

Unifying Regression and Uncertainty Quantification with Contrastive Spectral Representation Learning

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In this work, we discuss a contrastive representation learning framework, called
2 **Neural Conditional Probability (NCP)** [16], which introduces a new paradigm for
3 training deep **Neural Network (NN)** architectures for regression. This framework
4 enables high-quality regression estimates *and* parametric uncertainty quantification
5 without retraining or restrictive assumptions on the uncertainty distribution. **NCP**
6 learns high-dimensional data representations that are *linearly transferable* to re-
7 gression *and* uncertainty quantification tasks, backed by non-asymptotic statistical
8 learning guarantees linking representation quality to downstream performance.
9 Crucially, in equivariant regression contexts, the **NCP** framework can be adapted to
10 train any **Geometric Deep Learning (GDL)** architecture, resulting in a *disentangled*
11 equivariant representation learning algorithm with first-of-its-kind statistical guar-
12 antees for equivariant regression and symmetry-aware uncertainty quantification.

13 Introduction

14 A central problem in machine learning is modeling the conditional prob-
15 ability $\mathbb{P}(\mathbf{y}|\mathbf{x})$ of a target variable $\mathbf{y} \in \mathcal{Y}$ given an observed variable
16 $\mathbf{x} \in \mathcal{X}$. This underpins robust reasoning under *uncertainty* in safety-
17 critical applications such as medicine, finance, robotics, and physics
18 [12, 28, 35]. In these domains, the expected value of the target vari-
19 able $\mathbb{E}[\mathbf{y}|\mathbf{x}] := \mathbb{E}_{\mathbf{y} \sim \mathbb{P}(\mathbf{y}|\mathbf{x})}[\mathbf{y}]$ (i.e., regression) is often insufficient due
20 to high spread, multimodality, or asymmetry of the conditional distri-
21 bution (see Fig. 1). However, estimating conditional distributions in
22 high-dimensional settings remains challenging without resorting to en-
23 semble learning [10, 25] or imposing strong, often restrictive, inductive
24 biases—e.g., that $\mathbb{P}(\mathbf{y}|\mathbf{x})$ belongs to a parametric distribution family (e.g.,
25 Gaussian, Poisson, Mixture of Gaussians) for all $\mathbf{x} \in \mathcal{X}$ [1, 13, 21, 27].




Figure 1: Conditional distributions with uninformative expected values.

26 To tackle conditional probability modeling—and therefore reliable uncertainty quantification—using
27 deep **NNs**, several works in contrastive representation learning [16, 23, 26, 32] have shown that deep
28 **NNs** can be endowed with robust conditional probability and uncertainty modeling capabilities. This
29 is achieved when trained via a contrastive representation learning loss that aims to approximate the
30 **Pointwise Mutual Dependency (PMD)** ratio, $\kappa(\mathbf{x}, \mathbf{y}) = \frac{P_{\mathbf{x}\mathbf{y}}(d\mathbf{x}, d\mathbf{y})}{P_{\mathbf{x}}(d\mathbf{x})P_{\mathbf{y}}(d\mathbf{y})}$, between the joint distribution
31 $P_{\mathbf{x}\mathbf{y}}$ and the product of the marginals $P_{\mathbf{x}} \times P_{\mathbf{y}}$. The approach uses a bilinear model $\kappa_{\theta} : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$
32 (see Fig. 4), parameterized by two **NNs**, $\phi_{\theta} : \mathcal{X} \rightarrow \mathbb{R}^n$ and $\psi_{\theta} : \mathcal{Y} \rightarrow \mathbb{R}^n$, that learn high-
33 dimensional representations of the input and target variables, respectively [16].

34 In this preliminary work, we discuss how this contrastive learning framework, referred to as **NCP**,
35 could represent a paradigm shift in training deep **NNs** and **GDL** architectures for safety-critical
36 regression applications.

37 Neural Conditional Probability (NCP)

38 The **NCP** framework proposes to model conditional probability distributions by approximating the
 39 *conditional expectation operator* [2, 8, 29], $E_{y|x}: \mathcal{L}_y^2 \rightarrow \mathcal{L}_x^2$, a linear integral operator acting on the
 40 Hilbert spaces $\mathcal{L}_x^2 := \mathcal{L}_{P_x}^2(\mathcal{X}, \mathbb{R})$ and $\mathcal{L}_y^2 := \mathcal{L}_{P_y}^2(\mathcal{Y}, \mathbb{R})$ of square-integrable functions of the random
 41 variables $\mathbf{x}$ and $\mathbf{y}$, respectively. The action of this operator on any function $h \in \mathcal{L}_y^2$ returns the
 42 function’s *conditional expectation*:

$$[E_{y|x}h](\mathbf{x}) = \mathbb{E}[h(\mathbf{y})|\mathbf{x}=\mathbf{x}] := \int_{\mathcal{Y}} h(\mathbf{y})P_{y|x}(d\mathbf{y}|\mathbf{x}) = \int_{\mathcal{Y}} h(\mathbf{y}) \frac{P_{xy}(d\mathbf{y}, \mathbf{x})}{P_x(d\mathbf{x})} = \int_{\mathcal{Y}} h(\mathbf{y})\kappa(\mathbf{x}, \mathbf{y})P_y(d\mathbf{y}), \quad (1)$$

43 where the **PMD** $\kappa(\mathbf{x}, \mathbf{y}) := \frac{P_{xy}(d\mathbf{x}, d\mathbf{y})}{P_x(d\mathbf{x})P_y(d\mathbf{y})}$ becomes the kernel function defining the operator $E_{y|x}$ (see
 44 Fig. 3) and $P_{y|x}$ is the conditional probability measure [30].

45 The practical utility of $E_{y|x}$ lies in that it provides an infinite-dimensional linear model—in a nonlinear
 46 representation space—for computing conditional probabilities *and* expectations. To see this, note
 47 that for any $\mathbf{x} \in \mathcal{X}$ and any measurable set $\mathbb{B} \subset \mathcal{Y}$ we have that:

$$\mathbb{P}(\mathbf{y} \in \mathbb{B}|\mathbf{x}=\mathbf{x}) := \int_{\mathcal{Y}} \mathbb{1}_{\mathbb{B}}(\mathbf{y})P_{y|x}(d\mathbf{y}|\mathbf{x}) = [E_{y|x}\mathbb{1}_{\mathbb{B}}](\mathbf{x}), \quad \text{and} \quad \mathbb{E}[\mathbf{y}|\mathbf{x}=\mathbf{x}] := [E_{y|x}\mathbf{y}](\mathbf{x}). \quad (2)$$

48 Therefore, to provide estimates for (2), **NCP** seeks the best finite-dimensional approximation of $E_{y|x}$
 49 as a rank- r operator E_{θ} with matrix representation $\mathbf{E}_{\theta} \in \mathbb{R}^{r \times r}$ given by

$$\arg \min_{\theta} \|E_{y|x} - E_{\theta}\|_{\text{HS}}^2 \equiv \mathbb{E}_{\mathbf{x}}\mathbb{E}_{\mathbf{y}}(\kappa(\mathbf{x}, \mathbf{y}) - \kappa_{\theta}(\mathbf{x}, \mathbf{y}))^2, \quad \text{s.t. } \mathbb{E}_{\mathbf{x}}\mathbb{E}_{\mathbf{y}}\kappa_{\theta}(\mathbf{x}, \mathbf{y}) = 1 \text{ and } \text{rank}(\mathbf{E}_{\theta}) \leq r. \quad (3)$$

50 **Spectral representation learning** The optimal solution to Eq. (3), denoted $E_{\star}$, is the r -truncated
 51 **Singular Value Decomposition (SVD)** of $E_{y|x}$ [6, 15, 26], namely

$$[E_{\star}f](\mathbf{x}) = \sum_{i=1}^r \sigma_i \langle f, v_i \rangle_{P_y} u_i(\mathbf{x}), \quad \text{with} \quad \sigma_i u_i(\mathbf{x}) = [E_{y|x}v_i](\mathbf{x}), \quad \forall i \in [1, r], \quad (4)$$

52 where (σ_i, u_i, v_i) denotes the i^{th} singular value and left/right singular functions of $E_{y|x}$ [2, 15].

53 Consequently, **NCP** parameterizes E_{θ} by a bilinear model $\kappa_{\theta}(\mathbf{x}, \mathbf{y}) = 1 + \phi_{\theta}(\mathbf{x})^{\top} \mathbf{E}_{\theta} \psi_{\theta}(\mathbf{y})$, com-
 54 posed of two encoder **NNs** $\phi_{\theta}: \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_{\theta}: \mathcal{Y} \rightarrow \mathbb{R}^r$ that aim to approximate the spectral
 55 decomposition of the operator, i.e., the *span* of the top r (non-constant) left and right singular
 56 functions of $E_{y|x}$. See Fig. 4-left.

57 Since κ is generally unavailable analytically, (3) is solved via the regularized contrastive loss¹:

$$\begin{aligned} \mathcal{L}_{\gamma}(\theta) = & -2\mathbb{E}_{\mathbf{x}\mathbf{y}}\kappa_{\theta}(\mathbf{x}, \mathbf{y}) + \mathbb{E}_{\mathbf{x}}\mathbb{E}_{\mathbf{y}}\kappa_{\theta}(\mathbf{x}, \mathbf{y})^2 + 2\gamma(\|\mathbb{E}_{\mathbf{x}}\phi_{\theta}(\mathbf{x})\|_F^2 + \|\mathbb{E}_{\mathbf{y}}\psi_{\theta}(\mathbf{y})\|_F^2 \\ & + \|\text{Cov}(\phi_{\theta}) - \mathbf{I}_r\|_F^2 + \|\text{Cov}(\psi_{\theta}) - \mathbf{I}_r\|_F^2), \end{aligned} \quad (5)$$

58 where the first two regularization terms center the representations, ensuring that $\mathbb{E}_{\mathbf{x}}\mathbb{E}_{\mathbf{y}}\kappa_{\theta}(\mathbf{x}, \mathbf{y}) \approx 1$
 59 [15], while the last two enforce approximate orthonormality of the learned bases in $\mathcal{F}_{\mathbf{x}}^{\theta} :=$
 60 $\text{span}(\phi_{\theta}) \subset \mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{F}_{\mathbf{y}}^{\theta} := \text{span}(\psi_{\theta}) \subset \mathcal{L}_{\mathbf{y}}^2$ [13]—analogous to kernel methods [20].

61 Equivariant Neural Conditional Probability (eNCP)

62 In virtually all applications of **GDL** [5], symmetry priors result in the conditional $\mathbb{P}(\mathbf{y}|\mathbf{x})$ and marginal
 63 distributions $\mathbb{P}(\mathbf{x})$ and $\mathbb{P}(\mathbf{y})$ being invariant to symmetry transformations of the data (see Fig. 3), i.e:

$$\mathbb{P}(\mathbf{y}|\mathbf{x}) = \mathbb{P}(g \triangleright_{\mathcal{Y}} \mathbf{y} | g \triangleright_{\mathcal{X}} \mathbf{x}) \quad \text{and} \quad \mathbb{P}(\mathbf{x}) = \mathbb{P}(g \triangleright_{\mathcal{X}} \mathbf{x}) \quad \Rightarrow \quad \mathbb{P}(\mathbf{y}) = \mathbb{P}(g \triangleright_{\mathcal{Y}} \mathbf{y}), \quad \forall g \in \mathbb{G}, \quad (6)$$

64 where $\mathbb{G}$ denotes a compact symmetry group acting on the data spaces $\mathcal{X}$ and $\mathcal{Y}$ via the group actions,
 65 $\triangleright_{\mathcal{X}}: \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$, and $\triangleright_{\mathcal{Y}}: \mathbb{G} \times \mathcal{Y} \rightarrow \mathcal{Y}$, with $g \triangleright_{\mathcal{X}} \mathbf{x} \in \mathcal{X}$ and $g \triangleright_{\mathcal{Y}} \mathbf{y} \in \mathcal{Y}$ denoting linear, invertible
 66 transformations of $\mathbf{x}$ and $\mathbf{y}$ defined by $g \in \mathbb{G}$ (see Fig. 3).

67 These priors imply the $\mathbb{G}$ -invariance of the joint distribution $\mathbb{P}(\mathbf{x}, \mathbf{y})$, as well as the $\mathbb{G}$ -equivariance
 68 of conditional expectations (i.e., equivariant regression) and the $\mathbb{G}$ -invariance of the **PMD** kernel
 69 defining the conditional expectation operator (see Fig. 3 and [23]):

$$g \triangleright_{\mathcal{Y}} \mathbb{E}[\mathbf{y}|\mathbf{x}=\mathbf{x}] = \mathbb{E}[\mathbf{y}|\mathbf{x}=g \triangleright_{\mathcal{X}} \mathbf{x}] \quad \text{and} \quad \kappa(\mathbf{x}, \mathbf{y}) = \kappa(g \triangleright_{\mathcal{X}} \mathbf{x}, g \triangleright_{\mathcal{Y}} \mathbf{y}) \quad \forall g \in \mathbb{G}, \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}. \quad (7)$$

70 These symmetry priors represent powerful inductive biases for the **NCP** framework, summarized in
 71 the following and in [23].

¹Used in density-ratio fitting [30], representation learning [9, 34], and mutual information estimation [31].

Task	$f(x) := \mathbb{E}_y[y x=x] \approx \hat{f}_\theta(x)$	$\mathbb{P}[y \in \mathbb{B} x \in \mathbb{A}] \approx \hat{\mathbb{P}}_\theta[y \in \mathbb{B} x \in \mathbb{A}]$
Estimate	$\hat{\mathbb{E}}_y[y] + \phi_\theta(x)^\top \mathbf{E}_\theta \hat{\mathbb{E}}_y[\psi_\theta(y) \otimes y]$	$\hat{\mathbb{E}}_y[\mathbb{1}_\mathbb{B}] + \frac{\hat{\mathbb{E}}_x[\mathbb{1}_\mathbb{A}(x) \otimes \phi_\theta(x)]^\top \mathbf{E}_\theta \hat{\mathbb{E}}_y[\mathbb{1}_\mathbb{B}(y) \otimes \psi_\theta(y)]}{\hat{\mathbb{E}}_x[\mathbb{1}_\mathbb{A}(x)]}$
Guarantees	$\ f - \hat{f}_\theta\ _{\mathcal{L}_x^2} \lesssim \sqrt{\text{Var}[\ y\]} \left(\mathcal{E}_\theta^r + \frac{\ln(n_{\text{iso}}/\delta)}{(d_{\text{iso}}N)^{\frac{\alpha}{1+2\alpha}}} \right)$	$ \mathbb{P} - \hat{\mathbb{P}}_\theta \lesssim \sqrt{\frac{\mathbb{P}[y \in \mathbb{B}]}{\mathbb{P}[x \in \mathbb{G} \triangleright \mathbb{A}]}} \left(\mathcal{E}_\theta^r + \frac{\ln(n_{\text{iso}}/\delta)}{(d_{\text{iso}}N)^{\frac{\alpha}{1+2\alpha}}} \right)$

Table 1: Statistical guarantees for eNCP [23]. The error bounds are shaped by (i) the structure of the symmetry group $\mathbb{G}$ —the number of isotypic subspaces n_{iso} and the sum of the group’s irreducible representations dimension $d_{\text{iso}} = \sum_{k=1}^{n_{\text{iso}}} d_k$, which enlarge the effective sample size—, (ii) the quality of the learned representations $\mathcal{E}_\theta^r = \|\mathbb{E}_{y|x} - \mathbf{E}_\theta\|_{\text{op}} \leq \sqrt{\mathcal{L}_\gamma(\theta) - \mathcal{L}_\gamma(\star)}$, and (iii) the operator’s singular-value decay rate $\alpha > 0$. Note that $\mathbb{G} \triangleright_{\mathcal{X}} \mathbb{A} := \bigcup_{g \in \mathbb{G}} g \triangleright_{\mathcal{X}} \mathbb{A}$ denotes the group orbit of $\mathbb{A}$. NCP’s guarantees are recovered when $n_{\text{iso}} = d_{\text{iso}} = 1$.

Symmetric function spaces A significant property of the NCP and eNCP frameworks is their non-asymptotic statistical learning guarantees, summarized in Tab. 1

The $\mathbb{G}$ -invariance of the marginal probabilities implies that the function spaces $\mathcal{L}_x^2$ and $\mathcal{L}_y^2$ are symmetric Hilbert spaces with unitary group actions $\triangleright_{\mathcal{L}_x^2} : \mathbb{G} \times \mathcal{L}_x^2 \rightarrow \mathcal{L}_x^2$ and $\triangleright_{\mathcal{L}_y^2} : \mathbb{G} \times \mathcal{L}_y^2 \rightarrow \mathcal{L}_y^2$ (see details in App. F and Fig. 11). Consequently, whenever $\mathbb{G}$ is a compact symmetry group, a standard result from harmonic analysis [18] enables the orthogonal decomposition of these spaces 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 Thm. E.8), where each $\mathcal{L}_x^{2(k)}$ and $\mathcal{L}_y^{2(k)}$ denotes the subspaces that transform according to a unique quotient group $\mathbb{G}^{(k)} \subseteq \mathbb{G}$.

Equivariant conditional expectation operator Due to Eq. (7), $\mathbb{E}_{y|x}$ becomes a $\mathbb{G}$ -equivariant linear operator [23]. This means that $\mathbb{E}_{y|x}$ commutes with the group action on the function spaces; $g \triangleright_{\mathcal{L}_x^2} [\mathbb{E}_{y|x} h](\cdot) = \mathbb{E}_{y|x} [g \triangleright_{\mathcal{L}_y^2} h](\cdot)$ for all $g \in \mathbb{G}, h \in \mathcal{L}_y^2$, and consequently, can be decomposed (disentangled) into a direct sum of operators acting on the corresponding isotypic subspaces $[\mathbb{E}_{y|x} h](\cdot) = \sum_{k=1}^{n_{\text{iso}}} [\mathbb{E}_{y|x}^{(k)} h^{(k)}](\cdot)$, for all $h \in \mathcal{L}_y^2$ and $h^{(k)} \in \mathcal{L}_y^{2(k)}$, where each $\mathbb{E}_{y|x}^{(k)} : \mathcal{L}_y^{2(k)} \rightarrow \mathcal{L}_x^{2(k)}$ models $\mathbb{G}^{(k)}$ -equivariant conditional expectation.

Disentangled spectral representation learning The orthogonal decomposition of the function spaces and of $\mathbb{E}_{y|x}$ decomposes the representation learning objective in (3) into n_{iso} independent (disentangled) representation learning subproblems:

$$\begin{aligned} \arg \min_{\theta} \quad & \|\mathbb{E}_{y|x} - \mathbf{E}_\theta\|_{\text{HS}}^2 = \sum_{k=1}^{n_{\text{iso}}} \|\mathbb{E}_{y|x}^{(k)} - \mathbf{E}_\theta^{(k)}\|_{\text{HS}}^2 \equiv \mathbb{E}_x \mathbb{E}_y \sum_{k=1}^{n_{\text{iso}}} (\kappa^{(k)}(x, y) - \kappa_\theta^{(k)}(x, y))^2, \\ \text{s.t.} \quad & \mathbb{E}_x \mathbb{E}_y \kappa_\theta(x, y) = 1, \quad \text{and} \quad \kappa_\theta(g \triangleright_{\mathcal{X}} x, g \triangleright_{\mathcal{Y}} y) = \kappa_\theta(x, y), \quad \forall g \in \mathbb{G}, (x, y) \in \mathcal{X} \times \mathcal{Y}. \end{aligned} \quad (8)$$

Analogous to (4), the optimal truncation of each $\mathbb{E}_{y|x}^{(k)}$ is its truncated SVD [6]. Consequently, as explained in [23], Eq. (8) can be solved by parameterizing the approximated PMD as a bilinear model

$$\kappa_\theta(x, y) = \mathbb{1}_{P_x}(x) \mathbb{1}_{P_y}(y) + \sum_{k=1}^{n_{\text{iso}}} \kappa_\theta^{(k)}(x, y), \quad \kappa_\theta^{(k)}(x, y) := \phi_\theta^{(k)}(x)^\top \mathbf{E}_\theta^{(k)} \psi_\theta^{(k)}(y), \quad (9)$$

parameterized by two $\mathbb{G}$ -equivariant encoder NNs $\phi_\theta : \mathcal{X} \rightarrow \mathbb{R}^r$ and $\psi_\theta : \mathcal{Y} \rightarrow \mathbb{R}^r$, defined to expose the approximated isotypic subspaces, i.e., $\phi_\theta(\cdot) = [\phi_\theta^{(1)}(\cdot)^\top, \dots, \phi_\theta^{(n_{\text{iso}})}(\cdot)^\top]^\top$ and $\psi_\theta(\cdot) = [\psi_\theta^{(1)}(\cdot)^\top, \dots, \psi_\theta^{(n_{\text{iso}})}(\cdot)^\top]^\top$, and a $\mathbb{G}$ -equivariant truncated operator matrix parameterized in block-diagonal form $\mathbf{E}_\theta = \bigoplus_{k=1}^{n_{\text{iso}}} \mathbf{E}_\theta^{(k)}$ (see Fig. 4).

Statistical learning guarantees A significant property of the NCP and eNCP frameworks is their non-asymptotic statistical learning guarantees, summarized in Tab. 1, which provide first-of-its-kind learning guarantees for equivariant representations in the context of equivariant regression and conditional probability estimation. Crucially, these guarantees enable direct estimation of the benefits of leveraging symmetry priors for any finite symmetry group $\mathbb{G}$.

Experiments

We present two experiments evaluating the NCP and eNCP frameworks in $\mathbb{G}$ -equivariant regression and symmetry-aware uncertainty quantification.

$\mathbb{G}$ -Equivariant regression The goal is to predict a quadruped robot’s Center of Mass (CoM) linear $l \in \mathbb{R}^3$ and angular momenta $k \in \mathbb{R}^3$ given noisy observations of the robot’s generalized




Figure 2: **Left:** Test set sample efficiency for $\mathbb{G}$ -equivariant regression (MSE vs. training samples) when predicting linear and angular momentum of a quadruped robot’s center of mass (CoM) from noisy joint positions and velocities. **Right:** Uncertainty quantification via prediction of 90% confidence intervals (CI, light-red area) for the robot’s instantaneous work U_t and kinetic energy T_t during locomotion over rough terrain.

positions $\mathbf{q} \in \mathbb{R}^{12}$ and velocity coordinates $\dot{\mathbf{q}} \in \mathbb{R}^{12}$. We compare eNCP against NCP and two baselines—a standard Multi-Layer Perceptron (MLP) and an Equivariant MLP (eMLP)—all with equivalent architectural footprint, where MLP and eMLP are trained using Mean Squared Error (MSE) loss. The results in Fig. 2 show that NCP can beat MLP in terms of performance, while eNCP outperforms eMLP in both performance and sample complexity.

Symmetry-aware uncertainty quantification The goal is to provide uncertainty quantification for unavailable yet crucial state observables for robot control [4, 19]. Specifically, we use proprioceptive sensor readings to provide 90% confidence intervals for the robot’s Ground Reaction Forces (GRF) $\boldsymbol{\tau}_{\text{grf}} \in \mathbb{R}^{12}$, instantaneous work $U(\mathbf{q}, \dot{\mathbf{q}}, \boldsymbol{\tau}) \in \mathbb{R}$, and kinetic energy $T(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}$ while the robot traverses rough terrain. Robust estimates of these quantities are crucial for optimal control [4].

For eNCP and NCP models, quantile estimation is achieved by regressing the Conditional Cumulative Distribution Function (CCDF) for each dimension of $\mathbf{y} = [y_1, \dots]$ and applying linear search to extract quantiles (see Fig. 7 in the appendix). This involves discretizing each y_i range into N_b bins and estimating $\mathbb{P}(y_i \in \mathbb{A}_{i,n} | \mathbf{x} = \cdot) := [\mathbb{E}_{\mathbf{y}|\mathbf{x}} \mathbb{1}_{\mathbb{A}_{i,n}}](\cdot)$ for all $n \in [N_b]$ (see Tab. 1) in a single forward pass, where $\mathbb{A}_{i,n}$ consists of the first n bins. In contrast, baseline methods Conditional Quantile Regression (CQR) [7] and its equivariant adaptation Equivariant CQR (eCQR) directly regress quantiles for fixed coverage levels and require retraining for different coverage values.

	r-Coverage $\uparrow$	Coverage $\uparrow$
eNCP	99.5 $\pm$ 0.1%	95.0 $\pm$ 0.4%
NCP	99.5 $\pm$ 0.0%	56.9 $\pm$ 0.3%
eCQR	84.2 $\pm$ 0.7%	6.7 $\pm$ 1.2%
CQR	80.5 $\pm$ 3.7%	8.5 $\pm$ 0.9%

Table 2: Relaxed coverage, see (13), and Coverage, see (12), for the test-set confidence intervals in quadruped locomotion uncertainty estimation of $\boldsymbol{\tau}_{\text{grf}}$, U and T . Target coverage is: 90%.

Results in Tab. 2, Fig. 2 show NCP and eNCP outperforming their respective baselines, with eNCP being the sole method with empirical test coverage close to the desired value, making other models unreliable for practical applications.

Conclusions

This work discusses the NCP contrastive representation learning framework for training NN architectures for regression, enabling high-quality regression estimates and uncertainty quantification via conditional probability estimation, without requiring retraining or assumptions on the conditional distribution shape. NCP can be extended to leverage symmetry priors, resulting in a disentangled equivariant representation learning framework capable of training any equivariant NN architecture for equivariant regression and symmetry-aware uncertainty quantification with novel non-asymptotic statistical learning guarantees. We demonstrate this framework’s capabilities to replace standard training approaches on two preliminary robotics applications.

Discussion and future work From a practical perspective, the NCP framework simply requires changing the training loss from the standard MSE to the contrastive loss (5) and using an auxiliary NN model during training (see Fig. 4). Extensive experimental work—similar to that presented in [3] in the context of classification—is needed to elucidate in which applications the NCP framework can and cannot outperform standard training approaches and which engineering choices achieve optimal performance (e.g., batch size, orthonormalization technique, etc.).

References

- [1] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion*, 76:243–297, 2021.
- [2] Charles R Baker. Joint measures and cross-covariance operators. *Trans. Am. Math. Soc.*, 186: 273–273, 1973.
- [3] Han Bao, Yoshihiro Nagano, and Kento Nozawa. On the surrogate gap between contrastive and supervised losses. In *International conference on machine learning*, pages 1585–1606. PMLR, 2022.
- [4] 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.
- [5] 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.
- [6] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [7] 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.
- [8] Kenji Fukumizu, Francis R Bach, and Michael I Jordan. Dimensionality reduction for supervised learning with reproducing kernel hilbert spaces. *Journal of Machine Learning Research*, 5(Jan): 73–99, 2004.
- [9] 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.
- [10] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*.
- [11] 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.
- [12] Rafael Izbicki. *Machine Learning Beyond Point Predictions: Uncertainty Quantification*. imprint, 1st edition, 2025. ISBN 978-65-01-20272-3.
- [13] Rafael Izbicki and Ann B. Lee. Converting high-dimensional regression to high-dimensional conditional density estimation. 2017.
- [14] Anthony W. Knap. *Representation Theory of Semisimple Groups, An Overview Based on Examples (PMS-36)*. Princeton University Press, Princeton, 1986.
- [15] Vladimir R Kostic, Pietro Novelli, Riccardo Grazi, Karim Lounici, and Massimiliano Pontil. Learning invariant representations of time-homogeneous stochastic dynamical systems. In *The Twelfth International Conference on Learning Representations*, 2024.
- [16] Vladimir R Kostic, Gregoire Pacreau, Giacomo Turri, Pietro Novelli, Karim Lounici, and Massimiliano Pontil. Neural conditional probability for uncertainty quantification. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [17] Zicheng Liu, Steven J Gortler, and Michael F Cohen. Hierarchical spacetime control. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 35–42, 1994.

- [18] 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.
- [19] 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.
- [20] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1–2):1–141, 2017. URL <http://dx.doi.org/10.1561/22000000060>.
- [21] 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>.
- [22] 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.
- [23] Daniel Ordoñez-Apraez, Vladimir Kostić, Alek Fröhlich, Vivien Brandt, Karim Lounici, and Massimiliano Pontil. Equivariant representation learning for symmetry-aware inference with guarantees. *arXiv preprint arXiv:2505.19809*, 2025.
- [24] 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*, 0(0): 02783649241282422, 0. doi: 10.1177/02783649241282422. URL <https://doi.org/10.1177/02783649241282422>.
- [25] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019.
- [26] 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.
- [27] 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>.
- [28] Ralph C. Smith. *Uncertainty Quantification: Theory, Implementation, and Applications*. Society for Industrial and Applied Mathematics, January 2013. URL <http://dx.doi.org/10.1137/1.9781611973228>.
- [29] Le Song, Jonathan Huang, Alex Smola, and Kenji Fukumizu. Hilbert space embeddings of conditional distributions with applications to dynamical systems. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 961–968, 2009.
- [30] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- [31] 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.
- [32] 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.
- [33] Giulio Turrisi, Valerio Modugno, Lorenzo Amatucci, Dimitrios Kanoulas, and Claudio Semini. On the benefits of gpu sample-based stochastic predictive controllers for legged locomotion. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13757–13764, 2024. doi: 10.1109/IROS58592.2024.10801698.

- 243 [34] Ziyu Wang, Yucen Luo, Yueru Li, Jun Zhu, and Bernhard Schölkopf. Spectral representation
244 learning for conditional moment models. *arXiv preprint arXiv:2210.16525*, 2022.
- 245 [35] Larry Wasserman. *All of Nonparametric Statistics*. Springer Texts in Statistics. Springer, New
246 York, NY, 1 edition, May 2007.
- 247 [36] Maurice Weiler, Patrick Forré, Erik Verlinde, and Max Welling. *Equivariant and Coordinate In-*
248 *dependent Convolutional Networks*. World Scientific, 2023. URL [https://maurice-weiler.](https://maurice-weiler.gitlab.io/cnn_book/EquivariantAndCoordinateIndependentCNNs.pdf)
249 [gitlab.io/cnn_book/EquivariantAndCoordinateIndependentCNNs.pdf](https://maurice-weiler.gitlab.io/cnn_book/EquivariantAndCoordinateIndependentCNNs.pdf).

250

251 **Appendix**

252

253 **Table of Contents**

254	A Symbols and notation	9
255	B Symmetry constraints on conditional expectations	10
256	C $\mathbb{G}$-Equivariant bilinear NN architecture	10
257	D Experimental setup	11
258	E Background on group and representation theory	18
259	F Representation theory of symmetric function spaces	21

260

261

	Numbers and Arrays
x	A scalar, or scalar function $x(\cdot)$
$\mathbf{x}$	A vector, or vector-valued function $\mathbf{x}(\cdot)$
$\mathbf{x}_1 \oplus \mathbf{x}_2$	Direct sum (stacking) of vectors, such that $\mathbf{x}_1 \oplus \mathbf{x}_2 := \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix}$
$\mathbf{K}$	A matrix
$\mathbf{A} \oplus \mathbf{B}$	Direct sum of matrices, such that $\mathbf{A} \oplus \mathbf{B} := \begin{bmatrix} \mathbf{A} & \mathbf{O} \\ \mathbf{O} & \mathbf{B} \end{bmatrix}$
$\mathbf{K}$	A linear operator
$\mathbf{I}$	Identity matrix
$\delta_{i,j}$	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
$\mathbb{R}, \mathbb{C}$	The set of real and complex numbers
$\mathcal{X} \oplus \mathcal{Y}$	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 $\mathbf{x} \oplus \mathbf{y} \in \mathcal{X} \oplus \mathcal{Y}$
$\mathcal{L}_{P_{\mathbf{x}}}^2 := \mathcal{L}_{P_{\mathbf{x}}}^2(\mathcal{X}, \mathbb{R})$	The Hilbert space of square-integrable functions on $\mathcal{X}$ with respect to the measure $P_{\mathbf{x}}$, defined as $\mathcal{L}_{P_{\mathbf{x}}}^2(\mathcal{X}) := \{f \mid \int_{\mathcal{X}} f(\mathbf{x}) ^2 P_{\mathbf{x}}(d\mathbf{x}) < \infty\}$
$\langle f, f' \rangle_{P_{\mathbf{x}}}$	Inner product between f and f' in $\mathcal{L}_{P_{\mathbf{x}}}^2 \mathcal{X}$, defined as $\langle f, f' \rangle_{P_{\mathbf{x}}} := \int_{\mathcal{X}} f(\mathbf{x}) f'(\mathbf{x}) P_{\mathbf{x}}(d\mathbf{x})$

Group and Representation theory

$\mathbb{G}$	A symmetry group
g, g_1, g_a	A symmetry group element
$g \triangleright \mathbf{x}$	The (left) group action of g on $\mathbf{x}$ defined by $g \triangleright \mathbf{x} := \rho_{\mathcal{X}}(g)\mathbf{x}$
$\rho_{\mathcal{X}}$	A representation of the group $\mathbb{G}$ on the vector space $\mathcal{X}$, defined for a chosen basis of $\mathcal{X}$
$\bar{\rho}_k$	An irreducible representation Def. E.7 of the group $\mathbb{G}$
$\rho_{\mathcal{X}}(g)$	Representation of the group element g on the vector space $\mathcal{X}$, defined for a chosen basis $\mathcal{X}$
$\rho_{\mathcal{X}} \oplus \rho_{\mathcal{Y}}$	Direct sum of group representations, such that $\rho_{\mathcal{X}}(g) \oplus \rho_{\mathcal{Y}}(g) := \begin{bmatrix} \rho_{\mathcal{X}}(g) & \\ & \rho_{\mathcal{Y}}(g) \end{bmatrix}$
$\mathbb{G}\mathbf{x}$	The group orbit of $\mathbf{x}$, defined as $\mathbb{G}\mathbf{x} := \{g \triangleright \mathbf{x} \mid g \in \mathbb{G}\}$
$\gamma_{\mathbb{G}'}(A)$	The symmetry index of a set $A \subseteq \mathcal{X}$ w.r.t. probability distribution on $\mathcal{X}$ and group elements $\mathbb{G}' \subseteq \mathbb{G}$
$\mathbb{G}_a \times \mathbb{G}_b$	Direct product of groups $\mathbb{G}_a$ and $\mathbb{G}_b$
$\mathbf{U}(\mathcal{X})$	Unitary group on the vector space $\mathcal{X}$
$\mathbf{GL}(\mathcal{X})$	General Linear group on the vector space $\mathcal{X}$, a.k.a the space of invertible matrices in $\mathbb{R}^{ \mathcal{X} \times \mathcal{X} }$
$\mathbb{C}_n$	Cyclic group of order n
$\mathbb{K}_4$	Klein four-group

Probability Theory

$\mathbf{x} \sim \mathbb{P}(\mathbf{x})$	Random vector $\mathbf{x} \in \mathcal{X}$ has distribution $\mathbb{P}(\mathbf{x})$
$P_{\mathbf{x}}$	A probability measure on the space $\mathcal{X}$
$\mathbb{E}_{\mathbf{x}}[f(\mathbf{x})]$	Expectation of $f(\mathbf{x})$ with respect to $P_{\mathbf{x}}$
$\text{Cov}(f(\mathbf{x}))$	Variance of $f(\mathbf{x})$ with respect to $P_{\mathbf{x}}$, define as $\mathbb{E}_{\mathbf{x}}(f(\mathbf{x}) - \mathbb{E}_{\mathbf{x}}f(\mathbf{x}))^2$
$\text{Cov}(f(\mathbf{x}), h(\mathbf{y}))$	Covariance of $f(\mathbf{x})$ and $h(\mathbf{y})$ with respect to the joint distribution $P_{\mathbf{xy}}$, defined as $\mathbb{E}_{\mathbf{xy}}(f(\mathbf{x}) - \mathbb{E}_{\mathbf{x}}f(\mathbf{x}))(h(\mathbf{y}) - \mathbb{E}_{\mathbf{y}}h(\mathbf{y}))$
$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian distribution over $\mathbf{x}$ with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$

264 B Symmetry constraints on conditional expectations

265 Under the assumed symmetry priors in (6) the conditional expectation of $\mathbf{y}$ is a $\mathbb{G}$ -equivariant
 266 function/map. This property is depicted in Fig. 3-center and proved in the following proposition.

267 **Proposition B.1** ($\mathbb{G}$ -equivariant conditional expectations). *Let $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ be two vector*
 268 *valued random variables satisfying the symmetry priors of Eq. (6). Then, the conditional expectation*
 269 *of $\mathbf{y}$ given $\mathbf{x}$ is $\mathbb{G}$ -equivariant, since, for every $g \in \mathbb{G}$, $\mathbf{x} \in \mathcal{X}$,*

$$\begin{aligned} \mathbb{E}[\mathbf{y}|\mathbf{x} = g \triangleright_{\mathcal{X}} \mathbf{x}] &= g \triangleright_{\mathcal{Y}} \mathbb{E}[\mathbf{y}|\mathbf{x} = \mathbf{x}] \\ &= \int_{\mathcal{Y}} g \triangleright_{\mathcal{Y}} \mathbf{y} P_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|\mathbf{x}) \\ &= \int_{\mathcal{Y}} \mathbf{y} P_{\mathbf{y}|\mathbf{x}}(g^{-1} \triangleright_{\mathcal{Y}} d\mathbf{y}|\mathbf{x}) \\ &= \int_{\mathcal{Y}} \mathbf{y} P_{\mathbf{y}|\mathbf{x}}(d\mathbf{y}|g \triangleright_{\mathcal{X}} \mathbf{x}) \quad (\text{by Eq. (6)}) \\ &= \mathbb{E}[\mathbf{y}|\mathbf{x} = g \triangleright_{\mathcal{X}} \mathbf{x}]. \end{aligned}$$




Figure 3: Example of symmetric random variables $(\mathbf{x}, \mathbf{y}) \sim \mathcal{X} \times \mathcal{Y} \subset \mathbb{R} \times \mathbb{R}$, whose marginals $\mathbb{P}(\mathbf{x})$ and $\mathbb{P}(\mathbf{y})$; joint $\mathbb{P}(\mathbf{x}, \mathbf{y})$; and conditional $\mathbb{P}(\mathbf{y}|\mathbf{x})$ distributions are invariant to reflections of the data: $g_r \triangleright_{\mathcal{X}} \mathbf{x} = -\mathbf{x}$ and $g_r \triangleright_{\mathcal{Y}} \mathbf{y} = -\mathbf{y}$, where g_r denotes the reflection element of the reflection symmetry group $\mathbb{C}_2 := \{e, g_r | g_r^2 = e\}$. Consequently, the PMD $\kappa(\mathbf{x}, \mathbf{y})$ is $\mathbb{C}_2$ -invariant.

270 C $\mathbb{G}$ -Equivariant bilinear NN architecture

271 This section outlines how to construct a $\mathbb{G}$ -equivariant disentangled representation for the random
 272 variables $\mathbf{x}$ and $\mathbf{y}$ using **any** type of $\mathbb{G}$ -equivariant NN architecture backbone, such as MLP, CNNs,
 273 Transformers, and others.

274 Let $\mathbf{f}_{\theta} : \mathcal{X} \mapsto \mathbb{R}^r$ and $\mathbf{h}_{\theta} : \mathcal{Y} \mapsto \mathbb{R}^r$ be two $\mathbb{G}$ -equivariant NNs, whose outputs will be interpreted
 275 as the basis functions of the truncated symmetric function spaces $\mathcal{F}_{\mathbf{x}} \subset \mathcal{L}_{\mathbf{x}}^2$ and $\mathcal{F}_{\mathbf{y}} \subset \mathcal{L}_{\mathbf{y}}^2$. Assume,
 276 the group representations on $\mathcal{F}_{\mathbf{x}}$ and $\mathcal{F}_{\mathbf{y}}$ are constructed from multiplicities of the group's regular
 277 representation, $\rho_{\mathcal{F}_{\mathbf{x}}} = \bigoplus_{n=1}^{r/|\mathbb{G}|} \rho_{\text{reg}}$ and $\rho_{\mathcal{F}_{\mathbf{y}}} = \bigoplus_{n=1}^{r/|\mathbb{G}|} \rho_{\text{reg}}$ —as done usually in practice [36]. Since
 278 for (most) finite groups, the decomposition of ρ_{reg} into *irreps* is known or can be computed, we
 279 have access to the analytical change of basis $\mathbf{Q}_{\mathbf{x}} : \mathcal{F}_{\mathbf{x}} \mapsto \mathcal{F}_{\mathbf{x}}$ and $\mathbf{Q}_{\mathbf{y}} : \mathcal{F}_{\mathbf{y}} \mapsto \mathcal{F}_{\mathbf{y}}$ to transition to
 280 the isotypic basis. Consequently, we can directly parameterize the representations of the random
 281 variables in disentangled form as:

$$\phi_{\theta}(\cdot) = \mathbf{Q}_{\mathbf{x}}^{\top}(\mathbf{f}_{\theta}(\cdot) - \mathbb{E}_{\mathbf{x}}[\mathbf{f}_{\theta}(\mathbf{x})]), \quad \psi_{\theta}(\cdot) = \mathbf{Q}_{\mathbf{y}}^{\top}(\mathbf{h}_{\theta}(\cdot) - \mathbb{E}_{\mathbf{y}}[\mathbf{h}_{\theta}(\mathbf{y})]). \quad (10)$$

282 Given that during training these representations are not orthogonal, the truncated operator is param-
 283 eterized as the trainable $\mathbb{G}$ -equivariant matrix $\mathbf{E}_{\theta} = \bigoplus_k^{n_{\text{iso}}} \mathbf{E}_{\theta}^{(k)} = \bigoplus_k^{n_{\text{iso}}} \mathbf{O}^{(k)} \otimes \mathbf{I}_{d_k}$ with parameters

284 $\{\mathbf{O}^{(k)} \in \mathbb{R}^{m_k \times m_k}\}_{k=1}^{n_{\text{iso}}}$. Hence, the kernel of each the truncated operator is given in terms of the
 285 model free parameters by:

$$\kappa_{\theta}(\mathbf{x}, \mathbf{y}) = \mathbb{1}_{P_{\mathbf{x}}}(\mathbf{x})\mathbb{1}_{P_{\mathbf{y}}}(\mathbf{y}) + \sum_{k=1}^{n_{\text{iso}}} \sum_{s,t} \mathbf{O}_{s,t}^{(k)} \sum_{i,j} \phi_{s,i}^{\theta^{(k)}}(\mathbf{x}) \psi_{t,j}^{\theta^{(k)}}(\mathbf{y}). \quad (11)$$

286 Note that after training, the **SVD** of the learned operator can be computed by exploiting the constraints
 287 imposed by the operator’s $\mathbb{G}$ -equivariance (see Fig. 4). Importantly, once changed to the spectral
 288 basis, the group action on the approximated spectral basis matches that on the isotypic basis.




Figure 4: **Left:** **NCP**’s bilinear **NN** architecture. **Right:** **eNCP**’s $\mathbb{G}$ -equivariant bilinear **NN** architecture, featuring ϕ_{θ} and ψ_{θ} as $\mathbb{G}$ -equivariant **NNs** and $\mathbf{E}_{\theta}$ as a $\mathbb{G}$ -equivariant block-diagonal matrix. Each block is equivariant to a quotient group $\mathbb{G}^{(k)} \subseteq \mathbb{G}$ and is constrained to have singular spaces of dimension at least d_k —the dimension of the irreducible representation of type k .

289 D Experimental setup

290 In this section we provide details on the experimental setup. We first describe general design choices
 291 and hyperparameters and then provide details for each experiment.

292 **Sample efficiency experiments** For both the conditional expectation operator approximation and
 293 the $\mathbb{G}$ -equivariant regression experiments, we evaluate model performance by measuring sample
 294 efficiency/complexity. To do so, we partition the dataset $\mathbb{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$ into training, validation,
 295 and testing splits in proportions of 70%, 15%, and 15%, respectively. With fixed validation and
 296 testing sets, we iteratively train the models on increasing portions of the training set and report the
 297 test performance for each size.

298 For each training set size, we select the model checkpoint with the best validation loss to compute
 299 the test performance. Thus, these experiments quantify the generalization error (or true risk) and its
 300 evolution as a function of the training set size.

301 **NNs architectures and hyperparameters** To compare our equivariant representation learning
 302 framework with other contrastive and supervised methods, all (inference) models share a similar fixed
 303 architectural footprint. For the baseline models, the only hyperparameter tuned is the learning rate,
 304 whereas for the **NCP** and **eNCP** models we additionally tune the regularization weight γ in Eq. (5).
 305 Further details for each experiment are provided in the corresponding sections below.

306 **Code reproducibility** All experiments, plots and examples are provided in the open-access reposi-
 307 tory and python package [symm_rep_learn](#).

308 D.1 $\mathbb{G}$ -equivariant regression of robot’s CoM momenta

309 In this experiment, we evaluate the quality of the learned representations using the contrastive loss
 310 Eq. (5) alongside supervised learning baselines trained with the standard **MSE** loss. The task is
 311 a $\mathbb{G}$ -equivariant benchmark in robotics presented in [24], with the goal of predicting a quadruped
 312 robot’s **CoM** linear $\mathbf{l} \in \mathbb{R}^3$ and angular momenta $\mathbf{k} \in \mathbb{R}^3$ from noisy observations of the robot’s
 313 generalized positions $\mathbf{q} \in \mathbb{R}^{12}$ and velocity coordinates $\dot{\mathbf{q}} \in \mathbb{R}^{12}$. Consequently, the random variables
 314 are defined as $\mathbf{x} = \mathbf{q} + \epsilon_{\mathbf{q}} \oplus \dot{\mathbf{q}} + \epsilon_{\dot{\mathbf{q}}}$ and $\mathbf{y} = \mathbf{l} \oplus \mathbf{k}$, where $\epsilon_{\mathbf{q}} \in \mathbb{R}^{12}$ and $\epsilon_{\dot{\mathbf{q}}} \in \mathbb{R}^{12}$ are independent
 315 Gaussian noise terms that model sensor noise. The function computing the **CoM** momenta from these
 316 proprioceptive observations is highly non-linear and $\mathbb{G}$ -equivariant whenever $\mathbb{G}$ is a morphological
 317 symmetry group of the robot (see Fig. 5 and [24] for details).

318 The robot considered is the quadruped robot Solo (Fig. 5-right), which possesses a symmetry group
 319 of order 8: $\mathbb{G} = \mathbb{K}_4 \times \mathbb{C}_2$, as depicted in this animation showing 8 symmetric robot configurations
 320 along with their corresponding linear and angular momenta vectors.

321 **NN architectures** We configure all models under consideration (eNCP, NCP, eMLP, and MLP) to
 322 have an inference-time NN architecture with a similar footprint. In particular, the encoder network
 323 for $\mathbf{x}$ in NCP and eNCP is designed similarly to the NN used in MLP/eMLP. The idea is to test how
 324 a model with the same capacity performs on the downstream task of classification when trained using
 325 either the representation learning loss or a supervised learning loss. The backbone of all architectures
 326 is a standard multilayer perceptron consisting of three hidden layers, each with 512 units, followed
 327 by a final hidden layer containing 128 units. This final layer encodes the feature vector r for the NCP
 328 and eNCP models. Crucially, since $\mathbb{G}$ -equivariance enforces weight sharing in the NN architecture,
 329 the encoder NN for eNCP and eMLP comprises $\times 8$ fewer parameters than their symmetry-agnostic
 330 counterparts.




Figure 5: Example of morphological finite symmetry in robotics. **Left:** A humanoid robot with the reflectional symmetry group $\mathbb{G} \equiv \mathbb{C}_2$. **Right:** The quadruped robot Solo with the symmetry group $\mathbb{G} = \mathbb{K}_4 \times \mathbb{C}_2$ (only $\mathbb{K}_4$ is shown for clarity). The robot’s center of mass linear $\mathbf{l} \in \mathbb{R}^3$ and angular $\mathbf{k} \in \mathbb{R}^3$ momentum are depicted as orange and green vectors, respectively, for each symmetric configuration. Images adapted from Ordoñez-Apaez et al. [24] with author approval.

331 D.2 Uncertainty quantification via conditional quantile regression

332 The goal of these experiments benchmark is to learn the family of conditional distributions $\mathbb{P}(\mathbf{y} | \mathbf{x} =$
 333 $\cdot)$ for a bivariate random variable $\mathbf{y} = [y_0, y_1] \in \mathbb{R}^2$ given a scalar covariate $\mathbf{x} \in \mathbb{R}$. Once $\mathbb{P}(\mathbf{y} | \mathbf{x})$
 334 is recovered, the practitioner can estimate *conditional* $(1 - \alpha)$ -confidence regions by regressing
 335 the lower and upper conditional quantiles $q_{\alpha/2}(\mathbf{x})$, $q_{1-\alpha/2}(\mathbf{x})$ for any desired miscoverage level
 336 $\alpha \in (0, 1)$. In particular, a 95% confidence region corresponds to $\alpha = 0.05$, so the two quantiles of
 337 interest are $q_{0.025}(\mathbf{x})$ and $q_{0.975}(\mathbf{x})$. See Fig. 6 for a visual representation of the problem.




Figure 6: Synthetic experiment in uncertainty quantification, originally proposed by Feldman et al. [7]. The task is to predict the 95% confidence intervals (black bounding boxes) of a random variable $\mathbf{y} \in \mathbb{R}^2$ conditioned on a scalar random variable $\mathbf{x} \in \mathbb{R}$. **Left:** The marginal distribution $\mathbb{P}(\mathbf{y})$. **Middle:** The marginal distribution $\mathbb{P}(\mathbf{x})$. **Right:** Example conditional distributions $\mathbb{P}(\mathbf{y} | \mathbf{x} = \cdot)$ for different conditioning values.

338 **Conditional quantile regression models** We compare the NCP and proposed eNCP models to a
 339 standard baseline for parametric NN conditional quantile regression, namely CQR [7], which uses

two separate NNs to predict the lower and upper quantiles of the conditional distribution, trained with a pinball loss function (see [7] for details). Both models use MLP backbones with similar parameter counts, ensuring that improvements are solely due to the loss functions.

Furthermore, CQR can only be trained for specific confidence intervals, requiring retraining for different quantiles. In contrast, the NCP and eNCP models regress the CCDF of each dimension of $\mathbf{y}$ given $\mathbf{x}$. Thus, they can estimate conditional quantiles for any confidence interval via the quantile estimation algorithm from the CCDF described in Kostic et al. [15] without retraining. See details in Fig. 7.

Evaluation metrics: coverage and set size Let $\mathbb{C}_{1-\alpha}(\mathbf{x}) \subseteq \mathbb{R}^d$ denote a *prediction set* of nominal level $(1 - \alpha)$ produced by a conditional quantile regression model for the response $\mathbf{y} \in \mathbb{R}^d$ given the covariate $\mathbf{x} \in \mathbb{R}^p$. In all experiments we assess two complementary metrics.

- **Coverage.** The conditional *coverage* of $\mathbb{C}_{1-\alpha}$ is the probability that the true response is captured by the predicted region,

$$c_{1-\alpha}(\mathbf{x}) := \mathbb{P}(\mathbf{y} \in \mathbb{C}_{1-\alpha}(\mathbf{x}) \mid \mathbf{x}), \quad \text{with the target } c_{1-\alpha}(\mathbf{x}) \approx 1 - \alpha \quad \forall \mathbf{x}. \quad (12)$$

In practice we report the *marginal coverage* $\hat{\mathbb{E}}_{\mathbf{x}}[c_{1-\alpha}(\mathbf{x})]$, estimated on a large held-out sample; values above (resp. below) $1 - \alpha$ indicate over- (resp. under-) coverage.

- **Relaxed Coverage (r-Coverage).** The conditional *relaxed coverage* of $\mathbb{C}_{1-\alpha}$ is defined as the probability that each scalar component of the response lies within its corresponding predicted confidence interval. Formally, if $\mathbf{y} = [y_1, \dots, y_d]$ and $\mathbb{C}_{1-\alpha}(\mathbf{x})$ has corresponding marginal intervals $\mathbb{C}_{1-\alpha}^{(i)}(\mathbf{x})$ for $i \in \{1, \dots, d\}$, then

$$rc_{1-\alpha}(\mathbf{x}) := \prod_{i=1}^d \mathbb{P}(y_i \in \mathbb{C}_{1-\alpha}^{(i)}(\mathbf{x}) \mid \mathbf{x}), \quad (13)$$

with the target $rc_{1-\alpha}(\mathbf{x}) \approx 1 - \alpha$ for all $\mathbf{x}$. As with coverage, we report the *marginal relaxed coverage* $\hat{\mathbb{E}}_{\mathbf{x}}[rc_{1-\alpha}(\mathbf{x})]$.

- **Set size.** To quantify how informative the region is, we measure its *size* (volume) under the Lebesgue measure λ^d :

$$\text{Size}_{1-\alpha}(\mathbf{x}) := \text{vol}(\mathbb{C}_{1-\alpha}(\mathbf{x})). \quad (14)$$

Smaller sets correspond to sharper uncertainty estimates, provided the required coverage is met. For multidimensional responses the volume is expressed in the natural units of $\mathbb{R}^d$; for $d = 1$ it reduces to the interval length. As with coverage, we report the marginal expectation $\hat{\mathbb{E}}_{\mathbf{x}}[\text{Size}_{1-\alpha}(\mathbf{x})]$ so that models can be compared fairly across the entire input distribution.




Figure 7: Prediction of the 80% and 95% confidence intervals for the random variable $\mathbf{y}$ in experiment App. D.2 using the proposed eNCP model. The model estimates the CCDF by discretizing each dimension of $\mathbf{y} = [y_1, y_2]$ into 100 bins and computing the conditional probabilities $\mathbb{P}(y_i \in \mathbb{A}_n \mid \mathbf{x} = \cdot) := [\mathbb{E}_{\mathbf{y} \mid \mathbf{x}} \mathbb{1}_{\mathbb{A}_n}](\cdot)$ for all $n \in [100]$ based on the learned conditional expectation operator $\kappa_{\theta}(\mathbf{x}, \mathbf{y})$ (see Tab. 1). Here, $\mathbb{A}_n$ comprises the bins from the 0-th to the n -th. This yields the estimated CCDF for y_1 (center) and y_2 (right) at $\mathbf{x} = 2.7$. The CCDFs can then be used to estimate upper and lower quantiles for any confidence interval [15]. In practice, the eNCP model regresses 2×100 variables in a single forward pass. Thus, the final layer of the conditional quantile regression model is a linear layer of size $r \times (2 \times 100)$, where r is the number of features in the $\mathbf{y}$ representation.

367 D.2.1 Synthetic benchmark

368 The goal of these experiments is to learn the conditional distributions $\mathbb{P}(\mathbf{y} \mid \mathbf{x} = \cdot)$ for a bivariate
 369 random variable $\mathbf{y} = [y_0, y_1] \in \mathbb{R}^2$ given a scalar covariate $\mathbf{x} \in \mathbb{R}$. Following Feldman et al. [7], the
 370 covariate is sampled uniformly: $\mathbf{x} \sim \text{Unif}(0.8, 3.2)$, and the response variable $\mathbf{y}$ is produced by a
 371 non-linear transformation of auxiliary latent variables (see Fig. 6):

$$\begin{aligned} y_0 &= \frac{z}{\beta \mathbf{x}} + r \cos \phi, & z &\sim \text{Unif}(-\pi, \pi), \\ y_1 &= \frac{1}{2}(-\cos z + 1) + r \sin \phi + \sin \mathbf{x}, & \phi &\sim \text{Unif}(0, 2\pi), \\ & & r &\sim \text{Unif}(-0.1, 0.1). \end{aligned}$$

372 Here, $\beta > 0$ is a scaling constant.

373 The additive perturbation $r(\cos \phi, \sin \phi)$ yields heteroskedastic, anisotropic noise, whereas the
 374 $\frac{1}{2}(-\cos z + 1)$ and $\sin \mathbf{x}$ terms introduce strong non-monotonicity and interaction effects between $\mathbf{x}$
 375 and $\mathbf{y}$. As a result, the conditional quantile functions $\mathbf{x} \mapsto q_\tau(\mathbf{x})$ are highly non-linear, making this
 376 dataset an ideal low-dimensional experiment for conditional quantile regression methods.

377 **Results** The experiment results are depicted in Fig. 9. Where the **NCP** and **eNCP** models outperform
 378 the baseline **CQR** model in terms of both coverage and set size. Furthermore, Fig. 8 illustrates the
 379 basis functions learned by the **NCP** and **eNCP** models for the random variable $\mathbf{y} = [y_0, y_1]$. In
 380 contrast to the standard **NCP** model, the **eNCP** model incorporates symmetry priors, enabling a
 381 clean separation of its latent representation into two orthogonal subspaces: one corresponding to
 382 $\mathbb{C}_2$ -invariant functions and the other to functions that change sign under reflection.




Figure 8: **Left:** Learned basis functions from the **NCP** model for $\mathbf{y} = [y_0, y_1]$. **Right:** Learned basis functions from the **eNCP** model for $\mathbf{y}$. The marginal distribution of $\mathbf{y}$ exhibits reflection symmetry $g_r \triangleright \mathbf{y} = [-y_0, y_1]$ under $\mathbb{G} = \mathbb{C}_2$. Incorporating this prior, the **eNCP** model decomposes its latent space as $\mathcal{F}_{\mathbf{y}} = \mathcal{F}_{\mathbf{y}}^{\text{inv}} \oplus \mathcal{F}_{\mathbf{y}}^{(2)}$, with the first subspace capturing $\mathbb{C}_2$ -invariant functions and the second capturing those that change sign under reflection. The orthogonality of these subspaces allows independent optimization of the basis functions.




Figure 9: Results of a synthetic experiment in uncertainty quantification comparing **CQR**, **NCP**, and **eNCP** models. The task, originally proposed by Feldman et al. [7], is to predict the 95% confidence intervals of a random variable $\mathbf{y} \in \mathbb{R}^2$ conditioned on a scalar random variable $\mathbf{x} \in \mathbb{R}$. The conditional distributions $\mathbb{P}(\mathbf{y}|\mathbf{x} = \cdot)$ are shown in the left and fourth columns for different conditioning values, while the second-third and fifth-sixth columns display the **CCDF** predicted by the **eNCP** and **NCP** models, respectively. The **CQR** model directly regresses the upper and lower quantiles for each dimension of $\mathbf{y}$ and must be retrained if the confidence interval probability changes. In contrast, since the **NCP** and **eNCP** models estimate the **CCDF** for each dimension, these predictions can be easily adapted to any confidence interval probability by simply changing the threshold value.

	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 3: Validation and test set metrics for the prediction of 95% confidence intervals on observables of a quadruped robot traversing rough terrains (see App. D.2.2). Model performance is evaluated using three metrics: (i) relaxed coverage (r-Coverage) (Eq. (13)), (ii) coverage (Eq. (12)), and (iii) set size (Eq. (14)). The best results are highlighted in blue. Note that although the confidence interval volumes (set size) of the eCQR and CQR models are significantly smaller than those of the NCP and eNCP models, the former fail to achieve the expected 95% coverage on both the validation and test sets. In contrast, the eNCP model attains the best overall coverage, proving its effectiveness for uncertainty quantification. Importantly, the eNCP and NCP models can be adjusted, **without retraining**, to provide confidence intervals for any desired coverage level, whereas the CQR and eCQR models must be retrained for each new level.

D.2.2 Uncertainty quantification in quadruped legged locomotion

We test how well conditional-quantile models can recover the conditional 95% confidence regions of three physically meaningful observables produced by a simulated AlienGo quadruped walking over rough terrain (see Fig. 2) under varying friction coefficients. The dataset was collected using the Quadruped-PyMPC simulation framework and model predictive controller from [33].

The observables for which state-dependent uncertainty estimates are desired are $\mathbf{y}_t = [U_t, T_t, \boldsymbol{\tau}_t^{\text{grf}}]^\top$, with each component defined as follows:

- **G-invariant Kinetic Energy.** $T(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2} \dot{\mathbf{q}}^\top M(\mathbf{q}) \dot{\mathbf{q}} \in \mathbb{R}$, where $M(\mathbf{q})$ is the configuration-dependent inertia matrix. Noise is introduced through sensor measurement errors on the robot’s degree of freedom (DoF) position $\mathbf{q} \in \mathbb{R}^{12}$ and velocity $\dot{\mathbf{q}} \in \mathbb{R}^{12}$.
- **G-invariant Instantaneous Mechanical Work.** $U(\mathbf{q}, \dot{\mathbf{q}}, \boldsymbol{\tau}) \in \mathbb{R}$, representing the instantaneous mechanical work exerted or absorbed by the robot. This quantity depends on the actuator torques (typically measured with noisy, biased sensors) as well as the external forces (e.g. gravity, contact forces) that are not reliably measurable due to unobserved terrain parameters.
- **G-equivariant Ground-Reaction Forces** $\boldsymbol{\tau}_{\text{grf}} \in \mathbb{R}^{12}$, a fundamental quantity in quadruped control, whose reliable estimation and uncertainty quantification are critical for downstream tasks in robotics [17, 22].

The observables of interest are predicted using a suit of onboard proprioceptive sensory signals available at time t :

$$\mathbf{x}_t = [\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{a}_t, \mathbf{v}_t, \mathbf{v}_{t,\text{err}}, \boldsymbol{\omega}_t, \boldsymbol{\omega}_{t,\text{err}}, \mathbf{g}_t, \dot{\mathbf{p}}_{t,\text{feet}}, \boldsymbol{\tau}_t^{\text{cmd}}]^\top.$$

Specifically, $\mathbf{q}_t \in \mathbb{R}^{n_q}$ 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 $\mathbf{v}_{t,\text{err}} \in \mathbb{R}^3$ the command error base linear velocity; $\boldsymbol{\omega}_t \in \mathbb{R}^3$ 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; $\dot{\mathbf{p}}_{t,\text{feet}} \in \mathbb{R}^{12}$ 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.

Hence we design the experiments to compare models of similar footprint in number of parameters, while the loss used for training differs between the NCP and eNCP models w.r.t to the CQR and eCQR models.

NN architectures We configure all models considered eNCP, NCP, eCQR, and CQR to have an inference-time NN architecture of the similar footprint. The backbone of all architectures is a standard multilayer perceptron consisting of three hidden layers, each with 512 units, followed by

a final hidden layer containing 128 units. This final layer serves to encode the feature vector r for the **NCP** and **eNCP** models. Crucially, since $\mathbb{G}$ -equivariance enforces weight sharing in the **NN** architecture the encoder **NN** for **eNCP**, **eCQR** have $\times 2$ less parameters than their symmetry-agnostic counterparts.

Results. Given sensory input $\mathbf{x}$, the model predicts a set $\mathbb{C}_{0.95}(\mathbf{x}) \subseteq \mathbb{R}^{14}$ satisfying $\mathbb{P}(\mathbf{y} \in \mathbb{C}_{0.95}(\mathbf{x}) \mid \mathbf{x}) \approx 0.95$, while minimizing its volume $\hat{\mathbb{E}}_{\mathbf{x}}[\text{vol}(\mathbb{C}_{0.95}(\mathbf{x}))]$. Empirically high coverage implies that the true $\mathbb{G}$ -invariant kinetic energy, instantaneous mechanical work, and the $\mathbb{G}$ -equivariant 12-dimensional ground-reaction forces lie within the predicted confidence set. In contrast, relaxed coverage (r-Coverage) quantifies the reliability of the estimates on a per-dimension basis. **Tab. 3** summarizes the validation and test results for the **eNCP**, **NCP**, **CQR**, and **eCQR** models, and **Fig. 10** illustrates a trajectory of **GRF** and their respective 90% confidence intervals for each model. Both **CQR** and **eCQR** tend to produce confidence intervals of smaller volume but fail to achieve the desired coverage on the testing set, implying that the models' confidence intervals are not reliable and require further calibration through retraining or conformal calibration [7]. In contrast, the **eNCP** model achieves the desired coverage on the test set while producing confidence intervals of larger volume, hence yielding reliable confidence intervals.




Figure 10: Prediction of 90% confidence intervals (CI) for the ground-reaction forces $\tau_{\text{grf}} \in \mathbb{R}^{12}$ of a quadruped robot on rough terrain with varying friction. We compare the **eNCP**, **NCP**, **eCQR**, and **CQR** models based on relaxed coverage and set size (see **Tab. 3**). CIs are computed for each leg—front-right (FR), front-left (FL), hind-right (HR), and hind-left—along the x , y , and z axes. Forces outside the CI are highlighted in red, while those within appear in blue. Terrain variations cause significant variability in the x and y components due to differences in surface orientation and friction, whereas the z component is mainly influenced by local height changes that alter contact timing and produce short-duration high-impact forces.

E Background on group and representation theory

Group actions and representations This section provides a concise overview of the fundamental concepts in group and representation theory, which are used to define the symmetries of the random variables we consider in this work. For a comprehensive background on these topics in finite-dimensional vector spaces, see Weiler et al. [36]; for the infinite-dimensional case, consult Knapp [14]. These concepts will be referenced as needed in the main text. To begin, we define a group as an abstract mathematical object.

Definition E.1 (Group). A group is a set $\mathbb{G}$, endowed with a binary composition operator defined as:

$$\begin{aligned} (\circ) : \quad \mathbb{G} \times \mathbb{G} &\longrightarrow \mathbb{G} \\ (g_1, g_2) &\longrightarrow g_1 \circ g_2, \end{aligned} \quad (15a)$$

such that the following axioms hold:

$$\text{Associativity: } (g_1 \circ g_2) \circ g_3 = g_1 \circ (g_2 \circ g_3), \quad \forall g_1, g_2, g_3 \in \mathbb{G}, \quad (15b)$$

$$\text{Identity: } \exists e \in \mathbb{G} \text{ such that } e \circ g = g = g \circ e, \quad \forall g \in \mathbb{G}, \quad (15c)$$

$$\text{Inverses: } \forall g \in \mathbb{G}, \exists g^{-1} \in \mathbb{G} \text{ such that } g \circ g^{-1} = e = g^{-1} \circ g. \quad (15d)$$

We are primarily interested in symmetry groups, i.e., groups of transformations acting on a set $\mathcal{X}$. Each transformation is a bijection that leaves a fundamental property invariant. For example, if $\mathcal{X}$ represents states of a dynamical system, the invariant property is the state energy (see Fig. 5); if $\mathcal{X}$ is a data space, the preserved quantity is typically the probability density/distribution (see Fig. 3).

Definition E.2 (Group action on a set [36]). 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:

$$\begin{aligned} (\triangleright) : \quad \mathbb{G} \times \mathcal{X} &\longrightarrow \mathcal{X} \\ (g, \mathbf{x}) &\longrightarrow g \triangleright \mathbf{x} \end{aligned} \quad (16a)$$

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

$$\text{Identity: } e \triangleright \mathbf{x} = \mathbf{x}, \quad \forall \mathbf{x} \in \mathcal{X} \quad (16b)$$

$$\text{Associativity: } (g_1 \circ g_2) \triangleright \mathbf{x} = g_1 \triangleright (g_2 \triangleright \mathbf{x}), \quad \forall g_1, g_2 \in \mathbb{G}, \forall \mathbf{x} \in \mathcal{X}. \quad (16c)$$

We are primarily interested in studying symmetry transformations on sets with a vector space structure. In most practical cases, the group action on a vector space is linear, allowing symmetry transformations to be represented as linear invertible maps. These maps can be expressed in matrix form once a basis for the space is chosen.

Definition E.3 (Linear group representation). Let $\mathcal{X}$ 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 $\mathbb{GL}(\mathcal{X})$):

$$\begin{aligned} \rho_{\mathcal{X}} : \quad \mathbb{G} &\longrightarrow \mathbb{GL}(\mathcal{X}) \\ g &\longrightarrow \rho_{\mathcal{X}}(g), \end{aligned} \quad (17a)$$

such that the following properties hold:

$$\text{composition : } \rho_{\mathcal{X}}(g_1 \circ g_2) = \rho_{\mathcal{X}}(g_1) \rho_{\mathcal{X}}(g_2), \quad \forall g_1, g_2 \in \mathbb{G}, \quad (17b)$$

$$\text{inversion : } \rho_{\mathcal{X}}(g^{-1}) = \rho_{\mathcal{X}}(g)^{-1}, \quad \forall g \in \mathbb{G}. \quad (17c)$$

$$\text{identity : } \rho_{\mathcal{X}}(g \circ g^{-1}) = \rho_{\mathcal{X}}(e) = \mathbf{I}, \quad (17d)$$

Whenever the vector space is of finite dimension $n < \infty$, linear maps admit a matrix form $\rho_{\mathcal{X}}(g) \in \mathbb{R}^{n \times n}$, once a basis set $\mathbb{I}_{\mathcal{X}}$ for the vector space $\mathcal{X}$ is chosen. In this case, Eqs. (17b) to (17d) show how the composition and inversion of symmetry transformations translate to matrix multiplication and inversion, respectively. Moreover, $\rho_{\mathcal{X}}$ allows to express a (linear) group action (Def. E.2) as a matrix-vector multiplication:

$$\begin{aligned} (\triangleright) : \quad \mathbb{G} \times \mathcal{X} &\longrightarrow \mathcal{X} \\ (g, \mathbf{x}) &\longrightarrow g \triangleright \mathbf{x} := \rho_{\mathcal{X}}(g) \mathbf{x}. \end{aligned} \quad (17e)$$

Since the matrix form of linear maps depends on the choice of basis, we can relate different matrix representations of the same linear map through changes of basis. This leads us to the concept of equivalent group representations.

Definition E.4 (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 a change of basis $Q : \mathcal{X} \rightarrow \mathcal{X}$ such that*

$$\rho'_\mathcal{X}(g) = Q \rho_\mathcal{X}(g) Q^{-1}, \quad \forall g \in \mathbb{G}. \quad (18)$$

Equivalent representations arise when the same group action $(\triangleright) : \mathbb{G} \times \mathcal{X} \rightarrow \mathcal{X}$ 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 $Q_\mathbb{A}^\mathbb{B} : \mathcal{X} \rightarrow \mathcal{X}$ denote the change of basis from $\mathbb{A}_\mathcal{X}$ to $\mathbb{B}_\mathcal{X}$, so that $\mathbf{x}^\mathbb{B} = Q_\mathbb{A}^\mathbb{B} \mathbf{x}^\mathbb{A}$ 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

$$\begin{aligned} g \triangleright \mathbf{x}^\mathbb{B} &:= Q_\mathbb{A}^\mathbb{B}(g \triangleright \mathbf{x}^\mathbb{A}), \quad \forall g \in \mathbb{G}, \\ \rho_\mathcal{X}^\mathbb{B}(g) \mathbf{x}^\mathbb{B} &= Q_\mathbb{A}^\mathbb{B}(\rho_\mathcal{X}^\mathbb{A}(g) \mathbf{x}^\mathbb{A}) = (Q_\mathbb{A}^\mathbb{B} \rho_\mathcal{X}^\mathbb{A}(g) Q_\mathbb{A}^{\mathbb{B}-1}) \mathbf{x}^\mathbb{B}, \\ \rho_\mathcal{X}^\mathbb{B}(g) &= Q_\mathbb{A}^\mathbb{B} \rho_\mathcal{X}^\mathbb{A}(g) Q_\mathbb{A}^{\mathbb{B}-1}. \end{aligned} \quad (19)$$

To reveal the modular structure of symmetric vector spaces, we often change bases to decompose them into subspaces stable under the action of the group $\mathbb{G}$, termed $\mathbb{G}$ -stable subspaces. This decomposition mirrors how a symmetry group breaks down into subgroups and is essential for analyzing and simplifying group representations. We introduce the following definition.

Definition E.5 ($\mathbb{G}$ -stable and irreducible subspaces). *Let $\mathcal{X}$ be a vector space endowed with a group action $(\triangleright)$ 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,*

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

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.

Decomposing symmetric vector spaces into $\mathbb{G}$ -stable subspaces corresponds to decomposing the group representation associated with $\triangleright$ into smaller representations acting on these $\mathbb{G}$ -stable subspaces:

Definition E.6 (Decomposable representation). *Let $\mathcal{X}$ be a vector space with a group action $(\triangleright)$ 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 $Q_\mathbb{A}^\mathbb{B} : \mathcal{X} \rightarrow \mathcal{X}$ such that*

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

This shows that the decomposition $\rho_\mathcal{X} \sim \rho_{\mathcal{X}_1} \oplus \rho_{\mathcal{X}_2}$ corresponds to splitting the vector space into $\mathbb{G}$ -stable subspaces, $\mathcal{X} = \mathcal{X}_1 \oplus \mathcal{X}_2$. Moreover, if the representation is block-diagonal in some basis set $\mathbb{B}_\mathcal{X}$, then $\mathbb{B}_\mathcal{X}$ is the union of disjoint basis sets $\mathbb{B}_{\mathcal{X}_1}$ and $\mathbb{B}_{\mathcal{X}_2}$ for $\mathcal{X}_1$ and $\mathcal{X}_2$, respectively.

Definition E.7 (Irreducible representation). *Let $\mathcal{X}$ be a vector space endowed with a group action $(\triangleright)$ 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 (Def. E.5). That is, the only $\mathbb{G}$ -stable subspaces $\mathcal{X}' \subseteq \mathcal{X}$ are $\mathcal{X}' = \{\mathbf{0}\}$ and $\mathcal{X}' = \mathcal{X}$ itself.*

We have now equipped all the necessary tools to decompose symmetric vector spaces into their smallest building blocks: irreducible $\mathbb{G}$ -stable subspaces.

Irreducible representations are the fundamental building blocks for all representations of the group $\mathbb{G}$. Any unitary representation can be decomposed into a direct sum of irreducible representations, analogous to the prime factorization of integers. In terms of the vector spaces on which the group acts, this decomposition of the representation corresponds to decomposing the space into $\mathbb{G}$ -irreducible subspaces (Def. E.5):

Theorem E.8 (Isotypic decomposition of symmetric Hilbert spaces [14]). *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 $\bar{\rho}_k : \mathbb{G} \rightarrow \mathbb{U}(\mathcal{H}_k)$ that allow us to decompose $\mathcal{H}$ into a sum of orthogonal subspaces, denoted isotypic subspaces: $\mathcal{H} = \bigoplus_{1 \leq k \leq n_{\text{iso}}}^\perp \mathcal{H}_k$ where each $\mathcal{H}_k = \bigoplus_{j=1}^{m_k} \mathcal{H}_{k,j}$ is the sum of at most $m_k \leq \infty$ countably many subspaces isometrically isomorphic to $\mathcal{H}_k$.*

504 **Isotypic decomposition and disentangled representations** Whenever the symmetric vector space
 505 of interest defines a vector valued representation of some data, the isotypic decomposition of the
 506 representation space is intricately linked with the concept of *disentangled representations*

507 **Definition E.9** (Disentangled representation (Higgins et al. [11])). *A vector representation is called*
 508 *a disentangled representation with respect to a particular decomposition of a symmetry group into*
 509 *subgroups, if it decomposes into independent subspaces, where each subspace is affected by the*
 510 *action of a single subgroup, and the actions of all other subgroups leave the subspace unaffected.*

511 The *subspaces* of Def. E.9 refer to each of the isotypic subspaces $\mathcal{H}_i$, and the symmetry subgroups
 512 refer to the effective (matrix) group encoded by each irreducible representation $\bar{\rho}_k : \mathbb{G} \mapsto \mathbb{U}(\mathcal{H}_k)$.
 513 Which we denote in the main body as $\mathbb{G}^{(k)}$.

514 E.1 Maps between symmetric vector spaces

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

518 **Definition E.10** ($\mathbb{G}$ -equivariant and $\mathbb{G}$ -invariant maps). *Let $\mathcal{X}$ and $\mathcal{Y}$ be two vector spaces endowed*
 519 *with the same symmetry group $\mathbb{G}$, with the respective group actions $\triangleright_{\mathcal{X}}$ and $\triangleright_{\mathcal{Y}}$. A map $f : \mathcal{X} \mapsto \mathcal{Y}$ is*
 520 *said to be $\mathbb{G}$ -equivariant if it commutes with the group action, such that:*

$$\begin{aligned} g \triangleright_{\mathcal{Y}} \mathbf{y} = g \triangleright_{\mathcal{Y}} f(\mathbf{x}) = f(g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, g \in \mathbb{G}. \\ \rho_{\mathcal{Y}}(g)f(\mathbf{x}) = f(\rho_{\mathcal{X}}(g)\mathbf{x}) \end{aligned} \iff \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \downarrow f & & \downarrow f \\ \mathcal{Y} & \xrightarrow{\triangleright_{\mathcal{Y}}} & \mathcal{Y} \end{array} \quad (20a)$$

521 A specific case of $\mathbb{G}$ -equivariant maps are the $\mathbb{G}$ -invariant ones, which are maps that commute with
 522 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:

$$\begin{aligned} \mathbf{y} = g \triangleright_{\mathcal{Y}} f(\mathbf{x}) = f(g \triangleright_{\mathcal{X}} \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}, g \in \mathbb{G}. \\ \mathbf{y} = \rho_{\mathcal{Y}}(g)f(\mathbf{x}) = f(\rho_{\mathcal{X}}(g)\mathbf{x}) \end{aligned} \iff \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ & \searrow f & \downarrow f \\ & & \mathcal{Y} \end{array} \quad (20b)$$

$\curvearrowright \triangleright_{\mathcal{Y}}$

523 E.2 Structure of $\mathbb{G}$ -equivariant linear maps

524 **Definition E.11** (Homomorphism, Isomorphism, and $\mathbb{G}$ -equivariant linear maps). *Let $\mathcal{X}$ and $\mathcal{Y}$ be*
 525 *two vector spaces endowed with the same symmetry group $\mathbb{G}$, with the respective group actions*
 526 $\triangleright_{\mathcal{X}} : \mathbb{G} \times \mathcal{X} \mapsto \mathcal{X}$ 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*
 527 *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*
 528 *all $\mathbf{x} \in \mathcal{X}$. They are said to be $\mathbb{G}$ -isomorphic if the linear map is invertible. Graphically, $\mathcal{X}$ and $\mathcal{Y}$*
 529 *are $\mathbb{G}$ -homomorphic or $\mathbb{G}$ -isomorphic if the following diagrams commute:*

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \downarrow \mathbf{A} & & \downarrow \mathbf{A} \\ \mathcal{Y} & \xrightarrow{\triangleright_{\mathcal{Y}}} & \mathcal{Y} \end{array} \quad \mathbf{A} \in \text{Homo}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y}) \quad \text{or} \quad \begin{array}{ccc} \mathcal{X} & \xrightarrow{\triangleright_{\mathcal{X}}} & \mathcal{X} \\ \uparrow \mathbf{A}^{-1} & & \uparrow \mathbf{A}^{-1} \\ \mathcal{Y} & \xrightarrow{\triangleright_{\mathcal{Y}}} & \mathcal{Y} \end{array} \quad \mathbf{A} \in \text{Iso}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y}). \quad (21)$$

$\underbrace{\hspace{10em}}_{\text{Homomorphism}} \qquad \underbrace{\hspace{10em}}_{\text{Isomorphism}}$

530 Here, $\text{Homo}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y})$ denotes the space of $\mathbb{G}$ -equivariant linear maps between $\mathcal{X}$ and $\mathcal{Y}$, and
 531 $\text{Iso}_{\mathbb{G}}(\mathcal{X}, \mathcal{Y})$ denotes the space of $\mathbb{G}$ -equivariant invertible linear maps between $\mathcal{X}$ and $\mathcal{Y}$.

532 **Lemma E.12** (Schur's Lemma for unitary representations [14, Prop 1.5]). *Consider two Hilbert*
 533 *spaces, $\mathcal{H}$ and $\mathcal{H}'$, endowed with the irreducible unitary representations $\bar{\rho}_{\mathcal{H}} : \mathbb{G} \mapsto \mathbb{U}(\mathcal{H})$ and*
 534 $\bar{\rho}_{\mathcal{H}'} : \mathbb{G} \mapsto \mathbb{U}(\mathcal{H}')$, respectively. *Let $\mathbf{T} : \mathcal{H} \mapsto \mathcal{H}'$ be a linear $\mathbb{G}$ -equivariant operator such that*
 535 $\bar{\rho}_{\mathcal{H}'}\mathbf{T} = \mathbf{T}\bar{\rho}_{\mathcal{H}}$. *If the irreducible representations are not equivalent, i.e., $\bar{\rho}_{\mathcal{H}} \not\sim \bar{\rho}_{\mathcal{H}'}$, then $\mathbf{T}$ is the*

536 trivial (or zero) map. Conversely, if $\bar{\rho}_{\mathcal{H}} \sim \bar{\rho}_{\mathcal{H}'}$, then $\mathbb{T}$ is a constant multiple of an isomorphism
 537 (Def. E.11). Denoting $\mathbb{I}$ as the identity operator, this can be expressed as:

$$\bar{\rho}_{\mathcal{H}} \approx \bar{\rho}_{\mathcal{H}'} \iff \mathbf{0}_{\mathcal{H}'} = \mathbb{T} \mathbf{h} \mid \forall \mathbf{h} \in \mathcal{H} \quad (22a)$$

$$\bar{\rho}_{\mathcal{H}} \sim \bar{\rho}_{\mathcal{H}'} \iff \mathbb{T} = \alpha \mathbb{U}, \alpha \in \mathbb{C}, \mathbb{U} \cdot \mathbb{U}^H = \mathbb{I} \quad (22b)$$

$$\bar{\rho}_{\mathcal{H}} = \bar{\rho}_{\mathcal{H}'} \iff \mathbb{T} = \alpha \mathbb{I} \quad (22c)$$

538 For intuiion referer to the following blog post

539 F Representation theory of symmetric function spaces

540 In this section, we study symmetry group actions on infinite-dimensional function spaces and specify
 541 the conditions needed to approximate these spaces in finite dimensions. Specifically, given a set $\mathcal{X}$
 542 with a compact symmetry group $\mathbb{G}$ acting via $(\triangleright)$ (Def. E.2), the space of scalar-valued functions
 543 on $\mathcal{X}$, $\mathcal{F} = \{f \mid f : \mathcal{X} \mapsto \mathbb{R}\}$, becomes a symmetric function space. The action of a symmetry
 544 transformation on a function is defined as:

545 **Definition F.1** (Group action on a function space). *Let $\mathcal{X}$ be a set endowed with the symmetry group*
 546 *$\mathbb{G}$, and let $\mathcal{F}$ be the space of scalar-valued functions on $\mathcal{X}$. The (left) action of $\mathbb{G}$ on a function*
 547 *$f \in \mathcal{F}$ is defined as the composition of f with the inverse of the group element g^{-1} :*

$$\begin{aligned} (\triangleright_{\mathcal{F}}) : \mathbb{G} \times \mathcal{F} &\longrightarrow \mathcal{F} \\ (g, f) &\longrightarrow [g \triangleright_{\mathcal{F}} f](\mathbf{x}) := [f \circ g^{-1}](\mathbf{x}) = f(g^{-1} \triangleright \mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X}. \end{aligned} \quad (23a)$$

548 In other words, the point-wise evaluation of f on a g^{-1} -transformed set $\mathcal{X}$ is equivalent to the
 549 evaluation of the transformed function $g \triangleright_{\mathcal{F}} f \in \mathcal{F}$ on the original set $\mathcal{X}$ (see simple examples
 550 in Fig. 11). Any function space that is stable under the group action Eq. (23a) is referred to as a
 551 symmetric function space. Note that this action is compatible with the group composition and identity
 552 element $e \in \mathbb{G}$, such that the following properties hold:

$$\text{Identity: } e \triangleright_{\mathcal{F}} f(\cdot) = f(\cdot), \quad (23b)$$

$$\text{Associativity: } [(g_2 \circ g_1) \triangleright_{\mathcal{F}} f](\cdot) = [g_2 \triangleright_{\mathcal{F}} [g_1 \triangleright_{\mathcal{F}} f]](\cdot), \quad \forall g_1, g_2 \in \mathbb{G}. \quad (23c)$$




Figure 11: **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 $g_s \triangleright x = -x$ and on the function space $\mathcal{F} := \{f \mid f : \mathcal{X} \mapsto \mathbb{R}\}$ by $[g \triangleright_{\mathcal{F}} f](\mathbf{x}) = f(g \triangleright_{\mathcal{X}} \mathbf{x}) = f(-\mathbf{x})$. Hence we have that f_1 is a $\mathbb{G}$ -invariant function, $g_s \triangleright_{\mathcal{F}} f_1(x) = f_1(x)$ and f_2 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 $f(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{c}_1, 2) + \mathcal{N}(\mathbf{x}; \mathbf{c}_2, 1)$ 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 $g \triangleright \mathbf{x} = \rho_{\mathcal{X}}(g)\mathbf{x} = \mathbf{R}_g \mathbf{x}$, where $\mathbf{R}_g$ 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 $z \in \hat{\mathcal{F}}$, defined to be a member of the finite-dimensional symmetric function space $\hat{\mathcal{F}} := \text{span}(\mathbb{I}_{\hat{\mathcal{F}}})$, constructed from a basis set composed of the group orbit of the (arbitrarily chosen) function $f \in \mathcal{F}$, that is $\mathbb{I}_{\hat{\mathcal{F}}} := \mathbb{G} f = \{f, g_{120} \triangleright_{\mathcal{F}} f, g_{240} \triangleright_{\mathcal{F}} f\}$. This function space is $\mathbb{G}$ -stable by construction, since $\mathbb{G} \mathbb{I}_{\hat{\mathcal{F}}} = \mathbb{I}_{\hat{\mathcal{F}}}$. Note that the z -offsets are added for visualization purposes.

553 *Remark F.2.* From an algebraic perspective, the inversion g^{-1} (*contragredient representation*)
 554 emerges to ensure that the associativity property of the group action (Eq. (23c)) holds:

$$\begin{aligned} [(g_2 \circ g_1) \triangleright_{\mathcal{F}} f](\mathbf{x}) &= [g_2 \triangleright_{\mathcal{F}} [g_1 \triangleright_{\mathcal{F}} f]](\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{X} \\ f((g_2 \circ g_1)^{-1} \triangleright \mathbf{x}) &= [g_1 \triangleright_{\mathcal{F}} f](g_2^{-1} \triangleright \mathbf{x}) = f(g_1^{-1} \triangleright (g_2^{-1} \triangleright \mathbf{x})) \\ f((g_2 \circ g_1)^{-1} \triangleright \mathbf{x}) &= f((g_1 \circ g_2)^{-1} \triangleright \mathbf{x}). \end{aligned}$$

555 In the context of this work, we will study the scenario where the function space $\mathcal{F}$ is a separable
 556 Hilbert space and the group action of $\mathbb{G}$ on $\mathcal{F}$ is unitary, i.e., it preserves the inner product of the
 557 function space. This setup is crucial to enable us to approximate $\mathcal{F}$ and the group action on $\mathcal{F}$ in
 558 finite dimensions.

559 F.1 Unitary group representation on function spaces

560 Assume our symmetric set $\mathcal{X}$ is endowed with a measure space structure $(\mathcal{X}, \Sigma_{\mathcal{X}}, P_{\mathbf{x}})$, where
 561 $P_{\mathbf{x}} : \Sigma_{\mathcal{X}} \mapsto \mathbb{R}$ is the space measure. Then, consider a function space with a separable Hilbert space
 562 structure $\mathcal{F} := \mathcal{L}_{P_{\mathbf{x}}}^2 \mathcal{X}, \mathbb{R}$, and inner product $\langle f_1, f_2 \rangle_{P_{\mathbf{x}}} = \int_{\mathcal{X}} f_1(\mathbf{x}) f_2(\mathbf{x}) P_{\mathbf{x}}(d\mathbf{x})$ for all $f_1, f_2 \in \mathcal{F}$.
 563 Then, the action $\triangleright_{\mathcal{F}}$ of the group $\mathbb{G}$ on the function space $\mathcal{F}$ is termed unitary if it preserves the inner
 564 product of the function space:

$$\begin{aligned} \langle f_1, f_2 \rangle_{P_{\mathbf{x}}} &= \langle g \triangleright_{\mathcal{F}} f_1, g \triangleright_{\mathcal{F}} f_2 \rangle_{P_{\mathbf{x}}} \quad \forall f_1, f_2 \in \mathcal{F}, g \in \mathbb{G} \\ \int_{\mathcal{X}} f_1(\mathbf{x}) f_2(\mathbf{x}) P_{\mathbf{x}}(d\mathbf{x}) &= \int_{\mathcal{X}} (g \triangleright_{\mathcal{F}} f_1)(\mathbf{x}) (g \triangleright_{\mathcal{F}} f_2)(\mathbf{x}) P_{\mathbf{x}}(d\mathbf{x}) \\ &= \int_{\mathcal{X}} f_1(g^{-1} \triangleright \mathbf{x}) f_2(g^{-1} \triangleright \mathbf{x}) P_{\mathbf{x}}(d\mathbf{x}) \\ &= \int_{g \triangleright \mathcal{X} = \mathcal{X}} f_1(\mathbf{x}) f_2(\mathbf{x}) P_{\mathbf{x}}(g \triangleright d\mathbf{x}). \end{aligned} \tag{24}$$

565 That is, the group action is unitary if $P_{\mathbf{x}}$ is a $\mathbb{G}$ -invariant measure $P_{\mathbf{x}}(g \triangleright d\mathbf{x}) = P_{\mathbf{x}}(d\mathbf{x})$, $\forall g \in$
 566 $\mathbb{G}, d\mathbf{x} \subseteq \mathcal{X}$. Note that an $\mathbb{G}$ -invariant measure (and inner product) exists whenever $\mathbb{G}$ is finite,
 567 because for any measure $\eta : \Sigma_{\Omega} \mapsto \mathbb{R}$, we can use the group-average trick to obtain one, given by
 568 $P_{\mathbf{x}}(\mathbb{X}) = \sum_{g \in \mathbb{G}} \eta(g \triangleright \mathbb{X})$.²

569 The importance of the Hilbert space structure is that it enables the definition of a unitary group
 570 representation. Unitary representations have a well-studied modular structure that allows their
 571 decomposition (Thm. E.8) into $\mathbb{G}$ -stable subspaces (Def. E.5), which is crucial for approximating
 572 symmetric function spaces using a finite set of basis elements. Let $\mathbb{I}_{\mathcal{F}} = \{\phi_i \mid \phi_i \in \mathcal{L}_{P_{\mathbf{x}}}^2\}_{i \in \mathbb{N}}$ be an
 573 orthogonal basis for the function space $\mathcal{F} = \text{span}(\mathbb{I}_{\mathcal{F}})$, so that any function $f \in \mathcal{F}$ can be represented
 574 by its basis expansion coefficients $\alpha = [\langle \phi_i, f \rangle_{P_{\mathbf{x}}}]_{i \in \mathbb{N}}$, since $f_{\alpha}(\mathbf{x}) = \sum_{i \in \mathbb{N}} \langle \phi_i, f \rangle_{P_{\mathbf{x}}} \phi_i(\mathbf{x})$. In this
 575 basis, the group action of $\mathbb{G}$ on $\mathcal{F}$ defines a unitary group representation mapping group elements to
 576 unitary linear integral operators on $\mathcal{F}$, which can be expressed in matrix form.

577 **Definition F.3** (Unitary group representation on a function space). *Let $\mathcal{F} = \mathcal{L}_{P_{\mathbf{x}}}^2 \mathcal{X}, \mathbb{R}$ be a separable*
 578 *Hilbert space of scalar-valued functions on a set $\mathcal{X}$ endowed with the symmetry group $\mathbb{G}$. Let $\mathbb{I}_{\mathcal{F}}$ be*
 579 *an orthogonal basis set spanning $\mathcal{F}$. Then, the group action of $\mathbb{G}$ on $\mathcal{F}$ (Def. F.1) defines a unitary*
 580 *group representation mapping group elements to unitary linear integral operators on $\mathcal{F}$:*

$$\begin{aligned} \rho_{\mathcal{F}} : \mathbb{G} &\longrightarrow \mathbb{U}(\mathcal{F}) \\ g &\longrightarrow \rho_{\mathcal{F}}(g), \quad \text{s.t.} \quad \rho_{\mathcal{F}}(g)^* = \rho_{\mathcal{F}}(g^{-1}). \end{aligned} \tag{25}$$

581 Each unitary operator $\rho_{\mathcal{F}}(g) : \mathcal{F} \mapsto \mathcal{F}$ admits an infinite-dimensional matrix representation with
 582 entries $[\rho_{\mathcal{F}}(g)]_{i,j} := \langle \hat{f}_i, g \triangleright_{\mathcal{F}} \hat{f}_j \rangle_{P_{\mathbf{x}}}$, which characterize how the group action transforms the chosen
 583 basis functions. Consequently, once the group representation for a chosen basis set is defined, the
 584 group action on a function $f_{\alpha} \in \mathcal{F}$ can be expressed as an (infinite-dimensional) matrix transforma-
 585 tion of its basis expansion coefficients α , given by:

$$[g \triangleright_{\mathcal{F}} f_{\alpha}](\cdot) := \sum_{i \in \mathbb{N}} \langle \hat{f}_i, g \triangleright_{\mathcal{F}} f_{\alpha} \rangle_{P_{\mathbf{x}}} \hat{f}_i(\cdot) = \sum_{i \in \mathbb{N}} \left(\sum_{j \in \mathbb{N}} \langle \hat{f}_i, g \triangleright_{\mathcal{F}} \hat{f}_j \rangle_{P_{\mathbf{x}}} \underbrace{\langle \hat{f}_j, f \rangle_{P_{\mathbf{x}}}}_{\alpha_j} \right) \hat{f}_i(\cdot). \tag{26}$$

586

²Such a $\mathbb{G}$ -invariant measure exists for any (finite or continuous) compact group. See discussion.