathbb{G} -invariant (see discussion in assumption 6.2), approximation of the controlled transfer operator should be formulated as a symmetry-constrained learning problem (see discussion in chapter 5), as described in eq. 9.14.

In this context, the bilinear approximation of the controlled transfer operator in prop. 9.1 becomes a \mathbb{G} -equivariant bilinear model of the dynamics, as stated in the following proposition:

Proposition 9.2 (\mathbb{G} -equivariant bilinear dynamics of controlled transfer operator). *For a \mathbb{G} -equivariant controlled transfer operator $E_{s|sa}^{\Delta t}$, the bilinear dynamics model in prop. 9.1 is \mathbb{G} -equivariant. In particular, for every $g \in \mathbb{G}$ and $(\mathbf{a}_t, \mathbf{s}_t) \in \mathcal{A} \times \mathcal{S}$,*

$$[E_{s|sa}^{\Delta t} \psi_\theta](g \triangleright_{\mathcal{A}} \mathbf{a}_t, g \triangleright_{\mathcal{S}} \mathbf{s}_t) = \widehat{\psi}_\theta(g \triangleright_{\mathcal{S}} \mathbf{s}_{t+\Delta t}) \approx \mathbf{b}_\theta + \sum_{i=1}^{r_a} \mu_{\theta,i}(g \triangleright_{\mathcal{A}} \mathbf{a}_t) \mathbf{A}_{\theta,i} \psi_\theta(g \triangleright_{\mathcal{S}} \mathbf{s}_t)$$

Proof. Let $\psi_\theta(\mathbf{s}) \in \mathcal{Z}_s \subseteq \mathbb{R}^{r_s}$ and $\mu_\theta(\mathbf{a}) \in \mathcal{Z}_a \subseteq \mathbb{R}^{r_a}$ denote the \mathbb{G} -equivariant latent state and action representations of \mathbf{s} and \mathbf{a} , respectively (see section 9.2.1). Let the latent spaces carry group actions $\triangleright_{\mathcal{Z}_s}$ and $\triangleright_{\mathcal{Z}_a}$. Then the state-action tensor-product representation $\phi_\theta(\mathbf{a}, \mathbf{s}) := \mu_\theta(\mathbf{a}) \otimes \psi_\theta(\mathbf{s}) \in \mathcal{Z}_a \otimes \mathcal{Z}_s$ is \mathbb{G} -equivariant,

$$g \triangleright_{\mathcal{Z}_a \otimes \mathcal{Z}_s} \phi_\theta(\mathbf{a}, \mathbf{s}) = \phi_\theta(g \triangleright_{\mathcal{A}} \mathbf{a}, g \triangleright_{\mathcal{S}} \mathbf{s}), \quad \forall g \in \mathbb{G}, (\mathbf{a}, \mathbf{s}) \in \mathcal{A} \times \mathcal{S}.$$

where $\triangleright_{\mathcal{Z}_a \otimes \mathcal{Z}_s}$ denotes the group action on the tensor product of two \mathbb{G} -symmetric spaces (Steinberg, 2012).

When approximating the \mathbb{G} -equivariant $E_{\mathbf{s}|\mathbf{s}\mathbf{a}}^{\Delta t}$ with ϕ_θ and ψ_θ via eq. 9.14, the truncated matrix approximation satisfies $\mathbf{E}_\theta \in \text{Hom}_{\mathbb{G}}(\mathcal{Z}_{\mathbf{s}}, \mathcal{Z}_{\mathbf{a}} \otimes \mathcal{Z}_{\mathbf{s}}) \subset \mathbb{R}^{r_{\mathbf{a}} r_{\mathbf{s}} \times r_{\mathbf{s}}}$, that is, $g \triangleright_{\mathcal{Z}_{\mathbf{a}} \otimes \mathcal{Z}_{\mathbf{s}}} \mathbf{E}_\theta \cdot = \mathbf{E}_\theta(g \triangleright_{\mathcal{Z}_{\mathbf{s}}} \cdot)$, for all $g \in \mathbb{G}$. Therefore, the bilinear latent dynamics model in prop. 9.1 is \mathbb{G} -equivariant:

$$\begin{aligned} [E_{\mathbf{s}|\mathbf{s}\mathbf{a}}^{\Delta t} \psi_\theta](g \triangleright_{\mathcal{A}} \mathbf{a}_t, g \triangleright_{\mathcal{S}} \mathbf{s}_t) &\approx \mathbf{b}_\theta + \mathbf{A}_\theta \phi_\theta(g \triangleright_{\mathcal{A}} \mathbf{a}_t, g \triangleright_{\mathcal{S}} \mathbf{s}_t) \\ &= \mathbf{b}_\theta + \left(\mathbf{E}_\theta \mathbb{E}_{\mathbf{s}}[\psi_\theta(\mathbf{s}) \psi_\theta(\mathbf{s})^\top] \right)^\top g \triangleright_{\mathcal{Z}_{\mathbf{a}} \otimes \mathcal{Z}_{\mathbf{s}}} \phi_\theta(\mathbf{a}_t, \mathbf{s}_t) \\ &= \mathbf{b}_\theta + g \triangleright_{\mathcal{Z}_{\mathbf{s}}} \left(\mathbf{E}_\theta \mathbb{E}_{\mathbf{s}}[\psi_\theta(\mathbf{s}) \psi_\theta(\mathbf{s})^\top] \right)^\top \phi_\theta(\mathbf{a}_t, \mathbf{s}_t) \\ &= g \triangleright_{\mathcal{Z}_{\mathbf{s}}} \left(\mathbf{b}_\theta + \left(\mathbf{E}_\theta \mathbb{E}_{\mathbf{s}}[\psi_\theta(\mathbf{s}) \psi_\theta(\mathbf{s})^\top] \right)^\top \phi_\theta(\mathbf{a}_t, \mathbf{s}_t) \right), \\ &= g \triangleright_{\mathcal{Z}_{\mathbf{s}}} \widehat{\psi}_\theta(\mathbf{s}_{t+\Delta t}) = \widehat{\psi}_\theta(g \triangleright_{\mathcal{S}} \mathbf{s}_{t+\Delta t}). \end{aligned}$$

Here, we use that $\mathbf{b}_\theta = \int_{\mathcal{S}} \psi_\theta(\mathbf{s}) \mathbb{P}_{\mathbf{s}}(d\mathbf{s})$ is \mathbb{G} -invariant because $\mathbb{P}_{\mathbf{s}}$ is \mathbb{G} -invariant, and that the uncentred covariance commutes with the group action, i.e., $\mathbb{E}_{\mathbf{s}}[\psi_\theta(\mathbf{s}) \psi_\theta(\mathbf{s})^\top] \in \text{Hom}_{\mathbb{G}}(\mathcal{Z}_{\mathbf{s}}, \mathcal{Z}_{\mathbf{s}})$ (definition 2.7). This concludes the proof. \square

9.3.1 Closed-loop transfer/Koopman operators

A widely studied special case of the controlled transfer/Koopman operator is the closed-loop transfer operator, which arises when a stochastic control policy $\pi : \Sigma_{\mathcal{A}} \times \mathcal{S} \rightarrow [0, 1]$ is fixed. In this setting, the controlled system can be modeled as a closed-loop, autonomous, time-invariant dynamical system with dynamics:

$$\mathbf{s}_{t+\Delta t} \sim \mathbb{P}_{\mathbf{s}|\mathbf{s}}^{\Delta t}(\cdot | \mathbf{s}_t) := \int_{\mathcal{A}} \mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}(\cdot | \mathbf{a}, \mathbf{s}_t) \pi(d\mathbf{a} | \mathbf{s}_t). \quad (9.27)$$

This conditional probability distribution is modeled by a (potentially time-dependent) transfer operator $E_{\mathbf{s}|\mathbf{s}}^{t, \Delta t} : \mathcal{L}_{\mathbf{s}}^2 \rightarrow \mathcal{L}_{\mathbf{s}}^2$, acting on any state observable $h \in \mathcal{L}_{\mathbf{s}}^2$ as:

$$[E_{\mathbf{s}|\mathbf{s}}^{t, \Delta t} h](\mathbf{s}) := \widehat{h}(\mathbf{s}_{t+\Delta t}) = \mathbb{E}[h(\mathbf{s}_{t+\Delta t}) | \mathbf{s} = \mathbf{s}_t] = \int_{\mathcal{S}} h(\mathbf{s}') \kappa_{t, \Delta t}(\mathbf{s}', \mathbf{s}_t) \mathbb{P}_{\mathbf{s}}^t(d\mathbf{s}'), \quad (9.28)$$

where $\kappa_{t, \Delta t}(\mathbf{s}_{t+\Delta t}, \mathbf{s}_t) = \frac{d\mathbb{P}_{\mathbf{s}'\mathbf{s}}^{t, \Delta t}}{d(\mathbb{P}_{\mathbf{s}}^{t+\Delta t} \times \mathbb{P}_{\mathbf{s}}^t)}(\mathbf{s}_{t+\Delta t}, \mathbf{s}_t)$ is the PMD between the joint distribution of states at t and $t + \Delta t$, $\mathbb{P}_{\mathbf{s}'\mathbf{s}}^{t, \Delta t}$, and the product of marginals $\mathbb{P}_{\mathbf{s}}^{t+\Delta t} \times \mathbb{P}_{\mathbf{s}}^t$. The explicit time dependence highlights that, in general, the transfer operator varies with the evolution of $\mathbb{P}_{\mathbf{s}}^t$ and $\mathbb{P}_{\mathbf{s}'\mathbf{s}}^{t, \Delta t}$.

In most applications of transfer/Koopman operator theory in physics, it is assumed that there exists a *time-invariant* marginal state distribution $\mathbb{P}_{\mathbf{s}}$ and a Δt -joint distribution $\mathbb{P}_{\mathbf{s}'\mathbf{s}}^{\Delta t}$, so that the dynamics become a time-invariant linear system in infinite-dimensional space, defined by a single operator $E_{\mathbf{s}|\mathbf{s}}^{\Delta t}$ and its associated semigroup $\{E_{\mathbf{s}|\mathbf{s}}^{k\Delta t}\}_{k \in \mathbb{N}}$, which characterizes multi-step evolution.

From a control-theoretic perspective, the existence of such a time-invariant measure has a fundamental practical interpretation: the autonomous transfer/Koopman operator $E_{\mathbf{s}|\mathbf{s}}^{\Delta t}$ can be meaningfully approximated only for *stable* controlled dynamics. Recall that the world state $\mathbf{s} \in \mathcal{S}$ includes the robot state and all relevant variables needed to determine the robot's temporal evolution, including the control-task objective (see section 2.2). Thus, for any control task, a stochastic control policy π

induces one or more closed-loop invariant sets of the state space, analogous to robust control sets in the sense of Colonius et al. (2000).

Definition 9.2 (Closed-loop invariant set). *Let $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ be a controlled stochastic dynamical system (eq. 2.13), $\pi : \Sigma_{\mathcal{A}} \times \mathcal{S} \rightarrow [0, 1]$ a stochastic control policy, and $\mathbb{P}_{\mathbf{s}|\mathbf{s}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{S} \rightarrow [0, 1]$ the stochastic closed-loop dynamics induced by π (eq. 9.27). A measurable set $\mathbb{W} \subseteq \mathcal{S}$ is called a closed-loop invariant set of π if*

$$\mathbb{P}_{\mathbf{s}|\mathbf{s}}^{\Delta t}(\mathbb{W} \mid \mathbf{s}) = 1, \quad \forall \mathbf{s} \in \mathbb{W}.$$

The maximal closed-loop invariant set induced by π , denoted \mathbb{W}_{π} , is defined as the union of all disjoint closed-loop invariant sets of π :

$$\mathbb{W}_{\pi} := \bigcup \{ \mathbb{W} \subseteq \mathcal{S} : \mathbb{P}_{\mathbf{s}|\mathbf{s}}^{\Delta t}(\mathbb{W} \mid \mathbf{s}) = 1, \forall \mathbf{s} \in \mathbb{W} \}.$$

In practice, \mathbb{W}_{π} represents the subset of world states, under nominal or disturbed conditions, for which the control policy π can recover from perturbations and maintain the control-task objective. As discussed in Colonius et al. (2000), a closed-loop invariant set need not coincide with the basin of attraction of a stable limit cycle or fixed point; since under a given perturbation distribution, one invariant set may cover several basins of attraction of stable attractors. Moreover, the maximal invariant set need not be connected; it can consist of several disjoint invariant subsets, as illustrated in fig. 9.5.

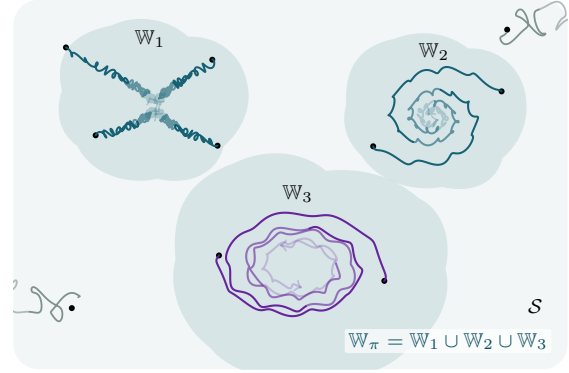


Figure 9.5 Illustration of a maximal closed-loop invariant set \mathbb{W}_{π} composed of three disjoint invariant subsets \mathbb{W}_1 , \mathbb{W}_2 , and \mathbb{W}_3 .

Since the stochasticity of $\mathbb{P}_{\mathbf{s}|\mathbf{a}\mathbf{s}}^{\Delta t}$ models uncertainty in the forcing terms acting on the robot (see section 2.2), i.e., disturbances in contact-rich environments, it is natural to assume that, under nominal operation of the control policy π for sufficiently long time horizons and control episodes, the closed-loop dynamics exhibit ergodic exploration of a time-invariant marginal state distribution $\mathbb{P}_{\mathbf{s}}$ supported on the maximal closed-loop invariant set \mathbb{W}_{π} , i.e., $\text{supp}(\mathbb{P}_{\mathbf{s}}) \subseteq \mathbb{W}_{\pi}$. This analysis enables us to formally define the transfer operator modelling closed-loop dynamics of robot control tasks in contact-rich environments:

Definition 9.3 (Closed-loop transfer operator). *Let $\mathbb{P}_{\mathbf{s}|\mathbf{s}}^{\Delta t} : \Sigma_{\mathcal{S}} \times \mathcal{S} \rightarrow [0, 1]$ be the stochastic closed-loop world dynamics induced by a control policy π (eq. 9.27), and \mathbb{W}_{π} its maximal closed-loop invariant set (definition 9.2). The closed-loop transfer operator $E_{\mathbf{s}|\mathbf{s}}^{\Delta t} : \mathcal{L}_{\mathbf{s}}^2 \rightarrow \mathcal{L}_{\mathbf{s}}^2$ is a linear integral operator (see definition 9.1) that is well-defined whenever the marginal distribution of state-next-state pairs $\mathbb{P}_{\mathbf{s}'|\mathbf{s}}^{\Delta t}$ is absolutely continuous with respect to a time-invariant marginal state distribution $\mathbb{P}_{\mathbf{s}}$ supported on a subset of \mathbb{W}_{π} , i.e., $\mathbb{P}_{\mathbf{s}'|\mathbf{s}}^{\Delta t} \ll \mathbb{P}_{\mathbf{s}} \times \mathbb{P}_{\mathbf{s}}$. Here, $\mathbb{P}_{\mathbf{s}}$ denotes a state distribution during nominal operation of π in the presence of disturbances.*

Thus, in the presence of severe disturbances that push the system outside \mathbb{W}_{π} , the unstable state-next-state transition dynamics $\mathbb{P}_{\mathbf{s}'|\mathbf{s}}^{\Delta t}$ may assign non-zero probability mass to states with zero probability under $\mathbb{P}_{\mathbf{s}}$, yielding $\mathbb{P}_{\mathbf{s}'|\mathbf{s}}^{\Delta t} \not\ll \mathbb{P}_{\mathbf{s}} \times \mathbb{P}_{\mathbf{s}}$ and making $E_{\mathbf{s}|\mathbf{s}}^{\Delta t}$ ill-defined for unstable dynamics outside \mathbb{W}_{π} . This formalism is the stochastic analog of the practical interpretation of the Koopman operator as a

linearization of the dynamics within the basin of attraction of a stable attractor (Brunton et al., 2021; Mauroy et al., 2020).

Remark: Section 9.3.1

In practice, as discussed in sections 8.3.1 and 9.3.2, the marginal distribution of states supported on the policy's maximal closed-loop invariant set \mathbb{W}_π can be empirically approximated by a dataset of stable and recoverable trajectories collected during closed-loop, stable/successful operation of a robotic control task.

Infinite-dimensional closed-loop linear dynamics models In general, the closed-loop dynamics $\mathbb{P}_{s|s}^{\Delta t}$ may be multimodal, since the control policy π can be arbitrary and stochastic. In such cases, conditional expectations are uninformative for multimodal distributions (see fig. 7.4); thus, the learned operator $\mathbb{E}_{s|s}^{\Delta t}$ should not be used for deterministic prediction, but rather to predict statistics such as probabilities of sets and quantiles of the next-state distribution, as described in section 7.3.1.

In practice, for sufficiently small time steps Δt , $\mathbb{P}_{s|s}^{\Delta t}$ can often be assumed unimodal and well approximated by a Gaussian for most states in \mathbb{W}_π (Pan et al., 2025). Under this assumption, conditional expectations of the next state represent the most likely next-state prediction. Thus, in the spirit of prop. 9.1, the closed-loop transfer operator $\mathbb{E}_{s|s}^{\Delta t}$ defines the following infinite-dimensional linear model of the closed-loop dynamics:

Proposition 9.3 (Linear dynamics of the closed-loop transfer operator). *Let $\psi_\theta : \mathcal{S} \rightarrow \mathbb{R}^{r_s}$ be a latent state representation. Applying table 9.1 (with $\phi_\theta := \psi_\theta$, since $\mathbb{E}_{s|s}^{\Delta t} : \mathcal{L}_s^2 \rightarrow \mathcal{L}_s^2$) to the observable ψ_θ yields the following approximate linear latent dynamics model:*

$$[\mathbb{E}_{s|s}^{\Delta t} \psi_\theta](s_t) := \widehat{\psi_\theta}(s_{t+\Delta t}) \approx \mathbf{A}_\theta \psi_\theta(s_t) + \mathbf{b}_\theta$$

where $\mathbf{b}_\theta := \mathbb{E}_s[\psi_\theta(s)] \in \mathbb{R}^{r_s}$, and $\mathbf{A}_\theta := (\mathbf{E}_\theta \mathbb{E}_s[\psi_\theta(s)\psi_\theta(s)^\top])^\top \in \mathbb{R}^{r_s \times r_s}$, with $\mathbf{E}_\theta \in \mathbb{R}^{r_s \times r_s}$ the matrix representation of the truncated operator \mathbb{E}_θ with $\phi_\theta = \psi_\theta$, and $\mathbb{E}_s[\psi_\theta(s)\psi_\theta(s)^\top] \in \mathbb{R}^{r_s \times r_s}$ the uncentered covariance, which equals Id_{r_s} if ψ_θ forms an orthonormal basis.

\mathbb{G} -equivariant closed-loop transfer operators

The symmetry priors discussed in section 9.1.1 transfer directly to the closed-loop transfer operator $\mathbb{E}_{s|s}^{\Delta t}$ of the world dynamics of a contact-rich control task for a robot featuring an environmental or morphological symmetry group \mathbb{G} (see chapters 3, 4 and 6), whenever the stochastic control policy π is \mathbb{G} -invariant. Note that, for these control tasks, any optimal stochastic control policy is \mathbb{G} -invariant (see corollary 8.2), and any (potentially suboptimal) stochastic control policy can be made \mathbb{G} -invariant via group averaging (see eq. 6.11 and assumption 8.1).

In such scenarios, we have that the maximal closed-loop invariant set \mathbb{W}_π is by construction \mathbb{G} -invariant (see fig. 9.6), that is

$$\mathbb{W}_\pi = g \triangleright_S \mathbb{W}_\pi, \quad \forall g \in \mathbb{G}.$$

Then, under the assumption that the initial state distribution of the control task is \mathbb{G} -invariant, $\mathbb{P}_s^0(\mathcal{S}) = \mathbb{P}_s^0(g \triangleright_S \mathcal{S})$ for all $g \in \mathbb{G}$ (see definition 8.1 and eq. 8.4c), by prop. 8.1 we get that the time-




Figure 9.6 A \mathbb{K}_4 -symmetric closed-loop invariant geometry. The upper-left seed configuration (top and lower invariant subsets from fig. 9.5) is reflected across the horizontal and vertical axes to generate the remaining quarters. Deterministic trajectory backbones are reflected by the group action, while stochastic perturbations are sampled independently per trajectory. Using the mini-cheetah notation, $\mathbb{G} = \mathbb{K}_4 = \{e, g_s, g_r, g_t | g_s^2 = g_r^2 = g_t^2 = e, g_s \circ g_t = g_r\}$, and each invariant set orbit is given by $\mathbb{G}W_i := \{g \triangleright_S W_i | g \in \mathbb{G}\}$ for $i \in \{1, 2\}$.

invariant marginal state distribution and the state-next-state distribution induced by π are \mathbb{G} -invariant:

$$\mathbb{P}_{S'|S}^{\Delta t}(\mathbb{S}', \mathbb{S}) = \mathbb{P}_{S'|S}^{\Delta t}(g \triangleright_S \mathbb{S}', g \triangleright_S \mathbb{S}), \quad \mathbb{P}_S(\mathbb{S}) = \mathbb{P}_S(g \triangleright_S \mathbb{S}), \quad \forall g \in \mathbb{G}, \mathbb{S}, \mathbb{S}' \subseteq \mathcal{S}. \quad (9.29)$$

Consequently, as discussed in section 9.1.1, the closed-loop transfer operator $E_{S|S}^{\Delta t}$ becomes a \mathbb{G} -equivariant operator, well defined on a \mathbb{G} -invariant subset of \mathbb{W}_π , satisfying:

$$g \triangleright_{\mathcal{L}_S^2} [E_{S|S}^{\Delta t} h](\cdot) = [E_{S|S}^{\Delta t} (g \triangleright_{\mathcal{L}_S^2} h)](\cdot) \iff \begin{array}{ccc} \mathcal{L}_S^2 & \xrightarrow{\triangleright_{\mathcal{L}_S^2}} & \mathcal{L}_S^2 \\ E_{S|S}^{\Delta t} \downarrow & & \downarrow E_{S|S}^{\Delta t} \\ \mathcal{L}_S^2 & \xrightarrow{\triangleright_{\mathcal{L}_S^2}} & \mathcal{L}_S^2 \end{array} \quad (9.30)$$

Similar to prop. 9.3, for sufficiently small time steps Δt the closed-loop transition dynamics $\mathbb{P}_{S|S}^{\Delta t}$ can be assumed unimodal and well approximated by a Gaussian. Under this assumption, the closed-loop transfer operator $E_{S|S}^{\Delta t}$ defines an infinite-dimensional \mathbb{G} -equivariant linear model of the closed-loop dynamics. Crucially, the \mathbb{G} -equivariance of the operator and of its low-rank approximation (see section 9.2.1) yields a disentangled time-invariant approximate linear dynamics model on each isotypic subspace of \mathcal{L}_S^2 (see discussion in section 9.1.1 and Ordoñez-Apraez et al. (2024)), summarized in the following proposition:

Proposition 9.4 (Disentangled closed-loop \mathbb{G} -equivariant linear dynamics models). *Let $\psi_\theta : \mathcal{S} \rightarrow \mathbb{R}^{r_s}$ be a latent state representation exposing the isotypic basis (see eq. 9.19),*

$$\psi_\theta(\cdot) = [\psi_\theta^{inv}(\cdot)^\top, \dots, \psi_\theta^{(n_{iso})}(\cdot)^\top]^\top \in \mathbb{R}^{r_s} \quad \text{with } \psi_\theta^{(k)}(\cdot) \in \mathbb{R}^{r_{k,s}}, \quad r_s = \sum_{k=1}^{n_{iso}} r_{k,s}.$$

Then, the closed-loop approximate linear dynamics model in prop. 9.3 becomes disentangled over isotypic subspaces, yielding independent linear dynamics models on each isotypic subspace:

$$\begin{aligned} [\mathbb{E}_{s|s}^{\Delta t} \psi_\theta](s_t) &:= \widehat{\psi}_\theta(s_{t+\Delta t}) \approx \mathbf{A}_\theta \psi_\theta(s_t) + \mathbf{b}_\theta \\ &\approx \left(\bigoplus_{k=1}^{n_{iso}} \mathbf{A}_\theta^{(k)} \psi_\theta^{(k)}(s_t) \right) + \mathbf{b}_\theta \\ &\approx \begin{bmatrix} \mathbf{A}_\theta^{inv} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_\theta^{(2)} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{A}_\theta^{(n_{iso})} \end{bmatrix} \begin{bmatrix} \psi_\theta^{inv}(s_t) \\ \psi_\theta^{(2)}(s_t) \\ \vdots \\ \psi_\theta^{(n_{iso})}(s_t) \end{bmatrix} + \begin{bmatrix} \mathbf{b}_\theta^{inv} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \end{aligned}$$

where $\mathbf{b}_\theta^{inv} := \mathbb{E}_s[\psi_\theta^{inv}(s)] \in \mathbb{R}^{r_{1,s}}$, the matrix representation of the truncated operator decomposes block-diagonally as $\mathbf{E}_\theta = \bigoplus_{k=1}^{n_{iso}} \mathbf{E}_\theta^{(k)}$, and each $\mathbf{A}_\theta^{(k)} := \left(\mathbf{E}_\theta^{(k)} \mathbb{E}_s[\psi_\theta^{(k)}(s)\psi_\theta^{(k)}(s)^\top] \right)^\top \in \mathbb{R}^{r_{k,s} \times r_{k,s}}$, is a $\mathbb{G}^{(k)}$ -equivariant matrix.

Proof. The result follows from the finite-rank approximation of the \mathbb{G} -equivariant operator $\mathbb{E}_{s|s}^{\Delta t} : \mathcal{L}_s^2 \rightarrow \mathcal{L}_s^2$ with finite-dimensional representations $\phi_\theta := \psi_\theta$ (see eqs. 9.14 and 9.20). In the isotypic basis of the latent representation space, the matrix representation \mathbf{E}_θ of the truncated operator is block diagonal and decomposes into blocks $\mathbf{E}_\theta^{(k)} \in \mathbb{R}^{r_{k,s} \times r_{k,s}}$, each representing the truncated $\mathbb{G}^{(k)}$ -equivariant operator $\mathbb{E}_{s|s}^{\Delta t(k)}$ (eq. 9.20).

By Schur's orthogonality relations (lemma 2.1), the uncentered covariance term is also block diagonal, decomposing into the uncentered covariance of the features in each isotypic subspace:

$$\mathbb{E}_s[\psi_\theta(s)\psi_\theta(s)^\top] = \bigoplus_{k=1}^{n_{iso}} \mathbb{E}_s[\psi_\theta^{(k)}(s)\psi_\theta^{(k)}(s)^\top].$$

With each $\mathbb{E}_s[\psi_\theta^{(k)}(s)\psi_\theta^{(k)}(s)^\top]$ commuting with the group action of $\mathbb{G}^{(k)}$ on the corresponding isotypic subspace of the latent representation space. Therefore, \mathbf{A}_θ decomposes block-diagonally.

Furthermore, the bias term $\mathbf{b}_\theta = \mathbb{E}_s[\psi_\theta(s)]$ is \mathbb{G} -invariant, owing to the \mathbb{G} -invariance of \mathbb{P}_s . Consequently, the expectation resides in the \mathbb{G} -invariant subspace of the latent representation space (see theorem 2.1), implying that $\mathbb{E}_s[\psi_\theta^{(k)}(s)] = \mathbf{0}$ for all $k \neq 1$. This concludes the proof. \square

Summary 9.3

Modelling the stochastic world dynamics of a robot control task in a contact-rich environment—represented by $\mathbb{P}_{s|sa}^{\Delta t}$ and its associated controlled transfer operator $\mathbb{E}_{s|sa}^{\Delta t}$ —leads to an infinite-dimensional bilinear dynamics model. This model can be approximated in finite dimensions, via eq. 9.10, as a latent bilinear system (see prop. 9.1):

$$[\mathbb{E}_{s|sa}^{\Delta t} \psi_\theta](a_t, s_t) := \widehat{\psi}_\theta(s_{t+\Delta t}) \approx \mathbf{b}_\theta + \sum_{i=1}^{r_a} \mu_{\theta,i}(a_t) \mathbf{A}_{\theta,i} \psi_\theta(s_t) \in \mathbb{R}^{r_s}.$$

where $\mu_{\theta,i} : \mathcal{A} \rightarrow \mathbb{R}^{r_a}$ and $\psi_{\theta} : \mathcal{S} \rightarrow \mathbb{R}^{r_s}$ denote the latent action and state representations, respectively.

If we are interested in modelling the world's closed-loop dynamics under a fixed control policy π , the closed-loop transfer operator $E_{s|s}^{\Delta t}$ defines an infinite-dimensional linear dynamics model that can be approximated in finite dimensions as (see prop. 9.3):

$$[E_{s|s}^{\Delta t} \psi_{\theta}](s_t) := \widehat{\psi_{\theta}}(s_{t+\Delta t}) \approx A_{\theta} \psi_{\theta}(s_t) + b_{\theta} \in \mathbb{R}^{r_s}.$$

This time-invariant linear dynamics model is well defined for *stable* closed-loop dynamics, i.e., on the support of a marginal state distribution \mathbb{P}_s supported on the policy's maximal closed-loop invariant (see definitions 9.2 and 9.3).

For robotic systems with a symmetry group \mathbb{G} , the controlled transfer operator $E_{s|sa}^{\Delta t}$ is \mathbb{G} -equivariant, resulting in the bilinear finite-dimensional approximation also being \mathbb{G} -equivariant (see prop. 9.2).

When modelling closed-loop dynamics, if the control policy π is \mathbb{G} -invariant, the closed-loop transfer operator $E_{s|s}^{\Delta t}$ is also \mathbb{G} -equivariant (see eq. 9.30), and its finite-dimensional approximation yields a block-diagonal linear dynamics model:

$$[E_{s|s}^{\Delta t} \psi_{\theta}](s_t) := \widehat{\psi_{\theta}}(s_{t+\Delta t}) \approx \begin{bmatrix} A_{\theta}^{\text{inv}} & 0 & \dots & 0 \\ 0 & A_{\theta}^{(2)} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A_{\theta}^{n_{\text{iso}}} \end{bmatrix} \begin{bmatrix} \psi_{\theta}^{\text{inv}}(s_t) \\ \psi_{\theta}^{(2)}(s_t) \\ \vdots \\ \psi_{\theta}^{n_{\text{iso}}}(s_t) \end{bmatrix} + \begin{bmatrix} b_{\theta}^{\text{inv}} \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

9.3.2 Modelling robot dynamics via closed-loop Koopman/Transfer operators

This section demonstrates how a closed-loop transfer operator can be employed to model the controlled dynamics of robotic systems in contact-rich environments, and how prop. 9.4 enables the learning of disentangled linear dynamics models for robots with a compact symmetry group \mathbb{G} (see fig. 9.7).

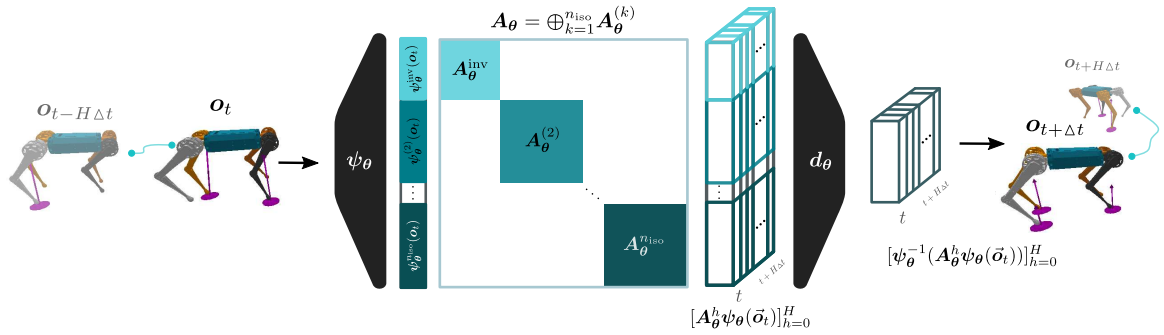


Figure 9.7 Diagram of the Equivariant Dynamics Autoencoder (eDAE) architecture (Ordoñez-Apraéz et al., 2024), which leverages the prior knowledge of the \mathbb{G} -equivariance of the closed-loop transfer operator $E_{s|s}^{\Delta t}$ to learn a finite-dimensional \mathbb{G} -equivariant approximation of the operator E_{θ} (see prop. 9.4). Where in the isotypic basis, the \mathbb{G} -equivariance of the truncated operator leads to its block-diagonal decomposition into blocks $E_{\theta}^{(k)}$, each representing a $\mathbb{G}^{(k)}$ -equivariant approximation of the operator $E_{s|s}^{\Delta t(k)}$ where $\mathbb{G}^{(k)} < \mathbb{G}$ is a quotient group of \mathbb{G} (see fig. 9.8). To use the model for multi-step prediction, one can simply power the approximated operator, enabling multi-step predictions in linear time with respect to the prediction horizon.

Koopman/Transfer operators under partial observability Note that the controlled and closed-loop transfer operators $E_{s|sa}^{\Delta t}$ and $E_{s|s}^{\Delta t}$ defined in section 9.3 are conditioned on the current state and action. When states are only partially observable, as is typical in robotics, analogous operators can be derived by following the same analysis as in section 9.3, but for the conditional distribution modelling the state observation dynamics $\mathbb{P}_{o|a\vec{o}}^{\Delta t} : \Sigma_{\mathcal{O}} \times \mathcal{A} \times \mathcal{O}^H \rightarrow [0, 1]$ (see eq. 2.15) as a function of the current action and past observation history $\vec{o} \in \mathcal{O}^H$.

For any (belief) control policy $\pi : \Sigma_{\mathcal{A}} \times \mathcal{O}^H \rightarrow [0, 1]$ (see eq. 8.10), the closed-loop observation dynamics and the observation history dynamics are given by:

$$\mathbf{o}_{t+\Delta t} \sim \mathbb{P}_{o|\vec{o}}^{\Delta t}(\cdot|\vec{o}_t) := \int_{\mathcal{A}} \mathbb{P}_{o|a\vec{o}}^{\Delta t}(\cdot|a, \vec{o}_t) \pi(da|\vec{o}_t) \implies \vec{o}_{t+\Delta t} \sim \mathbb{P}_{\vec{o}|\vec{o}}^{\Delta t}(\cdot|\vec{o}_t) \quad (9.31)$$

Here, $\mathbb{P}_{\vec{o}|\vec{o}}^{\Delta t}$ denotes the dynamics of observation trajectories induced by $\mathbb{P}_{o|\vec{o}}^{\Delta t}$ and the shift operator, which discards the oldest observation and appends the new observation to the history.

Analogous to the closed-loop state dynamics (eq. 9.27), $\mathbb{P}_{\vec{o}|\vec{o}}^{\Delta t}$ can be modeled by a closed-loop transfer operator $E_{\vec{o}|\vec{o}}^{\Delta t} : \mathcal{L}_{\vec{o}}^2 \rightarrow \mathcal{L}_{\vec{o}}^2$, where $\mathcal{L}_{\vec{o}}^2 := \mathcal{L}_{\mathbb{P}_{\vec{o}}}^2(\mathcal{O}^H)$ is the space of square-integrable functions over observation trajectories, and $\mathbb{P}_{\vec{o}}$ is a time-invariant marginal distribution of observation trajectories, induced by the time-invariant marginal distribution \mathbb{P}_s supported on a subset of the policy's maximal closed-loop invariant set \mathbb{W}_{π} (see definitions 9.2 and 9.3).

Approximating this operator in finite dimensions, using a latent state representation $\psi_{\theta} : \mathcal{O}^H \rightarrow \mathbb{R}^{r_s}$, yields the following *latent* linear dynamics model, analogous to the one in prop. 9.3, but defined over the observation history space:

$$[E_{\vec{o}|\vec{o}}^{\Delta t} \psi_{\theta}](\vec{o}_t) := \widehat{\psi}_{\theta}(\vec{o}_{t+\Delta t}) \approx \mathbf{A}_{\theta} \psi_{\theta}(\vec{o}_t) + \mathbf{b}_{\theta} \in \mathbb{R}^{r_s}. \quad (9.32)$$

To recover predictions of $\mathbf{o}_{t+\Delta t}$, a linear or non-linear decoder $\mathbf{d}_{\theta} : \mathbb{R}^{r_s} \rightarrow \mathcal{O}$ can be trained to map latent state representations $\widehat{\psi}_{\theta}(\vec{o}_{t+\Delta t})$ to the next observation $\widehat{o}_{t+\Delta t}$, as depicted in fig. 9.7.

For eq. 9.32 to provide meaningful predictions, the observation history horizon H and the time-discretization step Δt should be selected so that $\mathbb{P}_{\vec{o}|\vec{o}}^{\Delta t}$ is well approximated by a unimodal Gaussian distribution. In this regime, the conditional expectation in eq. 9.32 serves as a reliable predictor (see section 9.3). If Δt is too large, multimodality may arise due to the underlying nature of the dynamics and control policy; if H is too short, the observation history may be insufficient to disambiguate between multiple possible contact modes of the robot. Hence, these represent fundamental hyperparameters of the linear dynamics approximation that depend entirely on the complexity of the closed-loop dynamics.

Disentangled linear dynamics models for symmetric robots For world dynamics of contact-rich control tasks involving robots with an environmental or morphological symmetry group \mathbb{G} (see chapters 3 and 4), if the observation model is \mathbb{G} -invariant (see assumption 6.1), and the control policy π is \mathbb{G} -invariant (see assumption 8.1 and corollary 8.2), $\mathbb{P}_{\vec{o}|\vec{o}}^{\Delta t}$ becomes a \mathbb{G} -invariant conditional distribution, and consequently the closed-loop transfer operator $E_{\vec{o}|\vec{o}}^{\Delta t}$ becomes a \mathbb{G} -equivariant operator. Finite-dimensional approximations of this operator that preserve its \mathbb{G} -equivariance yield disentangled linear dynamics models, analogous to the one in prop. 9.4, but defined over the observation history space:

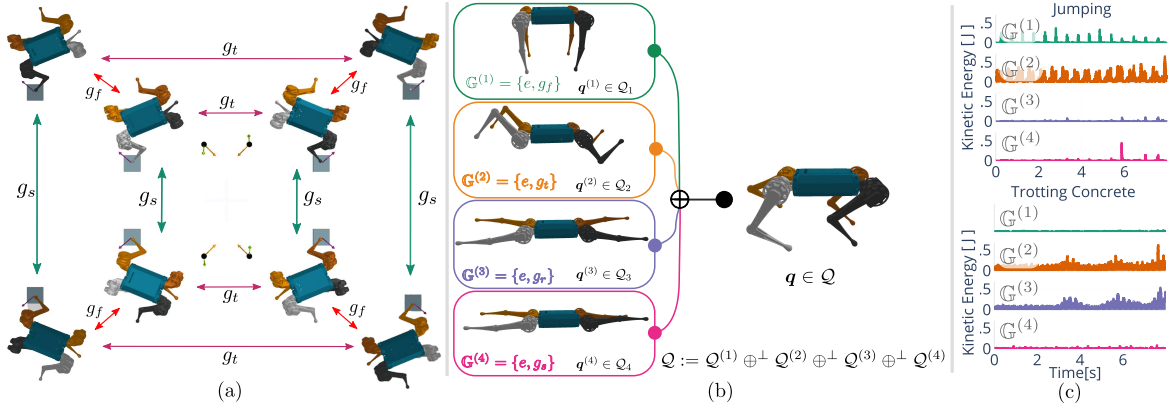


Figure 9.8 (a) Cayley diagram of the morphological symmetry group of the Mini-Cheetah quadruped robot, $\mathbb{G} = \{e, g_s, g_t, g_f, g_s \circ g_t, g_s \circ g_f, g_s \circ g_f \circ g_t, g_f \circ g_t | g_s^2 = g_t^2 = g_f^2 = e\}$, where g_s represents a sagittal plane reflection, g_t a transverse plane reflection, and g_f a frontal plane reflection. (b) The isotypic decomposition of the configuration manifold of the Mini-Cheetah, with each isotypic submanifold corresponding to the configuration space of a distinct NCM (see section 4.3). (c) Excitation of individual NCMs during real-world jumping and trotting gaits of the Mini-Cheetah, with the excitation being measured by the kinetic energy of the projection of the robot's configuration trajectory onto each isotypic submanifold (see section 4.3 for details).

$$[E_{\vec{\sigma}|\vec{\sigma}}^{\Delta t} \psi_{\theta}](\vec{\sigma}_t) := \widehat{\psi}_{\theta}(\vec{\sigma}_{t+\Delta t}) = \begin{bmatrix} \widehat{\psi}_{\theta}^{\text{inv}}(\vec{\sigma}_{t+\Delta t}) \\ \widehat{\psi}_{\theta}^{(2)}(\vec{\sigma}_{t+\Delta t}) \\ \vdots \\ \widehat{\psi}_{\theta}^{(n_{\text{iso}})}(\vec{\sigma}_{t+\Delta t}) \end{bmatrix} \approx \begin{bmatrix} \mathbf{A}_{\theta}^{\text{inv}} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{\theta}^{(2)} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{A}_{\theta}^{n_{\text{iso}}} \end{bmatrix} \begin{bmatrix} \psi_{\theta}^{\text{inv}}(\vec{\sigma}_t) \\ \psi_{\theta}^{(2)}(\vec{\sigma}_t) \\ \vdots \\ \psi_{\theta}^{n_{\text{iso}}}(\vec{\sigma}_t) \end{bmatrix} + \begin{bmatrix} \mathbf{b}_{\theta}^{\text{inv}} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \quad (9.33)$$

Here, $\mathbf{A}_{\theta}^{(k)}$ is the $\mathbb{G}^{(k)}$ -equivariant linear dynamics model acting on the k -th isotypic subspace of the latent representation space (see prop. 9.4 and section 9.1.1).

This global decomposition of the latent robot dynamics into independent linear models on each isotypic subspace is a direct stochastic generalization of the *dynamics' harmonic analysis* concept from section 4.3. There, the deterministic dynamics of robots with symmetry group \mathbb{G} are shown to admit an orthogonal decomposition of the configuration space into isotypic subspaces. Eq. 9.33 extends this to the stochastic setting, showing that \mathbb{G} -invariant stochastic world dynamics can be globally decomposed into independent linear models on each isotypic subspace of the latent space. If the decoder d_{θ} is linear, each $\mathbf{A}_{\theta}^{(k)}$ provides a linear dynamics model for the dynamics of each isotypic subspace (i.e., each NCM; see section 4.3).

As an example, consider the Mini-Cheetah quadruped robot, which possesses a $\mathbb{G} = \mathbb{C}_2 \times \mathbb{C}_2 \times \mathbb{C}_2 \equiv \mathbb{K}_4 \times \mathbb{C}_2$ symmetry group (see figs. 2.1 and 9.8-left). Its configuration manifold decomposes into four orthogonal isotypic subspaces (or submanifolds), each corresponding to the configuration space of a distinct NCM, as illustrated in fig. 9.8-middle. Each subspace is associated with a specific quotient group $\{\mathbb{G}^{(k)} < \mathbb{G}\}_{k=1}^4$ (see section 4.3 for details). For any \mathbb{G} -equivariant linear decoder d_{θ} , each latent linear subsystem $\mathbf{A}_{\theta}^{(k)}$ provides a high-dimensional linear approximation of the dynamics within its corresponding isotypic subspace, and thus of each NCM. Furthermore, the eigendecomposition of each $\mathbf{A}_{\theta}^{(k)}$ yields an independent dynamic mode decomposition of the dynamics within each isotypic subspace, thus representing a symmetry-aware mode decomposition of the robot's dynamics.

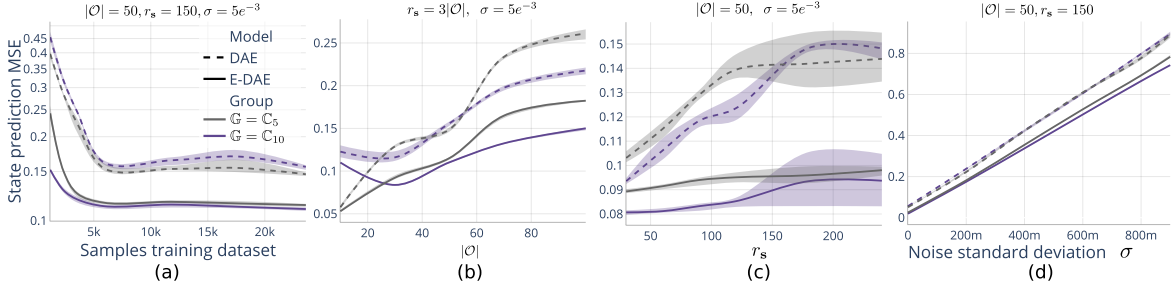


Figure 9.9 Test set prediction mean squared error (MSE) of DAE and eDAE for synthetic dynamical systems with symmetry groups \mathbb{C}_5 and \mathbb{C}_{10} , varying state dimension $|\mathcal{O}|$, latent model space dimension r_s , and noise variance σ . Solid lines and shaded areas represent the mean, maximum, and minimum prediction error among 4 training seeds. (a) Sample efficiency: test set MSE vs. training dataset size. (b) MSE as a function of observation space dimensionality. (c) MSE as a function of latent state dimension r_s . (d) MSE as a function of noise scale. Figure adapted from Ordoñez-Appraez et al. (2024).

In the next section, we study the impact of leveraging the \mathbb{G} -equivariance of the closed-loop transfer operator $E_{\vec{o}|\vec{o}}^{\Delta t}$ when approximating it with a finite-dimensional linear model.

Experimental validation

We present two experiments designed to quantify the impact of leveraging symmetry priors when approximating a \mathbb{G} -equivariant closed-loop transfer operator $E_{\vec{o}|\vec{o}}^{\Delta t}$ in finite dimensions, using a nonlinear latent state representation ψ_θ and a linear matrix A_θ . In all experiments, the truncated operator E_θ is learned by minimizing the multi-step dynamics autoencoder loss (Lusch et al., 2018; Ordoñez-Appraez et al., 2024):

$$\mathcal{L}_{\text{dae}}(A_\theta, H_{\text{pred}}, \gamma) = \sum_{h=0}^{H_{\text{pred}}} \|\mathbf{o}_{t+h\Delta t} - \mathbf{d}_\theta(A_\theta^h \psi_\theta(\vec{o}_t))\|^2 + \gamma \|\psi_\theta(\vec{o}_{t+h\Delta t}) - A_\theta^h \psi_\theta(\vec{o}_t)\|^2$$

This loss, balanced by γ , penalizes both the prediction error of state observables $\{\mathbf{o}_{t+h\Delta t}\}_{h=0}^{H_{\text{pred}}}$ and the prediction error of latent state representations $\{\psi_\theta(\vec{o}_{t+h\Delta t})\}_{h=0}^{H_{\text{pred}}}$ over a prediction horizon of H_{pred} steps. Multi-step predictions are obtained by powering the matrix representation of the truncated operator, with a nonlinear decoder \mathbf{d}_θ .

Remark: Section 9.3.2

The experiments in this section are adapted from Ordoñez-Appraez et al. (2024), a publication that preceded the development of the contrastive representation learning framework in Kostic et al. (2024); Ordoñez-Appraez et al. (2026) (see section 9.2), which introduces a contrastive loss for operator approximation in finite dimensions. A comparison between these approaches is left for future work, but it is worth noting that the contrastive training framework provides statistical learning guarantees that the dynamics autoencoder loss does not.

We compare three methods: Dynamics Autoencoder (DAE), Dynamics Autoencoder with data augmentation (DAEaug), and Equivariant Dynamics Autoencoder (eDAE), which differ in their use of standard or \mathbb{G} -equivariant latent state representation functions $\psi_\theta : \mathcal{O}^H \rightarrow \mathbb{R}^{r_s}$, linear maps A_θ , and decoders $\mathbf{d}_\theta : \mathbb{R}^{r_s} \rightarrow \mathcal{O}$. For all experiments, encoder and decoder NN architectures are matched in size, and we report the mean, min and max over four random training seeds.




Figure 9.10 Test set prediction MSE of DAE, DAE-Aug, and eDAE for closed-loop locomotion dynamics modelling of the Mini-Cheetah robot, comparing symmetry groups $\mathbb{G} = \mathbb{K}_4 \times \mathbb{C}_2$ and $\mathbb{G} = \mathbb{K}_4$, while varying training dataset size, prediction horizon, and latent model space dimension r_s . Solid lines and shaded areas represent the mean, maximum, and minimum prediction error among 4 training seeds. (a) Sample efficiency plot comparing test set MSE vs. size of the training dataset. (b) Models MSE while varying prediction horizon. (c) Models' MSE while varying dimensionality of the latent state representation r_s . (d) MSE of measurable state observables in original units. Figure adapted from Ordoñez-Appraez et al. (2024).

Synthetic symmetric dynamical systems with finite symmetry groups In this experiment, we generate synthetic nonlinear \mathbb{G} -symmetric dynamical systems with various finite symmetry groups \mathbb{G} . Each system is defined as a constrained stable linear system with stochastic perturbations:

$$\mathbf{o}_{t+\Delta t} = \mathbf{K}_{\Delta t} \mathbf{o}_t + \boldsymbol{\epsilon}_t, \quad \text{s.t. } \mathbf{C} \mathbf{o}_t \geq \mathbf{c}, \quad \forall \mathbf{o}_t \in \mathcal{O}, t \in \mathbb{R},$$

where $\mathbf{o} \in \mathcal{O}$ is a state observation, $\mathbf{K}_{\Delta t} \in \mathbb{R}^{|\mathcal{O}| \times |\mathcal{O}|}$ is a linear dynamics matrix, $\boldsymbol{\epsilon}_t \in \mathbb{R}^{|\mathcal{O}|}$ is a white-noise stochastic process with standard deviation σ , and $\mathbf{C} \in \mathbb{R}^{n_c \times |\mathcal{O}|}$, $\mathbf{c} \in \mathbb{R}^{n_c}$ are the parameters of n_c inequality hyperplane constraints. A system is \mathbb{G} -symmetric if $\mathbf{A}_{\Delta t}$ is a \mathbb{G} -equivariant matrix, i.e., $g \triangleright_{\mathcal{O}} \mathbf{A}_{\Delta t} \cdot = \mathbf{A}_{\Delta t} (g \triangleright_{\mathcal{O}} \cdot)$ for all $g \in \mathbb{G}$, and every hyperplane constraint is added in group orbits (see Ordoñez-Appraez et al. (2024) for details).

For these systems, $\mathbb{P}_{\sigma}^{\Delta t}$ is Gaussian, so we set the history horizon to $H = 1$, i.e., $\vec{\mathbf{o}}_t = \mathbf{o}_t$, and choose the discretization step Δt to ensure a reasonable signal-to-noise ratio.

These synthetic systems enable a systematic comparison of the impact of symmetry exploitation across different groups \mathbb{G} , observation space dimensions $|\mathcal{O}|$, latent state dimensions r_s , and noise levels σ . The results show that the eDAE architecture consistently outperforms the symmetry-agnostic model in terms of sample efficiency (fig. 9.9-a), compounding multi-step prediction error (fig. 9.9-b), sensitivity to the choice of latent state dimension r_s (fig. 9.9-c), and robustness to variations in the noise scale of the underlying dynamics (fig. 9.9-d).

Modelling quadruped closed-loop dynamics In this experiment, we learn a finite-dimensional approximation of the closed-loop transfer operator $\mathbb{E}_{\sigma}^{\Delta t}$ for quadruped legged locomotion, induced by a fixed \mathbb{G} -invariant stochastic control policy π that tracks a desired base velocity with a fixed trotting gait (Amatucci et al., 2022). The robot traverses uneven terrain, resulting in a distribution of disturbances that generate the policy's closed-loop invariant set \mathbb{W}_{π} (see definition 9.2). The time-invariant marginal distribution of state observation histories \mathbb{P}_{σ} represents the distribution of trajectories under nominal policy performance, empirically approximated by a \mathbb{G} -invariant dataset of successful episode recordings $\mathbb{D} = \{\vec{\mathbf{o}}_i\}_{i=1}^N$ composed of locomotion trajectories from the 8 different symmetric modes of the robot; see animation.

The Mini-Cheetah quadruped robot possesses a symmetry group of order 8: $\mathbb{G} = \mathbb{C}_2 \times \mathbb{C}_2 \times \mathbb{C}_2 \equiv \mathbb{K}_4 \times \mathbb{C}_2$ (see fig. 9.8). The state observation is defined as

$$\mathbf{o}_t = (\mathbf{q}_{\text{js}}, \dot{\mathbf{q}}_{\text{js}}, z_t, \mathbf{o}_t, \mathbf{v}_{\text{err},t}, \mathbf{w}_{\text{err},t}) \in \mathcal{O} \subseteq \mathbb{R}^{37}.$$

where $\mathbf{q}_{\text{js}} \in \mathcal{M} \subseteq \mathbb{R}^{12}$ and $\dot{\mathbf{q}}_{\text{js}} \in \mathcal{T}_q \mathcal{M} \subseteq \mathbb{R}^{12}$ are the joint-space generalized position and velocity coordinates, $z_t \in \mathbb{R}^1$ is the floating-base height, $\mathbf{R}_B \in \mathbb{SO}_3$ is the floating-base orientation, and $\mathbf{v}_{\text{err},t} \in \mathbb{R}^3$, $\mathbf{w}_{\text{err},t} \in \mathbb{R}^3$ are the control errors in the desired linear and angular base velocities, respectively. For this task, a history horizon of $H = 5$ is sufficient to disambiguate between the robot's different contact modes during normal locomotion.

The results show that the eDAE model consistently outperforms DAEaug and DAE in terms of sample efficiency (fig. 9.10-a), compounding multi-step prediction error (fig. 9.10-b), sensitivity to the choice of latent state dimension r_s (fig. 9.10-c), and overall observation prediction error (fig. 9.10-d). Notably, when comparing eDAE and DAEaug models trained with either the full symmetry group $\mathbb{K}_4 \times \mathbb{C}_2$ or the subgroup \mathbb{K}_4 (see fig. 7.1), we observe significant differences in sample efficiency and optimality gaps, highlighting the increasing benefits of leveraging larger symmetry groups to improve learned dynamics model performance.

Chapter 10

Discussion and Future Work

This thesis presents a theoretical and practical study of the symmetries of robot dynamics and their implications for the core learning tasks in robotics: dynamics modelling, state estimation, and control. In contrast to previous works on symmetries in robotics (Bloch et al., 1996; Murray et al., 2017; Wieber, 2006), which rely heavily on differential geometry and focus primarily on environmental/base symmetries of mechanical systems (see fig. 1.2), this thesis adopts a predominantly group-theoretic analysis of Lagrangian mechanics.

This perspective enables us to extend the notion of symmetry to a broader class of transformations, most notably morphological symmetries, which constitute the core focus of this thesis. It also reveals a new and largely unexplored connection between robot dynamics and harmonic analysis (Mackey, 1980), a field that provides a broad range of theoretical results and algorithms for the analysis of symmetric dynamical systems. While harmonic analysis has played a central role in modern particle physics and chemistry (Dresselhaus et al., 2007), its implications for robot modelling and control remain largely unexplored, opening a promising research direction for robotics (see section 4.3 and chapter 9).

Through substantial theoretical and empirical work, this thesis argues that environmental/base symmetries, and more importantly morphological symmetries, should be exploited as physics-informed priors to improve the efficiency and generalization of learning algorithms for modelling, state estimation, and control in robotics. Experimental results across these three core learning tasks consistently demonstrate that symmetry-aware learning algorithms significantly outperform their symmetry-agnostic counterparts, while the open-access software stack accompanying this thesis provides a practical toolkit for the adoption of these techniques in robotics research and applications.

Further work

Throughout this PhD project, we opened far more questions than the ones we closed. Below, I outline some of the core research directions that I am most excited to pursue next.

Computational design of robot morphologies The interplay of learning and morphology is an open field for research in symmetry-aware computational design. Large-scale locomotion learning experiments leveraging the bilateral reflection symmetry present in most animals in the phylogenetic

tree were performed by [Gupta et al. \(2021\)](#), opening a research line that connects learning difficulty to the morphology of the agent learning to solve the task. Throughout this thesis, we identified how the presence of morphological symmetries plays a crucial role in mitigating the curse of dimensionality ([Higgins et al., 2022](#)) and improving learning rates ([Ordoñez-Præez et al., 2026](#); [Wang et al., 2022b](#))—that is, in reducing the learning difficulty. While the majority of the industry is focused on producing robots with a humanoid form factor, I see no reason why we should not explore a completely different space of robot morphologies than those seen in nature. Especially because robot manufacturing possesses two key technologies that nature never had access to: (i) unbounded revolute joints, enabling not only wheels but also body motions without position constraints, and (ii) reproducibility and low cost of mass manufacturing. Unbounded revolute joints open a completely different design space of morphologies than those seen in nature, and reproducibility removes the key evolutionary pressure that guided morphology design in nature: resource efficiency. To guide the exploration of this new design space of artificial embodied intelligence, morphological symmetries provide the ideal mathematical language to connect morphology to learning and control tractability.

Transfer operator based control The original aim of this thesis’s work on Koopman/transfer operators was to address model-based predictive control in robotics. In practice, however, we were only able to tackle closed-loop dynamics modeling, due to two main obstacles: (i) the theory for approximating these operators from data was not mature enough when this project began, and (ii) practical methods for learning these operators offered little to no learning guarantees and performed poorly in high-dimensional, contact-rich settings. In recent years, substantial progress has been made on both fronts, with ([Kostic et al., 2024](#); [Ordoñez-Præez et al., 2026](#)) introducing learning algorithms that are now ready to be tested in control contexts. Notably, interpreting the low-rank approximation of the spectral decomposition of the transfer operator opens a promising avenue for future research, bridging operator approximation with the statistical learning theory of control and exploration [Tsiamis et al. \(2023\)](#). Furthermore, as discussed in [chapter 9](#), environmental and morphological symmetries play a pivotal role in both the approximation of the transfer operator and the decomposition of control problems into lower-dimensional subproblems, highlighting the substantial computational benefits of symmetry-aware control approaches.

References

- Farzad Abdolhosseini, Hung Yu Ling, Zhaoming Xie, Xue Bin Peng, and Michiel Van de Panne. On learning symmetric locomotion. In *Motion, Interaction and Games*, pages 1–10. 2019.
- Lorenzo Amatucci, Joon-Ha Kim, Jemin Hwangbo, and Hae-Won Park. Monte carlo tree search gait planner for non-gaited legged system control. In *International Conference on Robotics and Automation (ICRA)*, 2022.
- Charles R Baker. Joint measures and cross-covariance operators. *Trans. Am. Math. Soc.*, 186:273–273, 1973.
- Kaixi Bao, Chenhao Li, Yarden As, Andreas Krause, and Marco Hutter. Toward task generalization via memory augmentation in meta-reinforcement learning. *arXiv preprint arXiv:2502.01521*, 2025.
- Axel Barrau and Silvere Bonnabel. Invariant kalman filtering. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):237–257, 2018.
- Petar Bevanda, Max Beier, Shahab Heshmati-Alamdari, Stefan Sosnowski, and Sandra Hirche. Towards data-driven lqr with koopmanizing flows. *IFAC-PapersOnLine*, 55(15):13–18, 2022.
- Gerardo Bledt, Patrick M Wensing, Sam Ingersoll, and Sangbae Kim. Contact model fusion for event-based locomotion in unstructured terrains. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4399–4406. IEEE, 2018.
- Anthony M Bloch, Perinkulam S Krishnaprasad, Jerrold E Marsden, and Richard M Murray. Non-holonomic mechanical systems with symmetry. *Archive for rational mechanics and analysis*, 136(1):21–99, 1996.
- Johann Brehmer, Joey Bose, Pim de Haan, and Taco Cohen. Edgi: equivariant diffusion for planning with embodied agents. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23'*, pages 1–17, 2024.
- Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- Steven L Brunton, Marko Budišić, Eurika Kaiser, and J Nathan Kutz. Modern koopman theory for dynamical systems. *arXiv preprint arXiv:2102.12086*, 2021.
- Marco Camurri, Maurice Fallon, Stéphane Bazeille, Andreea Radulescu, Victor Barasuol, Darwin G Caldwell, and Claudio Semini. Probabilistic contact estimation and impact detection for state estimation of quadruped robots. *IEEE Robotics and Automation Letters*, 2(2):1023–1030, 2017.
- Justin Carpentier, Guilhem Saurel, Gabriele Buondonno, Joseph Mirabel, Florent Lamiroux, Olivier Stasse, and Nicolas Mansard. The pinocchio c++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives. In *2019 IEEE/SICE International Symposium on System Integration (SII)*, pages 614–619. IEEE, 2019.

- Rafael Cathomen, Mayank Mittal, Marin Vlastelica, and Marco Hutter. Divide, discover, deploy: Factorized skill learning with symmetry and style priors. *arXiv preprint arXiv:2508.19953*, 2025.
- Gabriele Cesa, Leon Lang, and Maurice Weiler. A program to build e (n)-equivariant steerable cnns. In *International Conference on Learning Representations*, 2021.
- Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research (IJRR)*, 2024.
- Taco S Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous spaces. *Advances in neural information processing systems*, 32, 2019.
- Fritz Colonius, Wolfgang Kliemann, and Lars Grüne. *The dynamics of control*, volume 202. Springer, 2000.
- Thomas Corbères, Carlos Mastalli, Wolfgang Merkt, Ioannis Havoutis, Maurice Fallon, Nicolas Mansard, Thomas Flayols, Sethu Vijayakumar, and Steve Tonneau. Perceptive locomotion through whole-body mpc and optimal region selection, 2024.
- John F Cornwell. *Group theory in physics: An introduction*. Academic press, 1997.
- Jiayu Ding and Zhenyu Gan. Breaking symmetries leads to diverse quadrupedal gaits. *IEEE Robotics and Automation Letters*, 2024.
- John D Dixon. Computing irreducible representations of groups. *Mathematics of Computation*, 24 (111):707–712, 1970.
- Mildred S Dresselhaus, Gene Dresselhaus, and Ado Jorio. *Group theory: application to the physics of condensed matter*. Springer Science & Business Media, 2007.
- Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- Roy Featherstone. *Rigid body dynamics algorithms*. Springer, 2008.
- Shai Feldman, Stephen Bates, and Yaniv Romano. Calibrated multiple-output quantile regression with representation learning. *Journal of Machine Learning Research*, 24(24):1–48, 2023.
- Marc Finzi, Gregory Benton, and Andrew G Wilson. Residual pathway priors for soft equivariance constraints. *Advances in Neural Information Processing Systems*, 34:30037–30049, 2021.
- Maxime Gautier. Dynamic identification of robots with power model. In *Proceedings of international conference on robotics and automation*, volume 3, pages 1922–1927. IEEE, 1997.
- Mario Geiger, Tess Smidt, Alby M., Benjamin Kurt Miller, Wouter Boomsma, Bradley Dice, Kostiantyn Lapchevskyi, Maurice Weiler, Michał Tyszkiewicz, Simon Batzner, Dylan Madisetti, Martin Uhrin, Jes Frellesen, Nuri Jung, Sophia Sanborn, Mingjian Wen, Josh Rackers, Marcel Rød, and Michael Bailey. Euclidean neural networks: e3nn, April 2022. URL <https://doi.org/10.5281/zenodo.6459381>.

- Maani Ghaffari, Ray Zhang, Minghan Zhu, Chien Erh Lin, Tzu-Yuan Lin, Sangli Teng, Tingjun Li, Tianyi Liu, and Jingwei Song. Progress in symmetry preserving robot perception and control through geometry and learning. *Frontiers in Robotics and AI*, 9:969380, 2022.
- Nicolas Gilardi, Samy Bengio, and Mikhail Kanevski. Conditional gaussian mixture models for environmental risk mapping. In *Proceedings of the 12th IEEE workshop on neural networks for signal processing*, pages 777–786. IEEE, 2002.
- Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. A survey of deep learning techniques for autonomous driving. *Journal of field robotics*, 37(3):362–386, 2020.
- F. Grimmering, A. Meduri, M. Khadiv, J. Viereck, M. Wüthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols, J. Fiene, A. Badri-Spröwitz, and L. Righetti. An open torque-controlled modular robot architecture for legged locomotion research. *IEEE Robotics and Automation Letters*, 5(2):3650–3657, 2020. doi: 10.1109/LRA.2020.2976639.
- Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via learning and evolution. *Nature communications*, 12(1):5721, 2021.
- Mustafa Hajj, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, Karthikeyan Natesan Ramamurthy, Tolga Birdal, Tamal K Dey, Soham Mukherjee, Shreyas N Samaga, et al. Topological deep learning: Going beyond graph data. *arXiv preprint arXiv:2206.00606*, 2022.
- Jeff Z. HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. In *Advances in Neural Information Processing Systems*, volume 34, pages 5000–5011. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/27debb435021eb68b3965290b5e24c49-Paper.pdf.
- Ross Hartley, Maani Ghaffari, Ryan M Eustice, and Jessy W Grizzle. Contact-aided invariant extended kalman filtering for robot state estimation. *The International Journal of Robotics Research*, 39(4):402–430, 2020.
- Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner. Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*, 2018.
- Irina Higgins, Sébastien Racanière, and Danilo Rezende. Symmetry-based representations for artificial and biological general intelligence. *Frontiers in Computational Neuroscience*, page 28, 2022.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Gábor Holló. Demystification of animal symmetry: Symmetry is a response to mechanical forces. *Biology Direct*, 12(1):11, 2017.
- Rafael Izbicki and Ann B. Lee. Converting high-dimensional regression to high-dimensional conditional density estimation. 2017.
- John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- Leslie Pack Kaelbling. The foundation of efficient robot learning. *Science*, 369(6506):915–916, 2020.

- Benjamin Katz, Jared Di Carlo, and Sangbae Kim. Mini cheetah: A platform for pushing the limits of dynamic quadruped control. In *2019 international conference on robotics and automation (ICRA)*, pages 6295–6301. IEEE, 2019.
- Seungyeon Kim, Byeongdo Lim, Yonghyeon Lee, and Frank C Park. Se (2)-equivariant pushing dynamics models for tabletop object manipulations. In *Conference on Robot Learning*, pages 427–436. PMLR, 2023.
- Leon Klein, Andreas Krämer, and Frank Noé. Equivariant flow matching. *Advances in Neural Information Processing Systems*, 36:59886–59910, 2023.
- Anthony W. Knap. *Representation Theory of Semisimple Groups, An Overview Based on Examples (PMS-36)*. Princeton University Press, Princeton, 1986.
- Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International conference on machine learning*, pages 2747–2755. PMLR, 2018.
- Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch-gordan nets: a fully fourier space spherical convolutional neural network, 2018. URL <https://arxiv.org/abs/1806.09231>.
- Vladimir Kostic, Gregoire Pacreau, Giacomo Turri, Pietro Novelli, Karim Lounici, and Massimiliano Pontil. Neural conditional probability for uncertainty quantification. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 60999–61039. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/705b97ecb07ae86524d438abac97a3e2-Paper-Conference.pdf.
- Vladimir R Kostic, Pietro Novelli, Riccardo Grazi, Karim Lounici, and Massimiliano Pontil. Learning invariant representations of time-homogeneous stochastic dynamical systems. *arXiv preprint arXiv:2307.09912*, 2023.
- Vikash Kumar, Abhishek Gupta, Emanuel Todorov, and Sergey Levine. Learning dexterous manipulation policies from experience and imitation. *arXiv preprint arXiv:1611.05095*, 2016.
- Cornelius Lanczos. *The variational principles of mechanics*. University of Toronto press, 2020.
- Leon Lang and Maurice Weiler. A wigner-eckart theorem for group equivariant convolution kernels. *arXiv preprint arXiv:2010.10952*, 2020.
- Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020.
- Zechu Li, Yufeng Jin, Daniel Ordoñez-Apaez, Claudio Semini, Puze Liu, and Georgia Chalvatzaki. Morphologically symmetric reinforcement learning for ambidextrous bimanual manipulation. In *The Conference on Robot Learning (CoRL)*. IEEE, 2025.
- Chien Erh Lin, Minghan Zhu, and Maani Ghaffari. Se3et: Se (3)-equivariant transformer for low-overlap point cloud registration. *IEEE Robotics and Automation Letters*, 2024.
- Tzu-Yuan Lin, Ray Zhang, Justin Yu, and Maani Ghaffari. Legged robot state estimation using invariant kalman filtering and learned contact events. In *5th Annual Conference on Robot Learning*, 2021.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modelling. *arXiv preprint arXiv:2210.02747*, 2022.

- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6): 248:1–248:16, October 2015.
- Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications*, 9, 2018.
- George W. Mackey. Harmonic analysis as the exploitation of symmetry—a historical survey. *Bulletin (New Series) of the American Mathematical Society*, 3(1.P1):543 – 698, 1980.
- Michael Maravgakis, Despina-Ekaterini Argiropoulos, Stylianos Piperakis, and Panos Trahanias. Probabilistic contact state estimation for legged robots using inertial information. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12163–12169. IEEE, 2023.
- Sergi Martinez, Steve Tonneau, and Carlos Mastalli. System identification under constraints and disturbance: A bayesian estimation approach. *arXiv preprint arXiv:2602.16358*, 2026.
- Carlos Mastalli, Rohan Budhiraja, Wolfgang Merkt, Guilhem Saurel, Bilal Hammoud, Maximilien Naveau, Justin Carpentier, Ludovic Righetti, Sethu Vijayakumar, and Nicolas Mansard. Crocodyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020a.
- Carlos Mastalli, Ioannis Havoutis, Michele Focchi, Darwin G. Caldwell, and Claudio Semini. Motion Planning for Quadrupedal Locomotion: Coupled Planning, Terrain Mapping, and Whole-Body Control. *IEEE Transactions on Robotics*, 36, 2020b.
- Alexandre Mauroy, Y Susuki, and Igor Mezic. *Koopman operator in systems and control*. Springer, 2020.
- Hrishik Mishra. *Employing Symmetry in Dynamics and Motion Control of Robotic Mechanisms*. PhD thesis, Technische Universität Wien, 2025.
- Mayank Mittal, Nikita Rudin, Victor Klemm, Arthur Allshire, and Marco Hutter. Symmetry considerations for learning task symmetric robot policies. *arXiv preprint arXiv:2403.04359*, 2024.
- Arnab Kumar Mondal, Vineet Jain, Kaleem Siddiqi, and Siamak Ravanbakhsh. Eqr: Equivariant representations for data-efficient reinforcement learning. In *International Conference on Machine Learning*, pages 15908–15926. PMLR, 2022.
- Richard M Murray, Zexiang Li, and S Shankar Sastry. *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- Thomas Nagler and Claudia Czado. Evading the curse of dimensionality in nonparametric density estimation with simplified vine copulas. *Journal of Multivariate Analysis*, 151:69–89, October 2016. URL <http://dx.doi.org/10.1016/j.jmva.2016.07.003>.
- Ylenia Nisticò, João Carlos Virgolino Soares, Lorenzo Amatucci, Geoff Fink, and Claudio Semini. Muse: A real-time multi-sensor state estimator for quadruped robots. *IEEE Robotics and Automation Letters*, 2025.
- Frank Noé, Alexandre Tkatchenko, Klaus-Robert Müller, and Cecilia Clementi. Machine learning for molecular simulation. *Annual review of physical chemistry*, 71:361–390, 2020.
- Emmy Noether. Invariante variationsprobleme. In *Gesammelte Abhandlungen-Collected Papers*, pages 231–239. Springer, 1983.

- Oskar Nordensfors, Fredrik Ohlsson, and Axel Flinth. Optimization dynamics of equivariant and augmented neural networks. *arXiv preprint arXiv:2303.13458*, 2023.
- Pietro Novelli, Marco Praticò, Massimiliano Pontil, and Carlo Ciliberto. Operator world models for reinforcement learning. *Advances in Neural Information Processing Systems*, 37:111432–111463, 2024.
- Daniel Ordoñez-Apraez, Antonio Agudo, Francesc Moreno-Noguer, and Mario Martin. An adaptable approach to learn realistic legged locomotion without examples. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 4671–4678. IEEE, 2022.
- David E Orin, Ambarish Goswami, and Sung-Hee Lee. Centroidal dynamics of a humanoid robot. *Autonomous robots*, 35(2):161–176, 2013.
- Chaoyi Pan, Giri Anantharaman, Nai-Chieh Huang, Claire Jin, Daniel Pfrommer, Chenyang Yuan, Frank Permenter, Guannan Qu, Nicholas Boffi, Guanya Shi, et al. Much ado about noising: Dispelling the myths of generative robotic control. *arXiv preprint arXiv:2512.01809*, 2025.
- Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Stefanos Pertigkiozoglou, Evangelos Chatzipantazis, Shubhendu Trivedi, and Kostas Daniilidis. Improving equivariant model training via constraint relaxation. *Advances in Neural Information Processing Systems*, 37:83497–83520, 2024.
- Stefanos Pertigkiozoglou, Mircea Petrache, Shubhendu Trivedi, and Kostas Daniilidis. Recurrent equivariant constraint modulation: Learning per-layer symmetry relaxation from data. *arXiv preprint arXiv:2602.02853*, 2026.
- Henri Poincaré. The principles of mathematical physics. *The Monist*, pages 1–24, 1905.
- Balaraman Ravindran and Andrew G Barto. Symmetries and model minimization in markov decision processes, 2001.
- Michael Reed. *Methods of modern mathematical physics: Functional analysis*. Elsevier, 2012.
- Sahand Rezaei-Shoshtari, Rosie Zhao, Prakash Panangaden, David Meger, and Doina Precup. Continuous mdp homomorphisms and homomorphic policy gradient. *Advances in Neural Information Processing Systems*, 35:20189–20204, 2022.
- Nicholas Roy, Ingmar Posner, Tim Barfoot, Philippe Beaudoin, Yoshua Bengio, Jeannette Bohg, Oliver Brock, Isabelle DePATIE, Dieter Fox, Dan Koditschek, et al. From machine learning to robotics: Challenges and opportunities for embodied intelligence. *arXiv preprint arXiv:2110.15245*, 2021.
- Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Conference on Robot Learning*, pages 91–100. PMLR, 2022.

- J Jon Ryu, Xiangxiang Xu, HS Erol, Yuheng Bu, Lizhong Zheng, and Gregory W Wornell. Operator svd with neural networks via nested low-rank approximation. *arXiv preprint arXiv:2402.03655*, 2024.
- Anastasiya Salova, Jeffrey Emenheiser, Adam Rupe, James P Crutchfield, and Raissa M D’Souza. Koopman operator and its approximations for systems with symmetries. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(9), 2019.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- DAVID W. Scott. Feasibility of multivariate density estimates. *Biometrika*, 78(1):197–205, 1991. URL <http://dx.doi.org/10.1093/biomet/78.1.197>.
- Jon M Selig. Lie groups and lie algebras in robotics. In *Computational Noncommutative Algebra and Applications*, pages 101–125. Springer, 2004.
- Parikshit Shah and Venkat Chandrasekaran. Group symmetry and covariance regularization. In *2012 46th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2012.
- Max Siebenborn, Daniel Ordoñez-Apraez, Sophie Lueth, Giulio Turrisi, Massimiliano Pontil, Claudio Semini, and Georgia Chalvatzaki. Morphologically equivariant flow matching for bimanual mobile manipulation. *arXiv preprint arXiv:2605.12228*, 2026. doi: 10.48550/arXiv.2605.12228. URL <https://arxiv.org/abs/2605.12228>.
- Joan Sola, Jeremie Deray, and Dinesh Atchuthan. A micro lie theory for state estimation in robotics. *arXiv preprint arXiv:1812.01537*, 2018.
- Benjamin Steinberg. *Representation theory of finite groups: an introductory approach*, volume 68. Springer, 2012.
- Maojiang Su, Jerry Yao-Chieh Hu, Sophia Pi, and Han Liu. On flow matching kl divergence. *arXiv preprint arXiv:2511.05480*, 2025.
- Zhi Su, Xiaoyu Huang, Daniel Ordoñez-Apraez, Yunfei Li, Zhongyu Li, Qiayuan Liao, Giulio Turrisi, Massimiliano Pontil, Claudio Semini, Yi Wu, et al. Leveraging symmetry in rl-based legged locomotion control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, et al. The limits and potentials of deep learning for robotics. *The International journal of robotics research*, 37(4-5):405–420, 2018.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza, Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*, 2025.
- Daniel Ordoñez-Apraez, Mario Martin, Antonio Agudo, and Francesc Moreno-Noguer. On discrete symmetries of robotics systems: A group-theoretic and data-driven analysis. In *Proceedings of the 19th Robotics: Science and Systems Conference (RSS)*, Daegu, Republic of Korea, 2023. URL <https://www.roboticsproceedings.org/rss19/index.html>.

- Daniel Ordoñez-Apaez, Vladimir Kostic, Giulio Turrisi, Pietro Novelli, Carlos Mastalli, Claudio Semini, and Massimiliano Pontil. Dynamics harmonic analysis of robotic systems: Application in data-driven koopman modelling. In *6th Annual Learning for Dynamics & Control Conference*, pages 1318–1329. PMLR, 2024.
- Daniel Ordoñez-Apaez, Giulio Turrisi, Vladimir Kostic, Mario Martin, Antonio Agudo, Francesc Moreno-Noguer, Massimiliano Pontil, Claudio Semini, and Carlos Mastalli. Morphological symmetries in robotics. *The International Journal of Robotics Research (IJRR)*, 2025.
- Daniel Ordoñez-Apaez, Vladimir Kostić, Alek Fröhlich, Vivien Brandt, Karim Lounici, and Massimiliano Pontil. Representation learning for equivariant inference with guarantees. In *The International Conference on Learning Representations*, 2026.
- Evangelos Theodorou, Yuval Tassa, and Emo Todorov. Stochastic differential dynamic programming. In *Proceedings of the 2010 American Control Conference*, pages 1125–1132. IEEE, 2010.
- Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.
- Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents. MIT Press, Cambridge, MA, 2005. ISBN 978-0-262-20162-9.
- Franki Nguimatsia Tiofack, Théotime Le Hellard, Fabian Schramm, Nicolas Perrin-Gilbert, and Justin Carpentier. Guided flow policy: Learning from high-value actions in offline reinforcement learning. *arXiv preprint arXiv:2512.03973*, 2025.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132 of *Proceedings of Machine Learning Research*, pages 1179–1206. PMLR, 16–19 Mar 2021.
- Yao-Hung Hubert Tsai, Han Zhao, Makoto Yamada, Louis-Philippe Morency, and Russ R Salakhutdinov. Neural methods for point-wise dependency estimation. In *Advances in Neural Information Processing Systems*, volume 33, pages 62–72. Curran Associates, Inc., 2020.
- Anastasios Tsiamis, Ingvar Ziemann, Nikolai Matni, and George J Pappas. Statistical learning theory for control: A finite-sample perspective. *IEEE Control Systems Magazine*, 43(6):67–97, 2023.
- Giacomo Turri, Luigi Bonati, Kai Zhu, Massimiliano Pontil, and Pietro Novelli. Self-supervised evolution operator learning for high-dimensional dynamical systems. *arXiv preprint arXiv:2505.18671*, 2025.
- Elise van der Pol, Daniel Worrall, Herke van Hoof, Frans Oliehoek, and Max Welling. Mdp homomorphic networks: Group symmetries in reinforcement learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4199–4210. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/2be5f9c2e3620eb73c2972d7552b6cb5-Paper.pdf.
- Dian Wang, Robin Walters, and Robert Platt. $SO(2)$ -equivariant reinforcement learning. In *International Conference on Learning Representations*, 2022a. URL https://openreview.net/forum?id=7F9cOhdvfk_.

- Dian Wang, Mingxi Jia, Xupeng Zhu, Robin Walters, and Robert Platt. On-robot learning with equivariant models. In Karen Liu, Dana Kulic, and Jeff Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 1345–1354. PMLR, 14–18 Dec 2023a. URL <https://proceedings.mlr.press/v205/wang23c.html>.
- Dian Wang, Xupeng Zhu, Jung Yeon Park, Mingxi Jia, Guanang Su, Robert Platt, and Robin Walters. A general theory of correct, incorrect, and extrinsic equivariance. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 40006–40029. Curran Associates, Inc., 2023b. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/7dc7793c89b93887e126a86f22ef63c6-Paper-Conference.pdf.
- Dian Wang, Stephen Hart, David Surovik, Tarik Kelestemur, Haojie Huang, Haibo Zhao, Mark Yeatman, Jiuguang Wang, Robin Walters, and Robert Platt. Equivariant diffusion policy. *arXiv preprint arXiv:2407.01812*, 2024.
- Dian Wang, Boce Hu, Shuran Song, Robin Walters, and Robert Platt. A practical guide for incorporating symmetry in diffusion policy. *Advances in Neural Information Processing Systems*, 38: 94105–94130, 2026.
- Rui Wang, Robin Walters, and Rose Yu. Data augmentation vs. equivariant networks: A theory of generalization on dynamics forecasting. *arXiv preprint arXiv:2206.09450*, 2022b.
- Ziyu Wang, Yucen Luo, Yueru Li, Jun Zhu, and Bernhard Schölkopf. Spectral representation learning for conditional moment models. *arXiv preprint arXiv:2210.16525*, 2022c.
- Joachim Weidmann. *Linear operators in Hilbert spaces*, volume 68. Springer Science & Business Media, 2012.
- Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. Coordinate independent convolutional networks—*isometry and gauge equivariant convolutions on riemannian manifolds*. *arXiv preprint arXiv:2106.06020*, 2021.
- Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. *Equivariant and Coordinate Independent Convolutional Networks*. 2023. URL https://maurice-weiler.gitlab.io/cnn_book/EquivariantAndCoordinateIndependentCNNs.pdf.
- James T. Wheeler. General coordinate covariance of the euler lagrange equations. *Classical Mechanics class notes*, 2014. URL <http://www.physics.usu.edu/Wheeler/ClassicalMechanics/CMCoordinateinvarianceofEulerLagrange.pdf>.
- Pierre-Brice Wieber. Holonomy and nonholonomy in the dynamics of articulated motion. In *Fast Motions in Biomechanics and Robotics: Optimization and Feedback Control*, pages 411–425. Springer, 2006.
- Samuel Willingham, Mårten Sjöström, and Christine Guillemot. Maximum a posteriori training of diffusion models for image restoration. In *2025 33rd European Signal Processing Conference (EUSIPCO)*, pages 1882–1886. IEEE, 2025.
- Xiaomi. Cyberdog2. <https://www.mi.com/cyberdog2>, 2024. Accessed: Mar. 2024.
- Fengze Xie, Sizhe Wei, Yue Song, Yisong Yue, and Lu Gan. Morphological-symmetry-equivariant heterogeneous graph neural network for robotic dynamics learning. *arXiv preprint arXiv:2412.01297*, 2024.
- YuQing Xie and Tess Smidt. A tale of two symmetries: Exploring the loss landscape of equivariant models. *arXiv preprint arXiv:2506.02269*, 2025.

- Jingyun Yang, Zi-ang Cao, Congyue Deng, Rika Antonova, Shuran Song, and Jeannette Bohg. Equibot: Sim (3)-equivariant diffusion policy for generalizable and data efficient learning. *arXiv preprint arXiv:2407.01479*, 2024.
- Raymond Yeh, Yuan-Ting Hu, and Alexander Schwing. Chirality nets for human pose regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- Wenhao Yu, Greg Turk, and C Karen Liu. Learning symmetric and low-energy locomotion. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018.
- Kevin Zakka. Mink: Python inverse kinematics based on MuJoCo, May 2025. URL <https://github.com/kevinzakka/mink>.
- Ray Zhang, Zheming Zhou, Min Sun, Omid Ghasemalizadeh, Cheng-Hao Kuo, Ryan M Eustice, Maani Ghaffari, and Arnie Sen. Correspondence-free se (3) point cloud registration in rkhs via unsupervised equivariant learning. In *European Conference on Computer Vision*, pages 68–86. Springer, 2024.
- Zhiyuan Zhang, Aditya Mohan, Seunggho Han, Wan Shou, Dongyi Wang, and Yu She. Equibim: Learning symmetry-equivariant policy for bimanual manipulation. *arXiv preprint arXiv:2603.08541*, 2026.
- Xupeng Zhu, Dian Wang, Guanang Su, Ondrej Biza, Robin Walters, and Robert Platt. On robot grasp learning using equivariant models. *Autonomous Robots*, 47(8):1175–1193, 2023.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.
- Martin Zinkevich and Tucker Balch. Symmetry in markov decision processes and its implications for single agent and multi agent learning. In *Proceedings of the 18th International Conference on Machine Learning*. Citeseer, 2001.

I

Glossary

cCDF conditional Cumulative Distribution Function.

cGMM Conditional Gaussian Mixture Model (Gilardi et al., 2002): A parametric model for benchmark conditional density estimation datasets. Generates random variables \mathbf{x} and \mathbf{y} of arbitrary dimensions with varying mutual information. Enables analytical computation of the **PMD** density ratio, unavailable in real-world datasets, allowing direct quantification of approximation error for the conditional expectation operator $E_{\mathbf{y}|\mathbf{x}}$ and its **PMD** density ratio .

CI Confidence Interval.

cLoRa Contrastive Low-Rank loss from (Kostic et al., 2024; Ryu et al., 2024) for operator and representation learning. Used in density-ratio fitting (Sugiyama et al., 2012), representation learning (HaoChen et al., 2021; Wang et al., 2022c), and mutual information estimation (Tsai et al., 2020) .

CNN Convolutional NN.

CNNaug CNN trained with data-augmentation.

CoM Center of Mass.

CoT cost-of-transport.

CQR Conditional Quantile Regression (Feldman et al., 2023): A multivariate neural network approach for regressing upper and lower quantiles at a specified miscoverage level α using pinball loss. Confidence intervals are typically calibrated post-training via conformal prediction, but calibration is omitted here for fair model comparison .

DAE Dynamics Autoencoder: Originally proposed in Lusch et al. (2018). .

DAEaug Dynamics Autoencoder with data augmentation.

DL Deep Learning.

DoF degree of freedom.

- DRF** Density Ratio Fitting (Tsai et al., 2020): A density ratio NN architecture that parameterizes the approximated PMD $\kappa_\theta : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ as a single NN. Consequently, this model cannot be used for downstream conditional probability estimation and regression—it is limited to estimating the mutual information between \mathbf{x} and \mathbf{y} (Tsai et al., 2020).
- eCNN** equivariant CNN.
- eCQR** Version of eCQR where the upper and lower parametric quantile functions are parameterized by \mathbb{G} -equivariant NNs.
- eDAE** Equivariant Dynamics Autoencoder: A \mathbb{G} -equivariant adaptation of the DAE architecture (Lusch et al., 2018), introduced in Ordoñez-Apraez et al. (2024).
- E-IPPO** Equivariant Independent Proximal Policy Optimization: Trains a single policy for one arm (22-DoF) and applies it to both arms. The symmetry effectively doubles the training data and includes task encoding in the observation to facilitate learning.
- eMLP** Equivariant MLP.
- eNCP** Equivariant Neural Conditional Probability.
- E-PPO** Equivariant Proximal Policy Optimization: A single policy with an equivariant network controlling the entire system jointly (44-DoF). Represents standard PPO with symmetry handling.
- eTransformer** equivariant transformer.
- GDL** Geometric Deep Learning: A field of machine learning that incorporates geometric priors into deep learning models (Bronstein et al., 2021).
- GRF** Ground Reaction Forces.
- iDRF** Invariant Density Ratio Fitting: This is a \mathbb{G} -invariant adaptation of the DRF model (Tsai et al., 2020) that parameterizes the approximated PMD κ_θ as a \mathbb{G} -invariant NN.
- IPPO** Independent Proximal Policy Optimization: Learns two fixed policies, each assigned to one arm. Due to scene randomization, each policy must learn both subtasks and select the appropriate one.
- MDP** Markov Decision Process.
- MLP** Multi-Layer Perceptron.
- MLPaug** MLP trained with data augmentation.
- MS** Morphological Symmetries.
- MSE** mean squared error.
- MTMA-POMDP** Multi-Task Multi-Agent POMDP.

NCM The term Normal Configuration Mode (NCM) of motion, introduced in this work, denotes a unique mode of symmetry-constrained motion of a robot configuration. The qualifier "Normal" underscores the orthogonality between different NCMs, drawing parallels with the concept of Normal Vibrational Modes in molecular dynamics (Dresselhaus et al., 2007).

NCP Neural Conditional Probability: A deep representation learning framework (Kostic et al., 2024) for conditional probability estimation and regression with statistical guarantees via operator theory (Baker, 1973). This framework is symmetry-agnostic .

NN Neural Network.

ODE ordinary differential equation.

PMD Pointwise Mutual Dependency Tsai et al. (2020): A pointwise dependency measure between random variables \mathbf{x} and \mathbf{y} , defined as $\kappa(\mathbf{x}, \mathbf{y}) = \frac{d\mathbb{P}_{\mathbf{xy}}(\mathbf{x}, \mathbf{y})}{d(\mathbb{P}_{\mathbf{x}}(\mathbf{x}) \times \mathbb{P}_{\mathbf{y}}(\mathbf{y}))} = \exp(\text{MI}(\mathbf{x}, \mathbf{y}))$.

POMDP Partially Observable Markov Decision Process.

PPO Proximal Policy Optimization.

PPOaug PPO with data-augmentation.

PPOsym PPO with hard equivariance / invariance symmetry constraints.

RL Reinforcement Learning.

s.o.t.a. state-of-the-art.

SDE stochastic differential equation.

SM-aug SYMDEX-aug: Uses the same training scheme as our method but replaces the equivariant network with on-policy data augmentation using group transformations, a common alternative in symmetry learning.

SM-c SYMDEX-c: Matches our method's policy structure but uses a global value function instead of separate ones, ablating the effect of global value functions commonly used in multi-agent cooperative settings.

SVD singular value decomposition.

SymDex SYMDEX: An IPPO variant with individual \mathbb{G} -invariant control policies per arm and separate \mathbb{G} -invariant value functions per arm, evaluating single arm task performance, and knowledge of the permutation symmetry between the tasks. .

UniMal Universal Animal design space.

List of figures

| | | |
|-----|--|----|
| 1.1 | Spatial symmetries of the Atlas humanoid | 2 |
| 1.2 | Taxonomy of symmetries in robot dynamics | 2 |
| 1.3 | Morphological symmetries of the Mini-Cheetah robot | 3 |
| 2.1 | Cayley diagrams for representative finite symmetry groups | 23 |
| 2.2 | Example of an endomorphism acting on a stable irreducible representation | 27 |
| 2.3 | Example of \mathbb{G} -invariant joint and conditional distributions | 34 |
| 3.1 | Example symmetry transformations for dynamically equivalent states | 38 |
| 3.2 | Example symmetry transformations of the Atlas humanoid robot | 39 |
| 3.3 | Symmetry breaking in the double pendulum | 41 |
| 3.4 | Geometric constraints induced by ignorable coordinates | 43 |
| 3.5 | Configuration manifold constraints of the Atlas humanoid robot | 44 |
| 3.6 | Configuration manifold constraints of the Solo quadruped robot | 45 |
| 4.1 | Environmental vs morphological group actions | 47 |
| 4.2 | Example modular kinematic trees | 51 |
| 4.3 | Permutations of the Mini-Cheetah leg-branches | 52 |
| 4.4 | Kinematic and inertial constraints induced by morphological symmetries. | 56 |
| 4.5 | Reflection-symmetric mass distribution equivalences | 57 |
| 4.6 | Projection of symmetry-related states onto isotypic subspaces | 59 |
| 4.7 | Decomposition of joint space dynamics into lower-dimensional sub-motions | 61 |
| 5.1 | Symmetry priors for scalar random variables | 64 |
| 5.2 | Parametric model approximation under a symmetry prior | 66 |
| 6.1 | Equivariant stochastic dynamics of the Solo quadruped robot | 72 |

| | | |
|------|--|-----|
| 6.2 | Symmetry transfer from state to observation transition distribution | 74 |
| 6.3 | Stochastic dynamics model graphical model | 75 |
| 6.4 | Model symmetrization under the Klein-four group | 77 |
| 7.1 | Symmetry-invariance of the optimal state estimation distribution | 81 |
| 7.2 | Sample efficiency for the regression of center of mass momenta | 85 |
| 7.3 | Static-friction-regime contact state classification results | 87 |
| 7.4 | Example uninformative conditional expectations. | 88 |
| 7.5 | Target conditional probability distribution and metrics | 92 |
| 7.6 | Confidence intervals for invariant observables in quadruped locomotion | 94 |
| 7.7 | Confidence intervals for equivariant ground-reaction forces | 95 |
| 8.1 | Reflectional symmetries of bimanual mobile manipulation | 98 |
| 8.2 | Control policies induced by invariant quality functions | 100 |
| 8.3 | Push-T reflection symmetry expert rollout | 108 |
| 8.4 | Flow-matching probability transport paths | 109 |
| 8.5 | Equivariant Transformer architecture for policy modeling | 111 |
| 8.6 | Benchmark environments for mobile-manipulation | 112 |
| 8.7 | Control performance versus expert demonstrations | 112 |
| 8.8 | Relative gain in spatial-generalization experiments | 114 |
| 8.9 | Training curves and snapshots for quadruped locomanipulation | 117 |
| 8.10 | Stand-and-turn qualitative and quantitative comparison | 118 |
| 8.11 | Zero-shot real-world deployment in quadruped locomanipulation | 119 |
| 8.12 | Dexterous bimanual manipulation symmetry-related tasks | 120 |
| 8.13 | Performance of SymDex and baselines in dexterous manipulation | 120 |
| 8.14 | Environment-policy rollout for the multi-arm task | 122 |
| 9.1 | Group action on function spaces and properties | 125 |
| 9.2 | Isotypic decomposition of a symmetric function space | 127 |
| 9.3 | ENCP vs NCP Neural Network architecture | 130 |
| 9.4 | Sample-efficiency plots for estimation error | 134 |
| 9.5 | Illustration of a maximal closed-loop invariant set | 139 |
| 9.6 | A symmetric closed-loop invariant geometry | 141 |
| 9.7 | Equivariant Dynamics AutoEncoder architecture | 143 |
| 9.8 | Cayley diagram and isotypic decomposition of the Mini-Cheetah | 145 |

| | | |
|------|--|-----|
| 9.9 | Multi-step prediction errors for long-horizon forecasting | 146 |
| 9.10 | Quantitative metrics and predicted trajectories for the Mini-Cheetah | 147 |

List of tables

| | | |
|-----|---|-----|
| 7.1 | Validation and test metrics for confidence intervals | 93 |
| 8.1 | Zero-shot evaluation in simulated mobile manipulation | 113 |
| 8.3 | Performance of distilled student policy over different parametrizations | 121 |
| 9.1 | Statistical learning guarantees of NCP | 128 |
| 9.2 | Statistical learning guarantees for ENCP | 132 |