

Dynamics Harmonic Analysis of Robotic Systems: Application in Data-Driven Koopman Modelling

Daniel Felipe Ordoñez-Apraez^{1,2}

Vladimir Kostic¹

Giulio Turrisi²

Pietro Novelli¹

Carlos Mastalli^{3*}

Claudio Semini²

Massimiliano Pontil^{1†}

DANIEL.ORDONEZ@IIT.IT

VLADIMIR.KOSTIC@IIT.IT

GIULIO.TURRISI@IIT.IT

PIETRO.NOVELLI@IIT.IT

C.MASTALLI@HW.AC.UK

CLAUDIO.SEMINI@IIT.IT

MASSIMILIANO.PONTIL@IIT.IT

¹ *Computational Statistics and Machine Learning - Istituto Italiano di Tecnologia*

² *Dynamic Legged Systems - Istituto Italiano di Tecnologia*

³ *Robot Motor Intelligence – Heriot-Watt University*

Abstract

We introduce the use of harmonic analysis to decompose the state space of symmetric robotic systems into orthogonal isotypic subspaces. These are lower-dimensional spaces that capture distinct, symmetric, and synergistic motions. For linear dynamics, we characterize how this decomposition leads to a subdivision of the dynamics into independent linear systems on each subspace, a property we term dynamics harmonic analysis (DHA). To exploit this property, we use Koopman operator theory to propose an equivariant deep-learning architecture that leverages the properties of DHA to learn a global linear model of the system dynamics. Our architecture, validated on synthetic systems and the dynamics of locomotion of a quadrupedal robot, exhibits enhanced generalization, sample efficiency, and interpretability, with fewer trainable parameters and computational costs.

Keywords: Symmetric dynamical systems, Harmonic analysis, Koopman operator, Robotics

1. Introduction

The current state-of-the-art in modelling, control, and estimation of robotic systems relies on the Lagrangian model of rigid-body dynamics, which represents the system’s state as a point in the space of generalized (or minimal) coordinates. This approach has fostered the development of efficient algorithms that leverage the system’s kinematic structure to enable recursive computations which are ubiquitous in methods for simulation, estimation, planning, and control in robotics (Featherstone, 2007). However, because the model’s dynamics are nonlinear, these methods need to cope with the challenges of nonlinear optimization, often through iterative local (or state-dependent) linearizations (Mayne, 1966; Li and Todorov, 2004). This can limit control policies to local minima, hinder the convergence of optimization methods, and bias the estimation of unobserved quantities.

The Koopman operator framework can potentially address these limitations by deriving globally linear models of robot dynamics, albeit in an infinite-dimensional function space (Brunton et al., 2022; Kostic et al., 2023). These models can be easily used with various estimation and control theory techniques (Mauroy et al., 2020) and can capture dynamic phenomena that impact the system’s evolution but are challenging to model analytically (Asada, 2023). However, building a robust data-driven model approximation in finite dimensions is a substantial machine-learning challenge. This

* Also with the Florida Institute for Human and Machine Cognition, US.

† Also with the Department of Computer Science, University College London, UK.

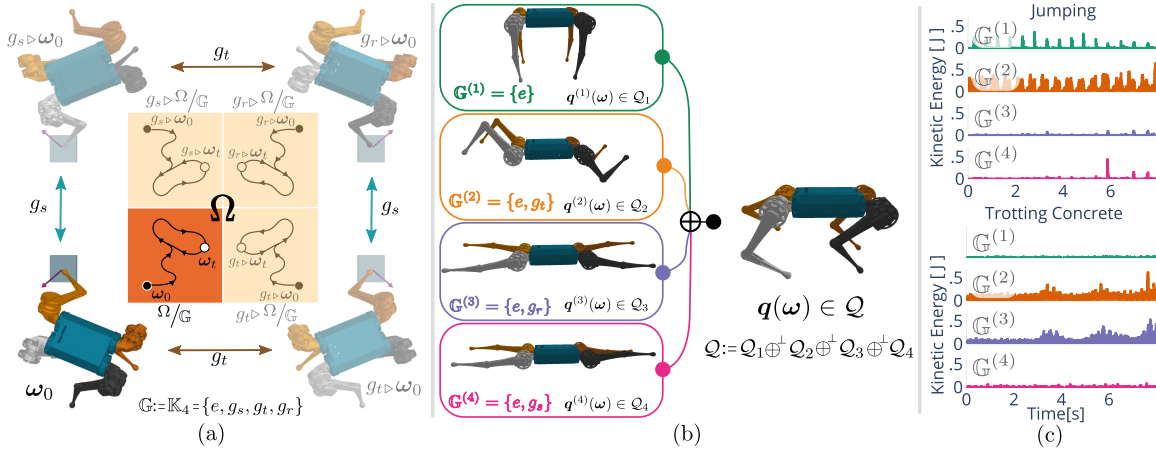


Figure 1: (a) Diagram of the discrete symmetry group $\mathbb{G} := \mathbb{K}_4$ of the mini-cheetah robot (see [animation](#)). Each symmetry $g \in \mathbb{G}$ relates states that evolve identically under physics laws; see [evolution](#) of states related by g_s . This results in the decomposition of the set of states Ω into g -transformed copies of the quotient set Ω/\mathbb{G} , encompassing all unique system states. (b) Isotypic decomposition of the robot’s space of generalized position coordinates Q into isotypic subspaces: $Q := \bigoplus_{i=1}^4 Q_i$. Each subspace, Q_i describes a space of symmetry-constrained synergistic motions. Consequently, any position configuration $q(\omega) \in Q$, can be decomposed into projections within these subspaces: $q(\omega) := \bigoplus_{i=1}^4 q^{(i)}(\omega)$ (see [animation](#)). (c) Joint-space Kinetic energy distribution across isotypic subspaces for two gait/motion trajectories in the real world: jumping and trotting (see [animation](#)). Both gaits primarily evolve within one or two lower-dimensional isotypic subspaces, with less significant subspaces engaged temporarily during disturbances.

study highlights the value of leveraging the state symmetries inherent in a system’s dynamics as a geometric prior, enhancing the approximation of the operator. This approach improves sample efficiency, generalization, and interpretability while reducing the number of trainable parameters and computational cost. Although our emphasis is on robotics, the proposed framework is broadly applicable to any dynamical system with discrete symmetry groups.

In robotics, the symmetries we aim to exploit are known as morphological symmetries ([Ordóñez-APraez et al., 2023](#)). These are discrete symmetry groups that capture the equivariance of the robot’s dynamics, arising from the duplication of rigid bodies and kinematic chains (see [fig. 1-a](#)). Inspired by the use of harmonic analysis of discrete symmetry groups in physics ([Dresselhaus et al., 2007](#)), we present the first application of these principles in robotics. Specifically, we employ the isotypic decomposition ([thm. 1](#)) to partition the state space of symmetric robotic systems into isotypic subspaces (see [fig. 1-b](#)). This allows us to represent any system motion (e.g., different locomotion gaits or manipulation movements) as a superposition of simpler, symmetric, synergistic motions¹; each evolving in a distinct isotypic subspace (see [fig. 1-c](#)). Moreover, we demonstrate that the isotypic decomposition leads to partitioning any (local or global) linear model of the system’s dynamics into independent linear subsystems, each characterizing the dynamics of an isotypic subspace. This offers numerous computational advantages and enhances the interpretability of the dynamics.

Contributions In summary, our work relies on the following contributions: (i) we introduce the use of harmonic analysis of discrete symmetry groups to robotics; (ii) we propose the concept of dynamics harmonic analysis (DHA), illustrating how local/global linear dynamics models decompose

1. These are the robotics analog of normal vibrational modes in molecular dynamics; see ([Dresselhaus et al., 2007, 8.3](#))

into linear subsystems within isotypic subspaces (sec. 3); (iii) leveraging on DHA, we propose the equivariant Dynamics Autoencoder (eDAE), a deep-learning architecture to approximate the Koopman operator (sec. 4), and report its strong performance on synthetic and robotics systems (sec. 5); and (iv) we provide an open-access [repository](#) enabling the application of harmonic analysis to a [library](#) of symmetric robots, the use of the eDAE architecture, and our experiments’ reproduction.

Related work For a data-driven approximation of Koopman operators [Noé and Nuske \(2013\)](#), [Lusch et al. \(2018\)](#), and [Kostic et al. \(2023\)](#) introduce symmetry-agnostic deep learning algorithms, which can be adapted to exploit DHA, as proposed in sec. 4. In the spirit of our work, [Steyert \(2022\)](#) studies the operator’s structure for systems evolving on manifolds, i.e., featuring continuous symmetry groups. For discrete symmetry groups, [Salova et al. \(2019\)](#) leverages the operator’s block-diagonal structure after using harmonic analysis on a non-learnable dictionary of observable functions to model Duffing oscillators. Lastly, [Mesbahi et al. \(2019\)](#); [Sinha et al. \(2020\)](#) provides a theoretical analysis of the operator’s structure to model symmetric dynamical systems.

In robotics, the linear nature of Koopman operator models makes them compatible with standard modelling, estimation, and control algorithms. This includes optimal and robust control ([Korda and Mezić, 2018](#); [Folkestad and Burdick, 2021](#); [Zhang et al., 2022](#)), active learning ([Abraham and Murphey, 2019](#)), and system identification and observer synthesis ([Bruder et al., 2020](#); [Surana, 2020](#)). Yet, symmetries remain unexploited in these Koopman-based approaches.

2. Preliminaries

This section overviews background material on dynamical systems modelling, Koopman models, and harmonic analysis needed to address the modelling of symmetric dynamical systems in sec. 3.

Modelling of dynamical systems In our analysis, distinguishing between a dynamical system and its numerical model is crucial. A dynamical system abstracts evolving real-world phenomena, such as a robot’s motion in an environment. Conversely, numerical models *approximate* the system’s dynamics as the time evolution of points (representing the system’s state) in a vector space. A dynamical system is typically denoted by a tuple $(\Omega, \mathbb{T}, \Phi_\Omega^{\Delta t})$, where Ω is the abstract set of system states $\omega \in \Omega$, the set \mathbb{T} represents time, and $\Phi_\Omega^{\Delta t} : \Omega \times \mathbb{T} \mapsto \Omega$ is the evolution map, determining the future state based on the present state and time. Given deterministic Newtonian mechanics governs our target systems, we focus on deterministic, Markovian, discrete-time systems. Here, $\mathbb{T} = \mathbb{N}_0$ and $\Phi_\Omega^{\Delta t}$ depends solely on the time difference between consecutive timesteps Δt , i.e., $\omega_{t+\Delta t} := \Phi_\Omega^{\Delta t}(\omega_t)$ for any $t \in \mathbb{T}, \omega \in \Omega$. Thus, we will denote a dynamical system by $(\Omega, \Phi_\Omega^{\Delta t})$.

Defining a numerical model of a dynamical system involves identifying a state representation vector-valued function $\mathbf{x} = [x_1, \dots, x_m] : \Omega \rightarrow \mathcal{X} \subseteq \mathbb{R}^m$, where the components, $x_j : \Omega \rightarrow \mathbb{R}, j \in [1, m]$, are observable functions that measure a relevant scalar quantity from the state (e.g., kinetic energy, joint position/velocity). This enables the representation of the state as a point in the model’s vector space $\mathbf{x}(\omega) \in \mathcal{X}$. The system’s evolution, represented as a trajectory $(\mathbf{x}(\omega_t))_{t \in \mathbb{T}}$, can be approximated by an evolution map $\mathbf{x}(\tilde{\omega}_{t+\Delta t}) := \Phi_{\mathcal{X}}^{\Delta t}(\mathbf{x}(\omega_t)) + \kappa(\omega_t)\Gamma_t$. Here, the predicted state representation $\mathbf{x}(\tilde{\omega}_{t+\Delta t})$ may differ from the true one $\mathbf{x}(\omega_{t+\Delta t})$ due to modelling errors and/or inaccessible observables. The influence of these factors is assumed to be captured by a mild white noise stochastic perturbation Γ_t , scaled by $\kappa(\omega_t)$ ([Lasota and Mackey, 1994](#), chpt. 10.5). When $\mathbf{x} : \Omega \rightarrow \mathcal{X}$ is injective, the optimality of the model $(\mathcal{X}, \Phi_{\mathcal{X}}^{\Delta t})$, over a given prediction horizon $H \in \mathbb{T}$, is quantified by its predictive error ([Mezić, 2021](#)), i.e.,

$$\text{err}_H(\mathbf{x}, \Phi_{\mathcal{X}}^{\Delta t}) = \int_{\Omega} \text{err}_{\omega_0, H}(\mathbf{x}, \Phi_{\mathcal{X}}^{\Delta t}) d\omega_0, \quad \text{given} \quad \text{err}_{\omega_0, H}(\mathbf{x}, \Phi_{\mathcal{X}}^{\Delta t}) := \sum_{h=1}^H \|\mathbf{x}(\tilde{\omega}_{h\Delta t}) - \mathbf{x}(\omega_{h\Delta t})\|^2. \quad (1)$$

Note that while modelling error can be locally minimal for specific ω_0 (r.h.s eq. (1)), optimal models exhibit uniformly small errors across all states and horizons (l.h.s eq. (1)).

Linear models and the Koopman operator A linear model $(\mathcal{Z}, \mathbf{K}_{\Delta t})$ is a dynamics model where the temporal evolution of state representations is characterized by an autonomous linear system, $\mathbf{z}(\tilde{\omega}_{t+\Delta t}) = \Phi_{\mathcal{Z}}^{\Delta t}(\mathbf{z}(\omega)) := \mathbf{K}_{\Delta t} \mathbf{z}(\omega)$, $\omega \in \Omega$, $\Delta t \in \mathbb{T}$. Here, $\mathbf{z} = [z_1, \dots] : \Omega \mapsto \mathcal{Z}$ is the state representation function, and $\mathbf{K}_{\Delta t} : \mathcal{Z} \mapsto \mathcal{Z}$ is a matrix evolving state representations in time by Δt . The linearity of these models ensures the existence of analytical solutions to the temporal dynamics of each state observable and guarantees the interpretability of the predictions through modal decomposition, making them fundamental to dynamical systems and control theory.

Although most dynamical systems of interest in robotics have been historically modeled with nonlinear analytic dynamics, one can also devise an optimal linear model in an infinite-dimensional space. This idea, originating from the seminal work of Koopman and Markov (Lasota and Mackey, 1994), proposes to represent the state as a point/function in a *space of functions* $\mathcal{F}_{\mathcal{Z}}$, and model the dynamics with a linear operator $\mathbf{K}_{\Delta t}$, defined by the flow $\Phi_{\Omega}^{\Delta t}$ and a time step $\Delta t \in \mathbb{T}$, which takes any function $z(\cdot) \in \mathcal{F}_{\mathcal{Z}}$ to $z(\Phi_{\Omega}^{\Delta t}(\cdot, \Delta t))$. Whenever the image of $\mathbf{K}_{\Delta t}$ is in the same space, $\mathbf{K}_{\Delta t} : \mathcal{F}_{\mathcal{Z}} \mapsto \mathcal{F}_{\mathcal{Z}}$ is a well-defined linear operator known as the *Koopman operator*, defined by

$$[\mathbf{K}_{\Delta t} z](\omega_t) := z(\Phi_{\Omega}^{\Delta t}(\omega_t)), \quad z \in \mathcal{F}_{\mathcal{Z}}, \quad \omega \in \Omega. \quad (2)$$

The requirement for the space $\mathcal{F}_{\mathcal{Z}}$ to be invariant under the flow $\Phi_{\Omega}^{\Delta t}$ is the characteristic that often renders it infinite-dimensional. As we discuss in sec. 4, machine learning can aid in finding finite-dimensional approximations of $\mathcal{F}_{\mathcal{Z}}$ and $\mathbf{K}_{\Delta t}$ (Lusch et al., 2018; Kostic et al., 2023).

Symmetry groups and their representations In the context of a dynamical systems, symmetries are defined as bijections on the state set Ω . The action of a symmetry transformation g on any state $\omega \in \Omega$ is defined by a map $(\triangleright) : \mathbb{G} \times \Omega \rightarrow \Omega$ to a symmetric state $g \triangleright \omega \in \Omega$ (see fig. 1-a). A set of symmetry transformations forms a group $\mathbb{G} = \{e, g_1, g_2, \dots\}$, that is closed under *composition*: $g_1 \circ g_2 \in \mathbb{G}$ for all $g_1, g_2 \in \mathbb{G}$, and *inversion*: $g^{-1} \in \mathbb{G} \mid g \in \mathbb{G}$ such that $g^{-1} \circ g = e$, where e denotes the identity transformation and (\circ) is the binary composition operation on \mathbb{G} .

As we study symmetries of numerical models in both finite-dimensional (Euclidean) spaces and infinite-dimensional function spaces, we assume \mathcal{H} to be a separable Hilbert space to accommodate both scenarios. This enable us to rely on the conventional concepts of inner product, orthogonality, and countably many basis elements (see appendix I.1). Consequently, symmetry transformations on \mathcal{H} are defined via a unitary group representation $\rho_{\mathcal{H}} : \mathbb{G} \rightarrow \mathbb{U}(\mathcal{H})$, mapping each $g \in \mathbb{G}$ to a unitary matrix/operator $\rho_{\mathcal{H}}(g) \in \mathbb{U}(\mathcal{H}) : \mathcal{H} \rightarrow \mathcal{H}$. Thus, the action of any $g \in \mathbb{G}$ on a point $\mathbf{h} \in \mathcal{H}$ is expressed as $g \triangleright \mathbf{h} := \rho_{\mathcal{H}}(g)\mathbf{h} \in \mathcal{H}$. When $\rho_{\mathcal{H}}$ exists, we say that \mathcal{H} is a symmetric space.

A map $f : \mathcal{H} \mapsto \mathcal{H}'$ between two symmetric spaces is denoted \mathbb{G} -equivariant if $f(\rho_{\mathcal{H}}(g)\mathbf{h}) = \rho_{\mathcal{H}'}(g)f(\mathbf{h})$, and \mathbb{G} -invariant if $f(\mathbf{h}) = f(\rho_{\mathcal{H}}(g)\mathbf{h})$, for all $g \in \mathbb{G}$. While the action of a symmetry transformation on \mathcal{H} is basis independent, the representation of $g \in \mathbb{G}$ depends on the chosen basis of \mathcal{H} . Therefore, applying a change of basis $\mathbf{Q} \in \mathbb{U}(\mathcal{H})$ results in a new point representation $\mathbf{h}_{\circ} = \mathbf{Q}\mathbf{h}$, and a new group representation $\rho_{\mathcal{H}, \circ}(\cdot) = \mathbf{Q}\rho_{\mathcal{H}}(\cdot)\mathbf{Q}^*$. Here, both representations describe the

same transformation: $\rho_{\mathcal{H}}(g)\mathbf{h} \iff \rho_{\mathcal{H},\circ}(g)\mathbf{h}_\circ$. Consequently, representations related by a basis change are termed equivalent, denoted $\rho_{\mathcal{H}} \sim \rho_{\mathcal{H},\circ}$. Lastly, writing $\rho_{\mathcal{H}} \sim \rho_{\mathcal{H}_1} \oplus \rho_{\mathcal{H}_2}$ implies that \mathcal{H} decomposes into orthogonal subspaces \mathcal{H}_1 and \mathcal{H}_2 , and $\rho_{\mathcal{H}}$ has block-diagonal structure.

Isotypic decomposition and its basis Our use of harmonic analysis is linked with the decision to work in a specific basis for the modelling space \mathcal{H} , the *isotypic basis*. In this basis, the unitary group representations $\rho_{\mathcal{H}}$ decomposes into a block-diagonal sum of multiple copies (*multiplicities*) of the group's k unique irreducible representations (*irrep*) $\{\bar{\rho}_i\}_{i=1}^k$. These are the indivisible building blocks of any group representation of \mathbb{G} . Each $\bar{\rho}_i : \mathbb{G} \rightarrow \mathbb{U}(\bar{\mathcal{H}}_i)$, describes a unique symmetry pattern, characterized by a subset of symmetry transformations within the broader group structure. The space $\bar{\mathcal{H}}_i$ associated to each *irrep* is the smallest finite-dimensional space capable of expressing the *irrep* symmetry pattern. For instance, our groups \mathbb{G} are often subgroups of the orthogonal group. Thus, we frequently work with *irreps* $\bar{\rho}_{tr}$ that describe a reflection symmetry, requiring a 1-dimensional space $\bar{\mathcal{H}}_{tr} \sim \mathbb{R}$ to act on, or with *irreps* $\bar{\rho}_{2\pi/a}$ that describe rotations by an angle $2\pi/a$, requiring a 2-dimensional space $\bar{\mathcal{H}}_{2\pi/a} \sim \mathbb{R}^2 \sim \mathbb{C}$ to act on. These representations are called irreducible because the spaces $\bar{\mathcal{H}}_i$ have no non-trivial invariant subspace to the actions of \mathbb{G} . That is, if \mathcal{V} is a subspace of $\bar{\mathcal{H}}_i$ and $\bar{\rho}_i(g)\mathcal{V} \subset \mathcal{V}$ for every $g \in \mathbb{G}$, then either $\mathcal{V} = \{0\}$ or $\mathcal{V} = \bar{\mathcal{H}}_i$.

The value of the isotypic basis lies in the fact that it allows decomposing the modelling space \mathcal{H} into an orthogonal sum of isotypic subspaces. Each subspace reflects a unique symmetry pattern of one of the group's *irreps*, hence the term *iso-typic* or *same-type*. This is a pivotal result in abstract harmonic analysis, succinctly captured by the Peter-Weyl Theorem, see (Knapp, 1986, thm 1.12) and (Golubitsky et al., 2012, Thm-2.5).

Theorem 1 (Isotypic Decomposition) *Let \mathbb{G} be a compact symmetry group and \mathcal{H} a symmetric separable Hilbert space with a unitary group representation $\rho_{\mathcal{H}} : \mathbb{G} \rightarrow \mathbb{U}(\mathcal{H})$. Then we can identify $k \leq |\mathbb{G}|$ irreducible representations $\bar{\rho}_i : \mathbb{G} \rightarrow \mathbb{U}(\bar{\mathcal{H}}_i)$ and change of basis $Q \in \mathbb{U}(\mathcal{H})$ such that $\rho_{\mathcal{H}} \sim \rho_{\mathcal{H}_1} \oplus \rho_{\mathcal{H}_2} \oplus \dots \oplus \rho_{\mathcal{H}_k}$, and each $\rho_{\mathcal{H}_i} \sim \bigoplus_{j=1}^{m_i} \bar{\rho}_i$ is composed of at most m_i countably many copies of the irreducible representation $\bar{\rho}_i$. This allows to decompose \mathcal{H} into orthogonal subspaces: $\mathcal{H} = \mathcal{H}_1 \oplus^\perp \mathcal{H}_2 \oplus^\perp \dots \oplus^\perp \mathcal{H}_k$. Where each $\mathcal{H}_i := \bigoplus_{j=1}^{m_i} \mathcal{H}_{i,j}$, composed of m_i subspaces isometrically isomorphic to $\bar{\mathcal{H}}_i$, is denoted as an isotypic subspace.*

3. Symmetries of dynamical systems

In the context of dynamical systems, a symmetry is a state transformation that results in another functionally equivalent state under the governing dynamics. From a modelling perspective, symmetries provide a valuable geometric bias, as identifying the dynamics of a single state suffices to capture the dynamics of all of its symmetric states (see fig. 1-a).

Definition 1 (Symmetric dynamical systems) *A dynamical system $(\Omega, \Phi_\Omega^{\Delta t})$ is \mathbb{G} -symmetric, if \mathbb{G} is a symmetry group of the set of states Ω , and the system's evolution map is \mathbb{G} -equivariant, i.e.,*

$$\Phi_\Omega^{\Delta t}(g \triangleright \omega, t) = g \triangleright \Phi_\Omega^{\Delta t}(\omega, t), \quad \forall g \in \mathbb{G}, t \in \mathbb{T}, \omega \in \Omega. \quad (3)$$

The symmetry group of Ω defines an equivalence relationship between any state $\omega \in \Omega$ and its set of symmetric states, denoted $\mathbb{G}\omega = \{g \triangleright \omega | g \in \mathbb{G}\}$. Given the \mathbb{G} -equivariance of the dynamics $\Phi_\Omega^{\Delta t}$, this equivalence implies that a set of symmetric states $\mathbb{G}\omega$ will evolve along a unique trajectory of motion, up to a symmetry transformation $g \in \mathbb{G}$ (see fig. 1-a). When the symmetry group is

discrete (or finite), the state set Ω decomposes into a union of symmetry-transformed copies of the quotient set Ω/\mathbb{G} , containing the system's unique states, that is $\Omega = \cup_{g \in \mathbb{G}} \{g \triangleright \Omega/\mathbb{G}\}$ (see [fig. 1-a](#)).

Modelling symmetric dynamical systems When designing a numerical model $\Phi_{\mathcal{X}}^{\Delta t}$ for a symmetric dynamical system, it is crucial to ensure that the modelling space \mathcal{X} inherits the group structure of Ω . This can be achieved by making the space invariant under the action of \mathbb{G} , i.e., $\mathbf{x}(g \triangleright \cdot) \in \mathcal{X}$ for all $g \in \mathbb{G}$. Thus, enabling the existence of a group representation $\rho_{\mathcal{X}} : \mathbb{G} \rightarrow \mathbb{U}(\mathcal{X})$. The significance of this design choice lies in the fact that the equivalence between symmetric states is translated into the corresponding equivalence of their representations in the modelling space:

$$\mathbb{G}\omega := \{g \triangleright \omega, \forall g \in \mathbb{G}\} \iff \mathbb{G}\mathbf{x}(\omega) := \{g \triangleright \mathbf{x}(\omega) := \rho_{\mathcal{X}}(g)\mathbf{x}(\omega) = \mathbf{x}(g \triangleright \omega), \forall g \in \mathbb{G}\}, \forall \omega \in \Omega \quad (4)$$

The symmetric structure of \mathcal{X} allows its decomposition into g -transformed copies of a quotient space \mathcal{X}/\mathbb{G} . This is a practical tool in data-driven applications, mitigating the effects of the curse of dimensionality ([Higgins et al., 2022](#)). As the following result shows, it also narrows the search space for the evolution map $\Phi_{\mathcal{X}}^{\Delta t}$ to the space of \mathbb{G} -equivariant ones.

Proposition 2 (Optimal models of \mathbb{G} -symmetric systems) *Let $\Phi_{\Omega}^{\Delta t}$ be a \mathbb{G} -symmetric dynamical system and $\Phi_{\mathcal{X}}^{\Delta t}$ an optimal model ([eq. \(1\)](#)). If \mathcal{X} is a \mathbb{G} -symmetric space, $\Phi_{\mathcal{X}}^{\Delta t}$ is \mathbb{G} -equivariant.*

Proof Let $\mathbb{G}\omega$ denote an arbitrary set of symmetric states and $\mathbb{G}\mathbf{x}(\omega)$ represent their symmetric representations on \mathcal{X} . Consider a non \mathbb{G} -equivariant evolution map, denoted as $\bar{\Phi}_{\mathcal{X}}^{\Delta t}$. Such map will yield varying prediction errors for the states in $\mathbb{G}\omega$. Consequently, it is possible to identify the state with the minimum prediction error, $\hat{g} \triangleright \omega \in \mathbb{G}\omega$, given $\hat{g} = \arg \min_{g \in \mathbb{G}} \text{err}_{g \triangleright \omega, H}(\mathbf{x}, \bar{\Phi}_{\mathcal{X}}^{\Delta t})$, as per [eq. \(1\)](#). Subsequently, we can construct a new map that replicates the predicted evolution of $\hat{g} \triangleright \omega$ for all symmetric states, defined as $\Phi_{\mathcal{X}}^{\Delta t}(\mathbf{x}(g \triangleright \omega)) := (g \circ \hat{g}^{-1}) \triangleright \bar{\Phi}_{\mathcal{X}}^{\Delta t}(\mathbf{x}(\hat{g} \triangleright \omega))$, $\forall g \in \mathbb{G}$. It is important to note that $\text{err}_{g \triangleright \omega, H}(\mathbf{x}, \Phi_{\mathcal{X}}^{\Delta t}) \leq \text{err}_{g \triangleright \omega, H}(\mathbf{x}, \bar{\Phi}_{\mathcal{X}}^{\Delta t})$ holds for all $g \in \mathbb{G}$. By iteratively applying this process for all $\omega \in \Omega/\mathbb{G}$, the resulting map will exhibit \mathbb{G} -equivariance. ■

Consistent with [def. 1](#), we denote models $(\mathcal{X}, \Phi_{\mathcal{X}}^{\Delta t})$ that possess a \mathbb{G} -symmetric modelling state space \mathcal{X} and a \mathbb{G} -equivariant evolution map as \mathbb{G} -symmetric models. A familiar example is the Lagrangian model of rigid-body dynamics. Since, for \mathbb{G} -symmetric robotic systems (e.g., the mini-cheetah in [fig. 1](#)), the modelling space $\mathcal{X} = \mathcal{Q} \times \mathcal{T}_q \mathcal{Q}$, defined by the space of generalized position \mathcal{Q} and velocity $\mathcal{T}_q \mathcal{Q}$ coordinates, is a symmetric vector space ([Ordóñez-Ápraez et al., 2023](#), III). The symmetry of the space is characterized by the group representation $\rho_{\mathcal{X}} := \rho_{\mathcal{Q}} \oplus \rho_{\mathcal{T}_q \mathcal{Q}}$, which describes the transformations shown in [fig. 1-a](#). Furthermore, the evolution map $\Phi_{\mathcal{X}}^{\Delta t}$, defined by the standard Euler-Lagrange equations of motion, features \mathbb{G} -equivariance ([Lanczos, 2012](#), VII.2).

A key property of (non-linear or linear) \mathbb{G} -symmetric models is the decomposition of the modelling state space into k isotypic subspaces $\mathcal{X} = \bigoplus_{i=1}^k \mathcal{X}_i$ (refer to [thm. 1](#)). This enables the projection of entire motion trajectories $(\mathbf{x}(\omega_t))_{t \in \mathbb{T}}$ onto each isotypic subspace $\mathbf{x}(\omega_t) := \mathbf{x}^{(1)}(\omega_t) \oplus^{\perp} \dots \oplus^{\perp} \mathbf{x}^{(k)}(\omega_t)$. Since each \mathcal{X}_i is a lower-dimensional space with a reduced number of symmetries, the trajectory's decomposition entails its characterization as a superposition of distinct synergistic motions $(\mathbf{x}^{(i)}(\omega_t))_{t \in \mathbb{T}} \mid \mathbf{x}^{(i)}(\omega_t) \in \mathcal{X}_i$, each constrained to feature the subset of symmetries of \mathcal{X}_i (see [fig. 1-b](#)). This understanding is instrumental in characterizing different system behaviors, such as different locomotion gaits or manipulation tasks, through their lower-dimensional projections onto each isotypic subspace, as detailed in [fig. 1-c](#).

For \mathbb{G} -symmetric linear dynamics models $(\mathcal{Z}, \mathbf{K}_{\Delta t})$, the isotypic decomposition ([thm. 1](#)) of $\mathcal{Z} = \bigoplus_{i=1}^k \mathcal{Z}_i = \bigoplus_{i=1}^k \bigoplus_{j=1}^{m_i} \mathcal{Z}_{i,j}$ and the associated group representation $\rho_{\mathcal{Z}} = \bigoplus_{i=1}^k \rho_{\mathcal{Z}_i} =$

$\bigoplus_{i=1}^k \bigoplus_{j=1}^{m_i} \bar{\rho}_i$ imply the dynamics' decomposition into k linear subsystems $\mathbf{K}_{\Delta t} = \bigoplus_{i=1}^k \mathbf{K}_{\Delta t, i}$. Each subsystem is \mathbb{G} -equivariant, $\rho_{\mathcal{Z}_i}(g) \mathbf{K}_{\Delta t, i} = \mathbf{K}_{\Delta t, i} \rho_{\mathcal{Z}_i}(g)$, $\forall g \in \mathbb{G}$, and evolves the state projections into isotypic subspaces independently. Moreover, each subsystem is block-decomposed into scalar multiples of the identity map $\mathbf{I}_i : \mathbb{R}^{|\bar{\rho}_i|} \mapsto \mathbb{R}^{|\bar{\rho}_i|}$, such that

$$\mathbf{K}_{\Delta t} = \begin{bmatrix} \mathbf{K}_{\Delta t, 1} & & \\ & \ddots & \\ & & \mathbf{K}_{\Delta t, k} \end{bmatrix}, \quad \mathbf{K}_{\Delta t, i} = \begin{bmatrix} c_{1,1} \mathbf{I}_i & \dots & c_{1,d_i} \mathbf{I}_i \\ \vdots & \ddots & \vdots \\ c_{d_i,1} \mathbf{I}_i & \dots & c_{d_i,d_i} \mathbf{I}_i \end{bmatrix}, \quad c_{i,j} \in \mathbb{R}, \quad d_i := |\mathcal{Z}_i|/|\bar{\rho}_i|, \forall i \in [1, k] \quad (5)$$

The block-diagonal structure of $\mathbf{K}_{\Delta t}$, and the block-structure of each $\mathbf{K}_{\Delta t, i}$ are geometric constraints that originate from Schur's Lemma (lemma 1), a standard result in harmonic analysis. These properties, essentially stated in Golubitsky et al. (2012, Thm 3.5) for finite-dimensional spaces, can be generalized to Hilbert spaces, leading to the following result:

Theorem 3 (Isotypic decomposition of symmetric linear models)

Let $(\mathcal{Z}, \mathbf{K}_{\Delta t})$ be a \mathbb{G} -symmetric linear model, \mathbb{G} be a finite group, and $\mathcal{Z} = \bigoplus_{i=1}^k \mathcal{Z}_i$ be the space isotypic decomposition. Then $\mathbf{K}_{\Delta t}$ is block-diagonal $\mathbf{K}_{\Delta t} = \bigoplus_{i=1}^k \mathbf{K}_{\Delta t, i}$, where each $\mathbf{K}_{\Delta t, i} : \mathcal{Z}_i \mapsto \mathcal{Z}_i$ is a \mathbb{G} -equivariant linear map characterizing the independent dynamics of each isotypic subspace.

This result has two primary applications in robot dynamics modelling. The first involves decomposing local linear models resulting from local (or state-dependent) linearizations of nonlinear dynamical models. These are widely used in the iLQR (Li and Todorov, 2004) and DDP (Mayne, 1966) algorithms, which are instrumental for trajectory optimization (Tassa et al., 2014; Mastalli et al., 2020), and state estimation (Alessandri et al., 2003; Kobilarov et al., 2015; Martinez et al., 2024). For these methods, the model decomposition facilitates the parallel optimization of the k local linear sub-models, presenting a promising direction for future research. The second application pertains to the approximation of \mathbb{G} -symmetric Koopman models, which we discuss next.

Dynamic harmonic analysis For a \mathbb{G} -symmetric dynamical system $\Phi_{\Omega}^{\Delta t}$, an injective state representation $\mathbf{x} : \Omega \mapsto \mathcal{Z}$, and \mathbb{G} -symmetric modelling space \mathcal{Z} , the Koopman operator $\mathbf{K}_{\Delta t}$ (eq. (2)) is \mathbb{G} -equivariant and globally optimal by construction. Hence, by thm. 3, it decomposes into k Koopman operators characterizing the dynamics of each isotypic subspace. Furthermore, for each eigenpair (λ, ψ) of $\mathbf{K}_{\Delta t}$ and every $g \in \mathbb{G}$, the symmetric function $\psi_g := \psi(g \triangleright (\cdot)) : \Omega \mapsto \mathbb{C}$ is also an eigenfunction of $\mathbf{K}_{\Delta t}$ in the same eigenspace,

$$\lambda \psi_g(\omega) = \lambda \psi(g \triangleright \omega) = [\mathbf{K}_{\Delta t} \psi](g \triangleright \omega) = \psi(\Phi_{\Omega}^{\Delta t}(g \triangleright \omega)) = \psi(g \triangleright \Phi_{\Omega}^{\Delta t}(\omega)) = [\mathbf{K}_{\Delta t} \psi_g](\omega), \quad (6)$$

implying that the Koopman eigenspaces are \mathbb{G} -symmetric spaces. Thus, applying the isotypic decomposition to each eigenspace (thm. 1) captures the relation of the temporal evolution of eigenfunctions with distinct symmetries (encoded by the *irreps*) via Koopman eigenvalues. This global decomposition of the dynamics in isotypic subspaces and its symmetry-aware spectral decomposition is referred to as *dynamics harmonic analysis* (DHA). As we show in secs. 4 and 5, DHA can be leveraged to learn data-driven approximations of the \mathbb{G} -equivariant Koopman operator.

4. \mathbb{G} -symmetric data-driven Koopman models

The Koopman operator formalism, while practically unfeasible, has inspired numerous data-driven models aiming to approximate the infinite-dimensional operator in finite-dimensions (Brunton et al., 2022). This process relies on a dataset of observations of state trajectories $\mathbb{D} = \{(\mathbf{x}(\omega_t))_{t=0}^H, \dots\}$

on a modelling space $\mathcal{X} \subseteq \mathbb{R}^m$. Where $\mathbf{x} = [x_1, \dots, x_m] : \Omega \mapsto \mathcal{X}$ is a state representation composed of physical observable functions $x_i : \Omega \mapsto \mathbb{R}$ that are measured or estimated from the state (e.g., position, momentum, energy). Then, the objective is to find a (latent) state representation vector-valued function $\mathbf{z} = [z_1, \dots, z_\ell] : \Omega \rightarrow \tilde{\mathcal{Z}} \subseteq \mathbb{R}^\ell$, that spans a finite-dimensional space of functions $\mathcal{F}_{\tilde{\mathcal{Z}}} := \{z_\alpha(\cdot) := \sum_{i=1}^\ell \alpha_i z_i(\cdot) = \langle \mathbf{z}(\cdot), \boldsymbol{\alpha} \rangle, \mid \boldsymbol{\alpha} \in \mathbb{R}^\ell\}$ (see [appendix I.1](#)), on which the Koopman operator $\mathbf{K}_{\Delta t}$ is approximated by a matrix $\mathbf{K}_{\Delta t}^* \in \mathbb{C}^{\ell \times \ell}$, given by $(\mathbf{K}_{\Delta t} z_\alpha)(\cdot) \approx z_{\mathbf{K}_{\Delta t}^* \boldsymbol{\alpha}}(\cdot) := \langle \mathbf{z}(\cdot), \mathbf{K}_{\Delta t}^* \boldsymbol{\alpha} \rangle = \langle \mathbf{K}_{\Delta t} \mathbf{z}(\cdot), \boldsymbol{\alpha} \rangle$ ([Kostic et al., 2022](#)). Where $\mathcal{F}_{\tilde{\mathcal{Z}}}$ represent the finite-dimensional approximation of the space of functions $\mathcal{F}_{\mathcal{Z}}$ on which $\mathbf{K}_{\Delta t}$ is defined (see [eq. \(2\)](#)), and $\boldsymbol{\alpha} = [\alpha_1, \dots] \in \mathbb{R}^\ell$ is the coefficient vector representation of z_α in a basis of $\mathcal{F}_{\tilde{\mathcal{Z}}}$.

Among different approaches to building Koopman data-driven models, we focus on the Dynamics Auto-Encoder (DAE) architecture ([Lusch et al., 2018](#)). This model parameterizes the matrix approximating the Koopman operator $\mathbf{K}_{\Delta t}^*$ as a trainable linear map, and the (latent) state representation $\mathbf{z}(\cdot) := (\mathbf{f}_\phi \circ \mathbf{x})(\cdot) : \Omega \mapsto \tilde{\mathcal{Z}}$ with an encoder neural network \mathbf{f}_ϕ . A decoder $\mathbf{f}_\phi^{-1} : \tilde{\mathcal{Z}} \mapsto \mathcal{X}$ is also defined to reconstruct states in the physical observable space. The cost function for DAE is composed of a reconstruction loss (encouraging injectivity of \mathbf{z}) and a state prediction error in both physical observable \mathcal{X} and latent $\tilde{\mathcal{Z}}$ spaces (encouraging the minimization of the modelling error):

$$\mathcal{L}(\omega_t, H) = \underbrace{\sum_{h=0}^H \|\mathbf{x}(\omega_{t+h\Delta t}) - \mathbf{f}_\phi^{-1}(\mathbf{K}_{\Delta t}^h \mathbf{z}(\omega_t))\|^2}_{\text{Reconstruction and err}_{\omega_t, H}(\mathbf{x}, \Phi_{\mathcal{X}}^{\Delta t})} + \underbrace{\gamma \|\mathbf{z}(\omega_{t+h\Delta t}) - \mathbf{K}_{\Delta t}^h \mathbf{z}(\omega_t)\|^2}_{\text{err}_{\omega_t, H}(\mathbf{z}, \mathbf{K}_{\Delta t})}$$

Where γ balances the modelling errors in $\tilde{\mathcal{Z}}$ and \mathcal{X} , and the prediction horizon is assumed $H \ll |\mathbb{D}|$.

The Equivariant DAE (eDAE) When modelling \mathbb{G} -symmetric dynamical systems ([def. 1](#)), the DAE architecture can be adapted to leverage the symmetry priors. First, to exploit the theoretical isotypic decomposition of the space of observable functions $\mathcal{F}_{\mathcal{Z}}$ ([prop. 2](#) and [thm. 1](#)), we must ensure that $\mathcal{F}_{\tilde{\mathcal{Z}}}$ is a \mathbb{G} -symmetric function space, such that it can be decomposed in isotypic subspaces $\mathcal{F}_{\tilde{\mathcal{Z}}} := \bigoplus_{i=1}^k \mathcal{F}_{\tilde{\mathcal{Z}}_i}$. This can be achieved by restricting the state representation $\mathbf{z} : \Omega \rightarrow \tilde{\mathcal{Z}}$ to be a \mathbb{G} -equivariant map, such that both $\tilde{\mathcal{Z}}$ and $\mathcal{F}_{\tilde{\mathcal{Z}}}$ are \mathbb{G} -symmetric spaces (see [appendix I.1.2](#)). In practice, this is done by constraining the encoder to be a \mathbb{G} -equivariant neural network, such that:

$$g \triangleright \mathbf{z}(\omega_t) = \rho_{\tilde{\mathcal{Z}}}(g)(\mathbf{f}_\phi \circ \mathbf{x})(\omega) = \mathbf{f}_\phi(\rho_{\mathcal{X}}(g)\mathbf{x}(\omega)) = \mathbf{f}_\phi(\mathbf{x}(g \triangleright \omega)) = \mathbf{z}(g \triangleright \omega_t), \forall g \in \mathbb{G}, \omega \in \Omega. \quad (7)$$

Where the group representation in the physical observable space $\rho_{\mathcal{X}}$ is assumed to be known from prior knowledge, and the latent group representation $\rho_{\tilde{\mathcal{Z}}}$ is defined to be equivalent to the direct sum of $\ell/|G|$ copies of the group regular representation, following ([Knapp, 1986](#), [thm. 1.12](#)).

Then, to exploit the theoretical \mathbb{G} -equivariance of $\mathbf{K}_{\Delta t}$ ([prop. 2](#) and [thm. 3](#)), we parameterize $\mathbf{K}_{\Delta t}$ as a \mathbb{G} -equivariant matrix, $\rho_{\tilde{\mathcal{Z}}}(g)\mathbf{K}_{\Delta t} = \mathbf{K}_{\Delta t}\rho_{\tilde{\mathcal{Z}}}(g)$, $\forall g \in \mathbb{G}$. Furthermore, if the latent group representation is defined on the isotypic basis $\rho_{\tilde{\mathcal{Z}}} = \bigoplus_{i=1}^k \rho_{\tilde{\mathcal{Z}}_i} = \bigoplus_{i=1}^k \bigoplus_{j=1}^{m_i} \bar{\rho}_i$ (see [thm. 1](#)), then the matrix $\mathbf{K}_{\Delta t}$ decomposes in block-diagonal form $\mathbf{K}_{\Delta t} = \bigoplus_{i=1}^k \mathbf{K}_{\Delta t, i}$, where each $\mathbf{K}_{\Delta t, i}$ is a \mathbb{G} -equivariant matrix characterizing the dynamics on the isotypic subspace $\tilde{\mathcal{Z}}_i$. This constraint on the learnable parameters of $\mathbf{K}_{\Delta t}$ is relevant geometric prior as it (i) ensures the learned latent dynamics model is \mathbb{G} -symmetric, (ii) exploits the orthogonality between distinct isotypic subspaces (eliminating spurious correlations between dimensions of $\tilde{\mathcal{Z}}$), and (iii) constraints the minimum dimension of the eigenspaces of each isotypic subspace operator $\mathbf{K}_{\Delta t, i}$ to match the dimension of the subspace irreducible representation dimension $|\bar{\rho}_i|$ ([eq. \(5\)](#)). This ensures that eigenfunctions of each isotypic subspace $\mathcal{F}_{\tilde{\mathcal{Z}}_i}$ are appropriately approximated in groups of symmetric functions by




Figure 2: Test set prediction mean square error (MSE) of learned Koopman models (DAE and eDAE) for synthetic systems with symmetry groups \mathbb{C}_5 and \mathbb{C}_{10} , varying state dimension $|\mathcal{X}|$, latent model space dimension $|\mathcal{Z}|$, and noise variance σ . Solid lines and shaded areas represent the mean, maximum, and minimum prediction error among 4 training seeds. (a) MSE vs. training samples. (b) MSE over varying state dimensions. (c) MSE over varying dimensionality of the latent model space. (d) MSE over varying noise variance σ .

their coefficient eigenvectors representations $\mathbb{G}\mathbf{v}^{(i)} = \{g \triangleright \mathbf{v}^{(i)} \mid \mathbf{K}_{\Delta t, i} \mathbf{v}^{(i)} = \lambda^{(i)} \mathbf{v}^{(i)}, \mathbf{v}^{(i)} \in \tilde{\mathcal{Z}}_i\}$, with temporal dynamics governed by an eigenvalue $\lambda^{(i)}$; refer to [appendix I.1.2](#).

5. Experiments and results

We conduct two experiments comparing equivariant and non-equivariant Koopman models using DAE and eDAE architectures. We set $\gamma = \sqrt{|\mathcal{X}|/|\mathcal{Z}|}$ in all experiments to balance the error across dimensions of \mathcal{X} and \mathcal{Z} . Lastly, to test modelling error and generalization, we use testing datasets with trajectories uniformly sampled across all state space quotient sets Ω/\mathbb{G} (see [fig. 1-a](#)).

Synthetic symmetric dynamical systems with finite state symmetry groups This experiment models synthetic nonlinear \mathbb{G} -symmetric systems with arbitrary groups \mathbb{G} . The systems are constrained stable linear systems with stochastic perturbations $\mathbf{x}(\omega_{t+\Delta t}) = \mathbf{A}_{\Delta t} \mathbf{x}(\omega_t) + \boldsymbol{\epsilon}_t, \mathbf{C} \omega_t \geq \mathbf{c}$, where $\mathbf{x}(\omega) \in \mathcal{X} \subseteq \mathbb{R}^m$ is the system state’s numerical representation, $\mathbf{A}_{\Delta t} \in \mathbb{R}^{m \times m}$ the linear dynamics matrix, $\boldsymbol{\epsilon} \in \mathbb{R}^m$ a white-noise stochastic process with standard deviation σ , and $\mathbf{C} \in \mathbb{R}^{n_c \times m}, \mathbf{c} \in \mathbb{R}^{n_c}$ the parameters of n_c inequality hyper-plane constraints. These systems are \mathbb{G} -symmetric if $\mathbf{A}_{\Delta t}$ is \mathbb{G} -equivariant $\rho_{\mathcal{X}}(g) \mathbf{A}_{\Delta t} = \mathbf{A}_{\Delta t} \rho_{\mathcal{X}}(g), \forall g \in \mathbb{G}$ and any constraint is also enforced for all symmetric states $\mathbf{C}_{k,:} g \triangleright \mathbf{x}(\omega) \geq \mathbf{c}_k, \forall k \in [1, n_c], g \in \mathbb{G}$.

These synthetic systems let us assess symmetry exploitation in learning Koopman models for arbitrary groups \mathbb{G} , system’s dimensionality $|\mathcal{X}|$, latent state dimensionality $|\mathcal{Z}|$, and noise standard deviation σ . The results show that the eDAE architecture provides superior models with better sample efficiency and generalization ([fig. 2-a](#)), reduced sensitivity to the dimensionality of \mathcal{X} ([fig. 2-b](#)) and \mathcal{Z} ([fig. 2-c](#)), and improved noise robustness ([fig. 2-d](#)).

Modelling quadruped closed-loop dynamics In this experiment, we investigate using a Koopman model for robot dynamics while quantifying the impact of symmetry exploitation. The focus is on modelling the closed-loop dynamics of the mini-cheetah quadruped robot’s locomotion on mildly uneven terrain. The training data comprises a few motion trajectories executed by a \mathbb{G} -equivariant model predictive controller that tracks a desired target base velocity with a fixed trotting periodic gait ([Amatucci et al., 2022](#)). As a result, the closed-loop dynamics are \mathbb{G} -equivariant ([Ordóñez-Ápraez et al., 2023](#), eq. (3-4)) and stable, with a limit-cycle trajectory describing the gait cycle and transient dynamics governed by the controller’s correction for tracking errors. The state of the sys-




Figure 3: The test prediction mean square error (MSE) of Koopman models (DAE, DAE_{aug}, eDAE) modelling the closed-loop dynamics of locomotion of the mini-cheetah robot. We compare models exploiting the system’s full symmetry group $G = \mathbb{K}_4 \times \mathbb{C}_2$ and the subgroup $G = \mathbb{K}_4$. Solid lines and shaded areas show the mean, max, and min prediction error among 4 training seeds. (a) MSE vs. training samples. (b) MSE vs. prediction horizon. (c) MSE for varying latent space dimensionality $|Z|$. (d) MSE of measurable state observables in original units.

tem is numerically represented as $\mathbf{x}(\omega_t) = [q_{js,t}, \dot{q}_{js,t}, z_t, \mathbf{o}_t, \mathbf{v}_{err,t}, \mathbf{w}_{err,t}] \in \Omega \subseteq \mathbb{R}^{37}$, composed of the joint-space generalized position $q_{js} \in \mathcal{Q}_{js} \subseteq \mathbb{R}^{12}$, and velocity $\dot{q}_{js} \in \mathcal{T}_q \mathcal{Q}_{js} \subseteq \mathbb{R}^{12}$ coordinates, base height $z_t \in \mathbb{R}^1$, base orientation quaternion $\mathbf{o} \in \mathbb{R}^4$, and the error in the desired linear and angular base velocities $\mathbf{v}_{err,t} \in \mathbb{R}^3$ and $\mathbf{w}_{err,t} \in \mathbb{R}^3$, respectively.

The symmetry group of the robot is $G = \mathbb{K}_4 \times \mathbb{C}_2$, of order 8 (see [symmetric states](#)). This group implies the decomposition of the system’s state set into 8 copies of the quotient set of unique states Ω / G ([sec. 3](#)), and the isotypic decomposition of the physical modelling space \mathcal{X} and the space of observable functions \mathcal{F}_Z into at most 8 isotypic subspaces ([thms. 1 and 3](#)). Therefore, for this dynamical system exploiting the symmetry prior is crucial to mitigate the curse of dimensionality ([Higgins et al., 2022](#)) and biases of the training dataset (which is collected from trajectories originating from one of the 8 quotient sets; see [animation](#)). The results demonstrate superior performance of the G -equivariant Koopman models (eDAE) over the models trained with data-augmentation DAE_{aug}, and without symmetry exploitation (DAE) in terms of sample-efficiency ([fig. 3-a](#)), forecasting error ([fig. 3-b-d](#)), and robustness to hyper-parameters variation ([fig. 3-c](#)). Furthermore, models exploiting the entire symmetry group $\mathbb{K}_4 \times \mathbb{C}_2$ consistently outperform those exploiting only the subgroup \mathbb{K}_4 . The G -symmetric eDAE models excel due to their ability to capture both transient and stable locomotion dynamics from an initial configuration ω_0 , and to generalize to symmetric states $G\omega_0$, potentially unseen during training; refer to [fig. 1-a](#) and [animation](#). Moreover, by utilizing DHA, these models can decompose the learned latent linear dynamics into the linear dynamics of each isotypic subspace, thereby accurately capturing the distinct roles (and relevance) of isotypic subspaces in locomotion dynamics (see [fig. 1-b,c](#)).

6. Conclusions

We introduced the use of harmonic analysis for decomposing and understanding the dynamics of symmetric robotic systems. By partitioning the state space into isotypic subspaces, we have shown how complex motions can be characterized as the superposition of lower-dimensional, symmetric and synergistic motions. This entails the decomposition of (local and global) linear models of the system’s dynamics into independent models for each subspace. Leveraging this, we learn a data-driven global linear model using a novel equivariant deep-learning architecture to approximate the Koopman operator. The method’s practical validity is evidenced by presenting the first successful attempt to learn a linear model of the closed-loop dynamics of a quadruped robot’s locomotion.

Acknowledgments

This work was supported in part by PNRR MUR Project PE000013 CUP J53C22003010006 "Future Artificial Intelligence Research (hereafter FAIR)", funded by the European Union – NextGenerationEU

References

- Ian Abraham and Todd D Murphey. Active learning of dynamics for data-driven control using koopman operators. *IEEE Transactions on Robotics*, 35, 2019.
- Angelo Alessandri, Marco Baglietto, and Giorgio Battistelli. Receding-horizon estimation for discrete-time linear systems. *IEEE Transactions on Automatic Control*, 48, 2003.
- 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.
- H Harry Asada. Global, unified representation of heterogenous robot dynamics using composition operators: A koopman direct encoding method. *IEEE/ASME Transactions on Mechatronics*, 2023.
- Daniel Bruder, Xun Fu, R Brent Gillespie, C David Remy, and Ram Vasudevan. Data-driven control of soft robots using koopman operator theory. *IEEE Transactions on Robotics*, 37, 2020.
- Steven L. Brunton, Marko Budišić, Eurika Kaiser, and J. Nathan Kutz. Modern koopman theory for dynamical systems. *SIAM Review*, 64, 2022.
- Mildred S Dresselhaus, Gene Dresselhaus, and Ado Jorio. *Group theory: application to the physics of condensed matter*. Springer Science & Business Media, 2007.
- R. Featherstone. *Rigid Body Dynamics Algorithms*. Springer-Verlag, Berlin, Heidelberg, 2007.
- Carl Folkestad and Joel W Burdick. Koopman nmpc: Koopman-based learning and nonlinear model predictive control of control-affine systems. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- Martin Golubitsky, Ian Stewart, and David G Schaeffer. *Singularities and Groups in Bifurcation Theory: Volume II*, volume 69. Springer Science & Business Media, 2012.
- Irina Higgins, Sébastien Racanière, and Danilo Rezende. Symmetry-based representations for artificial and biological general intelligence. *Frontiers in Computational Neuroscience*, 16, 2022.
- Anthony W. Knap. *Representation Theory of Semisimple Groups, An Overview Based on Examples (PMS-36)*. Princeton University Press, Princeton, 1986.
- Marin Kobilarov, Duy-Nguyen Ta, and Frank Dellaert. Differential dynamic programming for optimal estimation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93, 2018.

- Vladimir Kostic, Pietro Novelli, Andreas Maurer, Carlo Ciliberto, Lorenzo Rosasco, and Massimiliano Pontil. Learning dynamical systems via Koopman operator regression in reproducing kernel hilbert spaces. In Advances in Neural Information Processing Systems, 2022.
- Vladimir R Kostic, Pietro Novelli, Riccardo Grazi, Karim Lounici, and Massimiliano Pontil. Deep projection networks for learning time-homogeneous dynamical systems. arXiv preprint arXiv:2307.09912, 2023.
- Cornelius Lanczos. The variational principles of mechanics. Courier Corporation, 2012.
- Andrzej Lasota and Michael C. Mackey. Chaos, Fractals, and Noise, volume 97 of Applied Mathematical Sciences. Springer New York, 1994.
- Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In First International Conference on Informatics in Control, Automation and Robotics, volume 2. SciTePress, 2004.
- Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. Nature communications, 9, 2018.
- Sergi Martinez, Robert Griffin, and Carlos Mastalli. Multi-Contact Inertial Estimation and Localization in Legged Robots. 2024.
- Carlos Mastalli, Rohan Budhiraja, Wolfgang Merkt, Guilhem Saurel, Bilal Hammoud, Maximilien Naveau, Justin Carpentier, Ludovic Righetti, Sethu Vijayakumar, and Nicolas Mansard. Crocoddyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control. In IEEE International Conference on Robotics and Automation (ICRA), 2020.
- Alexandre Mauroy, Y Susuki, and I Mezić. Koopman operator in systems and control. Springer, 2020.
- David Mayne. A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems. International Journal of Control, 3, 1966.
- Afshin Mesbahi, Jingjing Bu, and Mehran Mesbahi. On modal properties of the koopman operator for nonlinear systems with symmetry. In American Control Conference (ACC), 2019.
- Igor Mezić. Koopman operator, geometry, and learning of dynamical systems. Not. Am. Math. Soc., 68, 2021.
- Frank Noé and Feliks Nuske. A variational approach to modeling slow processes in stochastic dynamical systems. Multiscale Modeling & Simulation, 11, 2013.
- Daniel F Ordonez-Apraez, Mario Martin, Antonio Agudo, and Francesc Moreno. On discrete symmetries of robotics systems: A group-theoretic and data-driven analysis. In Proceedings of Robotics: Science and Systems, Daegu, Republic of Korea, July 2023.
- 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, 2019.

Subhrajit Sinha, Sai Pushpak Nandanoori, and Enoch Yeung. Koopman operator methods for global phase space exploration of equivariant dynamical systems. IFAC-PapersOnLine, 53, 2020.

Vivian T Steyert. Uncovering Structure with Data-driven Reduced-Order Modeling. PhD thesis, Princeton University, 2022.

Amit Surana. Koopman framework for nonlinear estimation. The Koopman Operator in Systems and Control: Concepts, Methodologies, and Applications, 2020.

Yuval Tassa, Nicolas Mansard, and Emo Todorov. Control-limited differential dynamic programming. In IEEE International Conference on Robotics and Automation (ICRA), pages 1168–1175. IEEE, 2014.

Xinglong Zhang, Wei Pan, Riccardo Scattolini, Shuyou Yu, and Xin Xu. Robust tube-based model predictive control with koopman operators. Automatica, 137, 2022.

Appendix I. Supplementary material

Supplementary material can be found in <https://danfoa.github.io/DynamicsHarmonicsAnalysis/>