GEOMETRIC INVARIANCE OF GRADIENT-INDUCED COUPLING FOR CONTINUOUS 3D DYNAMICS

Anonymous authors
Paper under double-blind review

ABSTRACT

We develop an $E(3)$-equivariant graph neural network with continuous latent dynamics, using a shared continuous vector field across nodes and trajectories. Motivated by the regularities shared across diverse physical states, we use the same field across latent states. Because the same field is evaluated at different latent states, gradient updates induce cross-state coupling through the shared-field parameters. We characterize this *gradient-induced cross-state coupling* and introduce a geometric symmetry bias that makes it invariant under common latent transformations and dependent on relative latent configuration. Our experiments show that invariant coupling can be reflected in latent organization. Across 3D dynamics benchmarks, the continuous formulation remains competitive with established baselines in forecasting performance, even under random irregular temporal sampling. Importantly, the framework unifies continuous latent dynamics with explicit geometric control over cross-state coupling without sacrificing physical $E(3)$ equivariance.

1 INTRODUCTION

Learning physical dynamics requires models to represent interactions, geometric structure, and temporal evolution. Graph-based approaches provide a natural representation of interactions among physical entities (Battaglia et al., 2016; Sanchez-Gonzalez et al., 2019; 2020), while equivariant architectures incorporate geometric symmetries such as $E(3)$ equivariance (Thomas et al., 2018; Fuchs et al., 2020; Satorras et al., 2021). Recent work extends equivariant modeling to trajectory-level dynamics (Xu et al., 2024) and continuous-time dynamics while preserving geometric equivariance (Garcia Satorras et al., 2021; Liu et al., 2024; Knigge et al., 2024). In these approaches, the geometry and parameterization of the latent dynamics remain tied to the $E(3)$ action. This constraint limits the ability to design the latent geometry and parameterization independently of the $E(3)$ action.

We propose a modular framework built on an EGNN backbone (Satorras et al., 2021), where the latent dynamics operate on $E(3)$-invariant states while end-to-end $E(3)$ equivariance is preserved, allowing the latent dynamics to vary without redesigning the $E(3)$-equivariant physical pathway; see Figure 1. Our framework uses $E(3)$-invariant latent states to isolate the design of the latent dynamics from physical $E(3)$ equivariance, providing the ability to vary this design independently. Motivated by the fact that physical dynamics evolve *continuously* across diverse states while their governing regularities remain *shared*, we model nodes and trajectories with one shared continuous latent vector field over these $E(3)$-invariant latent states. Because the shared field is evaluated at different latent states, a gradient update induced at one state propagates through the shared parameters and alters the field evaluated at other states. We call this *gradient-induced cross-state coupling*.

Motivated by the $E(3)$ invariance of the latent states, we hypothesize that the influence of gradient updates across latent states should depend on their relative configuration rather than their absolute latent locations. We therefore introduce a symmetry-based inductive bias by imposing a geometric structure on the gradient-induced cross-state coupling, making it invariant under common transformations of latent states and dependent on relative latent configuration. We characterize this coupling by its parameter-gradient kernel, whose symmetry follows from the latent geometry and basis parameterization. We test this hypothesis by examining whether invariant relative-configuration bias is reflected in the organization of the learned latent representation. Our framework enables this test by allowing the latent design to vary independently of physical $E(3)$ equivariance.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

We compare three parameterizations of the shared continuous latent field. First, a polynomial basis on the latent space $\mathbb{R}^3$ yields generally non-invariant coupling. Second, a fixed Fourier basis on the latent space $\mathbb{R}^3$ yields translation-invariant coupling. Third, we place the latent flow on $\mathrm{SU}(2)$ and parameterize the shared field using the harmonic basis provided by the Peter–Weyl theorem (Knapp & Knapp, 1996), yielding coupling invariant under common latent group actions. The three-way comparison probes two factors: cross-state invariance and invariant geometry.

Cross-state invariance, latent organization, and forecasting performance are distinct properties and are therefore evaluated separately. Controlled systems with known physical variables enable direct examination of physical–latent organization, while established 3D dynamics benchmarks assess latent organization and forecasting across broader physical settings. To test the continuous formulation directly, we also evaluate random irregular temporal sampling.
Our contributions are:

- We propose a modular framework that uses $E(3)$-invariant latent states to isolate the design of the latent dynamics from physical $E(3)$ equivariance while preserving end-to-end $E(3)$ equivariance, with one shared continuous vector field across all nodes and trajectories.
- We characterize gradient-induced cross-state coupling arising from the shared field and introduce a geometric, symmetry-based inductive bias that makes this coupling invariant and dependent on relative latent configuration.

Related Work

Graph-based models provide relational structure for physical dynamics (Battaglia et al., 2016; Sanchez-Gonzalez et al., 2019; 2020), while equivariant architectures impose geometric symmetries such as $E(3)$ equivariance (Thomas et al., 2018; Fuchs et al., 2020; Satorras et al., 2021). Recent work extends these principles to trajectory-level dynamics (Xu et al., 2024) and continuous-time evolution (Garcia Satorras et al., 2021; Liu et al., 2024; Knigge et al., 2024), preserving geometric equivariance throughout dynamical modeling. Together, prior work extends $E(3)$ equivariance into the latent dynamics, tying their geometry and parameterization to the $E(3)$ action. In contrast, our framework uses $E(3)$-invariant latent states to isolate the design of the latent dynamics from physical $E(3)$ equivariance while preserving end-to-end $E(3)$ equivariance.

A complementary perspective concerns the role of parameterization in the function space. Sparse identification and geometric dictionaries use structured bases (Brunton et al., 2016; Champion et al., 2019; Huang et al., 2025), while harmonic and representation-theoretic constructions provide symmetry-adapted function spaces on groups (Cohen & Welling, 2016; Kondor & Trivedi, 2018; Cohen et al., 2019; 2018; Kondor et al., 2018; Weiler et al., 2018; Lang & Weiler, 2020; Xu et al., 2022b). Neural tangent kernel (NTK) theory relates parameter gradients to kernels governing function updates (Jacot et al., 2018; Arora et al., 2019; Li et al., 2019). Equivariant NTKs characterize group structure in these kernels (Misof et al., 2024), alongside work on equivariance induced by augmentation and ensembles (Gerken et al., 2022; Gerken & Kessel, 2024). Together, prior work connects structured parameterization to both the dynamical function space and the kernel induced by parameter gradients. Rather than specifying the kernel directly, we impose geometric structure through the latent geometry and basis parameterization, from which the parameter-gradient kernel and its symmetry arise implicitly.

2 Model Formulation

We consider three realizations of a shared continuous latent field: **R3-Poly**, **R3-Fourier**, and **SU2-PW**. R3-Poly and R3-Fourier evolve on $\mathbb{R}^3$ and differ only in the basis used to parameterize the shared field, whereas SU2-PW places the latent state on the compact Lie group $\mathrm{SU}(2)$ and parameterizes the field using the harmonic basis provided by the Peter–Weyl theorem.

2.1 Latent Geometries and Shared-Field Parameterizations

All three realizations evolve over a common latent phase interval $\zeta \in [0, \pi]$, discretized into K integration steps with $\Delta\zeta = \pi/K$. Physical observation and query times are linearly mapped to this latent phase interval, so that phase differences are proportional to physical time differences. Using

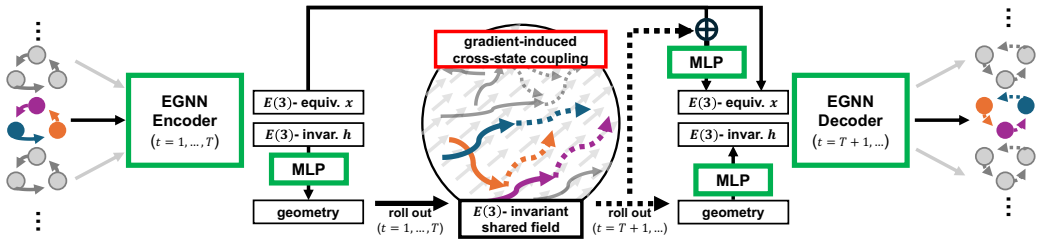

Figure 1: An overview of the model and its gradient-induced cross-state coupling .

the Lie-group logarithm $\mathrm{Log} : \mathrm{SU}(2) \to \mathfrak{su}(2)$, the corresponding distances on these latent domains are defined as

$$d_{\mathrm{R3}}(z_1, z_2) = \|z_1 - z_2\|_2, \ z_1, z_2 \in \mathbb{R}^3, \quad d_{\mathrm{SU2}}(q_1, q_2) = \left\|\mathrm{Log}(q_1^{-1} q_2)\right\|_2, \ q_1, q_2 \in \mathrm{SU}(2). \tag{1}$$

Euclidean latent dynamics. For $b \in \{\mathrm{R3\text{-}Poly}, \mathrm{R3\text{-}Fourier}\}$, node i has a latent state $z_i(\zeta) \in \mathbb{R}^3$. Let $\xi_b : \mathbb{R}^3 \to \mathbb{R}^3$ denote the shared field governing the nodewise latent trajectories, parameterized by the basis map ϕ_b. The evolution is

$$\frac{dz_i(\zeta)}{d\zeta} = \xi_b(z_i(\zeta)), \qquad z_i^{k+1} = z_i^k + \Delta\zeta\, \xi_b(z_i^k), \qquad k = 0, \ldots, K-1. \tag{2}$$

For R3-Poly, the shared field uses a complete degree-five monomial basis with S_{P} features:

$$\bar{\phi}_{\mathrm{P}}(z) = \left(z_1^a z_2^b z_3^c\right)_{a+b+c \le 5}, \qquad \phi_{\mathrm{P}}(z) = \sqrt{S_{\mathrm{P}}}\, \frac{\bar{\phi}_{\mathrm{P}}(z)}{\|\bar{\phi}_{\mathrm{P}}(z)\|_2}, \qquad \xi_{\mathrm{P}}(z) = W_{\mathrm{P}} \phi_{\mathrm{P}}(z). \tag{3}$$

The normalization is used to improve training stability and mitigate numerical instability.
For R3-Fourier, we use a fixed real Fourier basis with frequencies $\{\boldsymbol{\omega}_m\}_{m=1}^M \subset \mathbb{R}^3$:

$$\phi_{\mathrm{F}}(z) = \left[1, \left(\sqrt{2}\cos(\boldsymbol{\omega}_m^\top z), \sqrt{2}\sin(\boldsymbol{\omega}_m^\top z)\right)_{m=1}^M\right]^\top, \qquad \xi_{\mathrm{F}}(z) = W_{\mathrm{F}} \phi_{\mathrm{F}}(z). \tag{4}$$

Compact group latent dynamics. For SU2-PW, node i has a latent state $q_i(\zeta) \in \mathrm{SU}(2)$ with latent phase ζ. Writing L_q for left translation and $\mathrm{Exp} : \mathfrak{su}(2) \to \mathrm{SU}(2)$ for the Lie-group exponential map, let $\xi_{\boldsymbol{C}} : \mathrm{SU}(2) \to \mathfrak{su}(2) \simeq \mathbb{R}^3$ denote the shared Lie-algebra-valued velocity field governing the nodewise latent trajectories in left-trivialized coordinates. The evolution is

$$\frac{dq_i(\zeta)}{d\zeta} = \left(L_{q_i(\zeta)}\right)_* \xi_{\boldsymbol{C}}(q_i(\zeta)), \qquad q_i^{k+1} = q_i^k \,\mathrm{Exp}\left(\Delta\zeta\, \xi_{\boldsymbol{C}}(q_i^k)\right), \qquad k = 0, \ldots, K-1. \tag{5}$$

This rollout preserves the group-valued latent state, so every update remains in $\mathrm{SU}(2)$.
To parameterize this field, we use the harmonic structure of $\mathrm{SU}(2)$ given by the Peter–Weyl theorem.

Theorem 1 (Peter–Weyl theorem (Knapp & Knapp, 1996)). *If G is a compact group, the linear span of the matrix coefficients for all finite-dimensional irreducible unitary representations of G is dense in $L^2(G)$.*

With normalized Haar measure, Schur orthogonality implies that the normalized Wigner D-matrix coefficients form an orthonormal harmonic basis of $L^2(\mathrm{SU}(2))$. $D^{(j)} : \mathrm{SU}(2) \to U(d_j) \subset \mathbb{C}^{d_j \times d_j}$ denotes the complex-valued spin-j irreducible unitary representation, with $d_j = 2j + 1$.

$$\boldsymbol{\phi}_J^{\mathbb{C}}(q) = \bigoplus_{j \le J} \sqrt{d_j}\ \mathrm{vec}\, D^{(j)}(q) \in \mathbb{C}^{M_J}, \qquad M_J = \sum_{j \le J} d_j^2. \tag{6}$$

Because the shared field represents a left-trivialized latent velocity, its output must lie in the real Lie algebra $\mathfrak{su}(2) \simeq \mathbb{R}^3$. We parameterize this field using the complex Wigner D-matrix basis as

$$\xi_{\boldsymbol{C}}(q) = \sum_{j \le J} \sqrt{d_j} \sum_{m,n=-j}^{j} D_{mn}^{(j)}(q)\, \boldsymbol{C}_{mn}^{(j)} \in \mathfrak{su}(2) \simeq \mathbb{R}^3, \qquad \boldsymbol{C}_{mn}^{(j)} \in \mathbb{C}^3. \tag{7}$$

Since the Wigner D-matrix coefficients are generally complex, we design the trainable spectral coefficients $\boldsymbol{C}_{mn}^{(j)}$ to satisfy a conjugate-symmetry constraint that enforces the real-valued Lie-algebra output. Thus, conjugate-paired terms combine to yield a real-valued field; see Appendix A.1.

2.2 E(3)-EQUIVARIANT PHYSICAL ENCODING AND RECONSTRUCTION

Invariant latent encoding. For a graph-structured observation sequence $\{X(t_\ell)\}_{\ell=1}^{T}$ of N nodes, let $\boldsymbol{x}_i(t_\ell)$ and $\boldsymbol{v}_i(t_\ell)$ denote the position and velocity of node $i = 1, \dots, N$ at observation time t_ℓ. For $g = (\boldsymbol{R}, \boldsymbol{p}) \in E(3) = O(3) \ltimes \mathbb{R}^3$, $\boldsymbol{x}_i \mapsto \boldsymbol{R}\boldsymbol{x}_i + \boldsymbol{p}$, $\boldsymbol{v}_i \mapsto \boldsymbol{R}\boldsymbol{v}_i$. We encode $X(t)$ with the EGNN encoder $E^{\mathrm{EGNN}}_{\boldsymbol{\theta}_{\mathrm{enc}}}$ (Satorras et al., 2021). Its scalar hidden features $\boldsymbol{h}_i^{\mathrm{enc}}$ are $E(3)$-invariant, while its spatial coordinates $\boldsymbol{x}_i^{\mathrm{enc}}$ are $E(3)$-equivariant:

$$\big(\boldsymbol{h}_i^{\mathrm{enc}}, \boldsymbol{x}_i^{\mathrm{enc}}\big) = E^{\mathrm{EGNN}}_{\boldsymbol{\theta}_{\mathrm{enc}}}(X(t))_i, \quad \boldsymbol{h}_i^{\mathrm{enc}}(g{\cdot}X) = \boldsymbol{h}_i^{\mathrm{enc}}(X), \quad \boldsymbol{x}_i^{\mathrm{enc}}(g{\cdot}X) = \boldsymbol{R}\,\boldsymbol{x}_i^{\mathrm{enc}}(X) + \boldsymbol{p}. \tag{8}$$

The features $\boldsymbol{h}_i^{\mathrm{enc}}$ are mapped to the corresponding latent domain for each realization using branch-specific MLPs. For the two Euclidean realizations,

$$z_{i,b}(t) = \mathrm{MLP}^{\mathbb{R}^3}_{\mathrm{enc},b}(\boldsymbol{h}_i^{\mathrm{enc}}) \in \mathbb{R}^3, \qquad z_{i,b}(g{\cdot}X) = z_{i,b}(X), \qquad b \in \{\text{R3-Poly}, \text{R3-Fourier}\}. \tag{9}$$

while for SU2-PW an MLP latent head maps the invariant features to a unit quaternion,

$$\widetilde{\boldsymbol{q}}_i(t) = \mathrm{MLP}^{\mathrm{SU}(2)}_{\mathrm{enc}}(\boldsymbol{h}_i^{\mathrm{enc}}), \qquad q_i(t) = \frac{\widetilde{\boldsymbol{q}}_i(t)}{\|\widetilde{\boldsymbol{q}}_i(t)\|_2} \in S^3 \simeq \mathrm{SU}(2), \qquad q_i(g{\cdot}X) = q_i(X). \tag{10}$$

Since $\boldsymbol{h}_i^{\mathrm{enc}}$ is $E(3)$-invariant, so are $z_{i,b}$ and q_i in Eqs. 9 and 10. Because these initial latent states are $E(3)$-invariant and ξ_b and $\xi_{\boldsymbol{C}}$ are evaluated only on these states, the rollouts in Eqs. 2 and 5 remain $E(3)$-invariant. See Appendix A.3 for details.

Equivariant physical reconstruction. Although the latent dynamics are $E(3)$-invariant, reconstruction in physical space requires equivariant spatial information. We therefore retain the equivariant coordinates from Eq. 8 as a separate spatial pathway:

$$\boldsymbol{s}_i(t_T) = \boldsymbol{x}_i^{\mathrm{enc}}(t_T) - \boldsymbol{x}_i(t_T), \quad \boldsymbol{s}_i(g{\cdot}X, t_T) = (\boldsymbol{R}\boldsymbol{x}_i^{\mathrm{enc}} + \boldsymbol{p}) - (\boldsymbol{R}\boldsymbol{x}_i + \boldsymbol{p}) = \boldsymbol{R}\,\boldsymbol{s}_i(t_T). \tag{11}$$

The displacement $\boldsymbol{s}_i$ removes global translation so that the spatial correction transforms as a vector, retaining only the $O(3)$ action. Let $\ell_{i,\beta}(\tau)$ denote the evolved latent state for realization β at physical query offset τ, with β indexing the three realizations. The spatial pathway constructs $\boldsymbol{o}_{i,\beta}(\tau)$:

$$\boldsymbol{o}_{i,\beta}(\tau) = \alpha_{i,\beta}(\tau)\,\boldsymbol{s}_i(t_T), \qquad \alpha_{i,\beta}(\tau) = f_{\mathrm{sp},\beta}(\boldsymbol{h}_i^{\mathrm{enc}}(t_T), \ell_{i,\beta}(\tau), \tau) \in \mathbb{R}. \tag{12}$$

Since $f_{\mathrm{sp},\beta}$ is an MLP-based scalar function of $E(3)$-invariant inputs, $\alpha_{i,\beta}$ is $E(3)$-invariant. Since $\boldsymbol{s}_i$ is $E(3)$-equivariant, so is $\boldsymbol{o}_{i,\beta}$. We augment $\boldsymbol{x}_i(t_T) + \tau\,\boldsymbol{v}_i(t_T)$ with $\boldsymbol{o}_{i,\beta}(\tau)$ to introduce an equivariant spatial correction. A branch-specific MLP maps each evolved latent state to the decoder context $\boldsymbol{h}^{\mathrm{ctx}}_{i,\beta}(\tau) := \mathrm{MLP}_{\mathrm{dec},\beta}(\ell_{i,\beta}(\tau))$. The EGNN decoder combines the $E(3)$-equivariant physical coordinates with the $E(3)$-invariant latent contexts through graph message passing:

$$\widehat{\boldsymbol{x}}_{i,\beta}(t_T + \tau) = D^{\mathrm{EGNN}}_{\boldsymbol{\theta}_{\mathrm{dec},\beta}}\Big(\big\{\boldsymbol{h}^{\mathrm{ctx}}_{j,\beta}(\tau)\big\}_{j=1}^{N}, \{\boldsymbol{x}_j(t_T) + \tau\,\boldsymbol{v}_j(t_T) + \boldsymbol{o}_{j,\beta}(\tau)\}_{j=1}^{N}\Big)_i. \tag{13}$$

Therefore, the encoder–rollout–decoder map $\mathcal{F}_{\Theta,\beta} : (X, \tau) \mapsto \widehat{X}_\beta(t_T + \tau)$ is $E(3)$-equivariant:

$$\mathcal{F}_{\Theta,\beta}\big(g{\cdot}X; \tau\big) = g{\cdot}\mathcal{F}_{\Theta,\beta}\big(X; \tau\big), \qquad g \in E(3). \tag{14}$$

2.3 GRADIENT-INDUCED CROSS-STATE COUPLING

For a differentiable function $\xi_{\boldsymbol{\theta}}$, gradient updates at one latent state can influence function values at other states through the parameters shared across those evaluations (Jacot et al., 2018).

Proposition 1 (Gradient-induced cross-state coupling)**.** *Let $\boldsymbol{a}_\alpha$ denote the learning signal associated with a latent state ℓ_α. Under a gradient step $\Delta\boldsymbol{\theta} = -\eta\nabla_{\boldsymbol{\theta}}\mathcal{L}$, the first-order change of the field at a query state ℓ, by the chain rule and a first-order Taylor expansion, is*

$$\Delta\xi_{\boldsymbol{\theta}}(\ell) = -\eta \sum_{\alpha} \boldsymbol{\Theta}_{\boldsymbol{\theta}}(\ell, \ell_\alpha)\,\boldsymbol{a}_\alpha + O(\eta^2), \quad \boldsymbol{\Theta}_{\boldsymbol{\theta}}(\ell, r) = \mathcal{J}_{\boldsymbol{\theta}}(\ell)\mathcal{J}_{\boldsymbol{\theta}}(r)^\top, \quad \mathcal{J}_{\boldsymbol{\theta}}(\ell) = \frac{\partial \xi_{\boldsymbol{\theta}}(\ell)}{\partial \boldsymbol{\theta}}. \tag{15}$$

Here, $\boldsymbol{\Theta}_{\boldsymbol{\theta}}(\ell, r)$ is the cross-state kernel induced by $\xi_{\boldsymbol{\theta}}$, quantifying the first-order interaction between its evaluations $\xi_{\boldsymbol{\theta}}(\ell)$ and $\xi_{\boldsymbol{\theta}}(r)$ through shared parameters; see Appendix A.2 for the derivation.

The shared fields defined in Eqs. 3, 4, and 7 are indexed by β as

$$\big(\ell,\, \xi_\beta(\ell),\, \phi_\beta(\ell),\, \mathcal{W}_\beta\big) = \begin{cases} \big(z,\, \xi_{\mathrm{P}}(z),\, \phi_{\mathrm{P}}(z),\, W_{\mathrm{P}}\big), & \beta = \text{R3-Poly}, \\ \big(z,\, \xi_{\mathrm{F}}(z),\, \phi_{\mathrm{F}}(z),\, W_{\mathrm{F}}\big), & \beta = \text{R3-Fourier}, \\ \big(q,\, \xi_{\boldsymbol{C}}(q),\, \boldsymbol{\phi}_J(q),\, \{\boldsymbol{C}^{(j)}_{mn}\}_{j,m,n}\big), & \beta = \text{SU2-PW}. \end{cases}$$

For each realization β, let $\boldsymbol{\theta}_\beta$ denote the independent real parameters underlying the shared-field coefficient collection $\mathcal{W}_\beta$, with Jacobian $\mathcal{J}_{\boldsymbol{\theta}_\beta}(\ell) = \partial\xi_\beta(\ell)/\partial\boldsymbol{\theta}_\beta$. Since the shared fields $\xi_\beta(\ell)$ are linear in these independent parameters, the Jacobian $\mathcal{J}_{\boldsymbol{\theta}_\beta}$ is independent of the coefficient values $\mathcal{W}_\beta$. Hence, for updates of the shared-field coefficients, the relation in Proposition 1 is exact, with no $O(\eta^2)$ remainder. The gradient-induced cross-state kernel $\boldsymbol{\Theta}_\beta(\ell, r)$ therefore reduces to the scalar basis kernel $K_\beta(\ell, r)$ times $\boldsymbol{I}_3$:

$$\boldsymbol{\Theta}_\beta(\ell, r) = K_\beta(\ell, r)\boldsymbol{I}_3, \quad K_\beta(\ell, r) = \phi_\beta(\ell)^\top \phi_\beta(r), \quad \Delta\xi_\beta(\ell) = -\eta \sum_\alpha K_\beta(\ell, \ell_\alpha)\, \boldsymbol{a}_\alpha. \tag{16}$$

For R3-Poly, because the basis is constructed from the absolute Euclidean latent coordinates, the kernel is generally not invariant under a common translation:

$$K_{\mathrm{P}}(z + a, r + a) \neq K_{\mathrm{P}}(z, r), \qquad a \in \mathbb{R}^3. \tag{17}$$

For R3-Fourier, the kernel depends only on the relative displacement $z - r$ and is invariant under a common translation:

$$K_{\mathrm{F}}(z, r) = 1 + 2\sum_{m=1}^{M} \cos\big(\boldsymbol{\omega}_m^\top (z - r)\big) = K_{\mathrm{F}}(z + a, r + a), \qquad a \in \mathbb{R}^3. \tag{18}$$

For SU2-PW, let $\boldsymbol{\phi}_J(q) \in \mathbb{R}^{M_J}$ denote the real feature map induced by the conjugate-symmetric parameterization. As shown in Appendix A.1, realification preserves the basis inner product. For $g_L, g_R \in \mathrm{SU}(2)$, define $T_{g_L,g_R}(q) = g_L q g_R$. Then $T_{g_L,g_R}(q)^{-1} T_{g_L,g_R}(r) = g_R^{-1}(q^{-1}r)g_R$. Hence,

$$\begin{aligned} K_J(q, r) &= \boldsymbol{\phi}_J(q)^\top \boldsymbol{\phi}_J(r) = \sum_{j \le J} d_j\, \chi_j(q^{-1}r) = \kappa_J(q^{-1}r), \\ K_J(T_{g_L,g_R}(q), T_{g_L,g_R}(r)) &= \kappa_J\big(g_R^{-1}(q^{-1}r)g_R\big) = \kappa_J(q^{-1}r) = K_J(q, r). \end{aligned} \tag{19}$$

Here, $\chi_j(g) = \operatorname{Tr} D^{(j)}(g)$ is the character, and the second line follows from the conjugation invariance of irreducible characters. Thus, the SU2-PW kernel depends only on the relative group configuration $q^{-1}r$ and is invariant under common latent group actions.

R3-Poly induces generally non-invariant coupling, whereas R3-Fourier and SU2-PW induce invariant coupling through relative Euclidean and group configurations, respectively. Thus, R3-Poly versus R3-Fourier probes the effect of cross-state invariance within $\mathbb{R}^3$, while R3-Fourier versus SU2-PW tests whether the geometry used to realize invariant coupling also matters.

3 Evaluation Framework

We evaluate three aspects of the learned dynamics: the realized invariance of cross-state coupling, whether the relative-configuration bias is reflected in latent organization, and forecasting performance. These are evaluated separately rather than treated as a causal chain.

Comparison Setup. We compare the three realizations introduced: R3-Poly, R3-Fourier, and SU2-PW. Their shared-field capacities are chosen to be nearly matched. R3-Poly uses the complete degree-five monomial basis on $\mathbb{R}^3$, giving 56 features and $3 \times 56 = 168$ trainable shared-field coefficients. R3-Fourier uses $M = 27$ fixed frequency vectors, sampled once isotropically and held fixed during training, giving $1+2M = 55$ real features and $3\times55 = 165$ trainable coefficients. SU2-PW retains the Wigner D-matrix coefficients for $j \in \{0, \frac{1}{2}, 1, \frac{3}{2}, 2\}$, giving $\sum_{j\le2}(2j+1)^2 = 55$ basis features and $3 \times 55 = 165$ independent real trainable coefficients after realification.

Realized-support invariance defect. We evaluate cross-state kernel invariance on latent states visited by the learned dynamics. To probe realized-support invariance, we apply a common translation

T_b, $z \mapsto z + a$ for the two $\mathbb{R}^3$ realizations and a common group action $q \mapsto g_L q g_R$ for SU2-PW to pairs of learned latent states. We define the normalized realized-support invariance defect

$$H(\ell, r; T_b) = \frac{|K_b(T_b(\ell), T_b(r)) - K_b(\ell, r)|}{K_b(\ell, \ell)}, \qquad b \in \{\text{R3-Poly}, \text{R3-Fourier}, \text{SU2-PW}\}. \quad (20)$$

We report the mean and 95th percentile of H, with the mean summarizing the average realized defect and H_{95} capturing its upper tail; lower values indicate smaller invariance defect. R3-Fourier and SU2-PW are invariant by construction, so their measured defects should remain at numerical precision. For R3-Poly, H measures how strongly its translation non-invariance is exposed.

Latent organization. To examine whether the relative-configuration bias is reflected in the learned latent representation, we characterize latent organization through distances and neighborhoods among dynamical states. Using the distances in Eq. 1, let $\ell_{u,n}$ denote the latent state of node n in complete state u. We define

$$D(u,v) = \left[\frac{1}{N}\sum_{n=1}^{N} d(\ell_{u,n}, \ell_{v,n})^2\right]^{1/2}, \qquad (21)$$
$$G(u) = \left(d(\ell_{u,n}, \ell_{u,m})\right)_{n<m}, \qquad D_G(u,v) = \frac{1}{\sqrt{E}}\left\|G(u) - G(v)\right\|_2, \quad E = \binom{N}{2}.$$

D measures nodewise differences between complete latent states, while $G(u)$ collects all within-state pairwise distances and D_G compares their relational geometry. When explicit physical variables are available, physical–latent organization is quantified using Spearman rank correlation and distance correlation (dCor) (Székely et al., 2007).

Temporal locality and forecasting. For a trajectory, we quantify temporal locality by the percentile rank (PR) of the true next state among the other latent states from the same trajectory or source. Using D gives D-PR and using D_G gives G-PR; we report their means across temporal neighbor pairs, with lower PR indicating stronger temporal locality and more temporally coherent latent organization. Forecasting is evaluated by final-state MSE (F-MSE) and all-state MSE (A-MSE).

4 Experiments on Physical Dynamics

We evaluate whether structural differences in cross-state coupling are reflected in the organization of learned latent dynamics across two complementary settings. Controlled systems provide known physical variables, allowing direct evaluation of physical–latent organization in terms of the latent distances D and D_G. Standard dynamics benchmarks do not provide such explicit control variables, so we assess latent temporal locality and forecasting performance. Within each trajectory, states evolve under the same physical dynamics, so temporal adjacency provides a natural reference for latent organization. Realized-support invariance defect is evaluated in both settings. Implementation details and hyperparameters are provided in Appendix B.1.

4.1 Controlled Physical Dynamics

Control protocols. We construct two controlled systems with known physical variables. To eliminate arbitrary global-frame and node-order cues, we randomly apply a common rigid motion $\mathbf{x} \mapsto R\mathbf{x} + b$ and a node permutation. Three non-collinear reference nodes, fixed throughout the canonical trajectory, define the geometric reference plane relative to which the controlled motion is specified. Each system uses observations from a single time point as input and predicts future states sampled randomly at irregular times.

In the acceleration control, one node is already moving under constant rectilinear acceleration a at the input, so its observed velocity carries information about a (Figure. 2(a)). In the impulse control, the input is a pre-impulse state in which two three-node isosceles triangles A and B are at rest relative to the reference (Figure. 2(e)). Each is a scaled copy of the same template with unit surface density, $\rho = 1$, so its area equals its mass; the same fixed impulse J produces $V_A = +(J/m_A)\hat{\mathbf{z}}$ and $V_B = -(J/m_B)\hat{\mathbf{z}}$. After the impulse, both triangles move rectilinearly at constant velocity.

Across both systems, the controlled variables and node roles are not supplied to the model, and forecasting is nearly perfect, making predictive performance effectively non-discriminative. The

two controls retain rectilinear dynamics but differ in physical structure: Acceleration varies a single physical factor, whereas Impulse introduces two subsystem-dependent factors.

Table 1: Invariance defect H on the controlled systems. $< 10^{-15}$ reported as 0.

Control	R3-Poly		R3-Fourier		SU2-PW	
	Mean H	H_{95}	Mean H	H_{95}	Mean H	H_{95}
Acceleration	6.80×10^{-11}	2.94×10^{-10}	0	0	0	0
Impulse	1.57×10^{-2}	5.98×10^{-2}	0	0	0	0

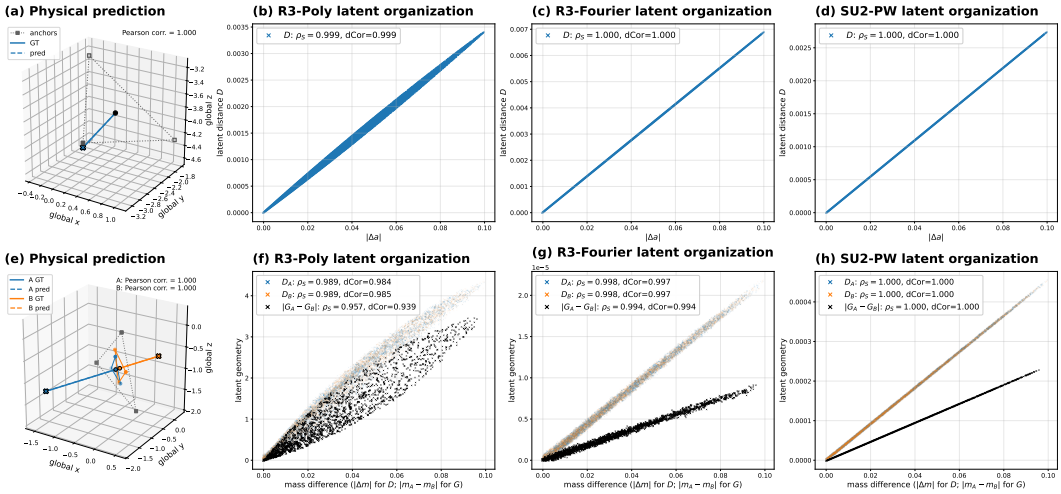

Figure 2: Acceleration and Impulse control.

Controlled-dynamics results. The controlled results are summarized in Table 1 and Figure 2. Latent organization is quantified by Spearman correlation (ρ_s) and distance correlation (dCor). In Acceleration, H is negligible across all realizations, and latent organization is uniformly near-perfect in both ρ_s and dCor. In Impulse, R3-Poly alone exhibits a clear invariance defect, while latent organization becomes progressively stronger from R3-Poly to R3-Fourier to SU2-PW.

The single-factor Acceleration control provides a baseline with uniformly strong latent organization across all realizations. In contrast, when the physical structure becomes subsystem-dependent in Impulse, the non-invariant realization exhibits weaker organization of relative latent configuration than the invariant realizations. Together, these results support the hypothesis that invariance in cross-state coupling can be reflected in the organization of relative latent configuration.

4.2 Standard Dynamics Benchmarks

Benchmark protocols. We follow established experimental setups for the three physical-dynamics benchmarks. For the charged N-body system, we use the 3D charged-particle benchmark derived from the interacting-particle system of Kipf et al. (2018) and subsequently adopted in 3D equivariant dynamics studies (Fuchs et al., 2020; Satorras et al., 2021; Xu et al., 2024). For CMU Motion Capture, we follow the Walk/Run setup used in prior equivariant dynamics studies (Huang et al., 2022; Han et al., 2022; Xu et al., 2024). For MD17 (Chmiela et al., 2017), we follow the dynamics-prediction setup and split of Huang et al. (2022) and Xu et al. (2024), including the latter's heavy-atom restriction, and construct two-hop molecular graphs as in prior geometric molecular models (Shi et al., 2021; Xu et al., 2022a). We use the standard setting of each benchmark as the baseline and additionally evaluate random irregular temporal sampling, as described in Appendix B.2. All results in Table 2 are obtained under random irregular temporal sampling.

Physical regimes. N-body consists of charged particles evolving under Coulomb interactions. CMU Run consists of articulated human motion with recurrent coordination through the gait cycle. MD17

Table 2: Invariance defect H and temporal locality on standard benchmarks. $< 10^{-15}$ reported as 0.

Dataset	Invariance defect						Mean D-PR			Mean G-PR		
	R3-Poly		R3-Fourier		SU2-PW		R3-Poly	R3-Fourier	SU2-PW	R3-Poly	R3-Fourier	SU2-PW
	Mean H	H_{95}	Mean H	H_{95}	Mean H	H_{95}						
N-body	9.12×10^{-4}	5.34×10^{-3}	0	0	0	0	0.1859	0.1850	0.1814	0.1859	0.1964	0.2145
CMU Run	3.63×10^{-3}	1.16×10^{-2}	0	0	0	0	0.0403	0.0161	0.0384	0.0441	0.0160	0.0395
Benzene	0	0	0	0	0	0	0.00447	0.00662	0.00324	0.00272	0.00357	0.00238
Aspirin	1.31×10^{-3}	8.11×10^{-3}	0	0	0	0	0.00222	0.00213	0.00220	0.00223	0.00214	0.00214

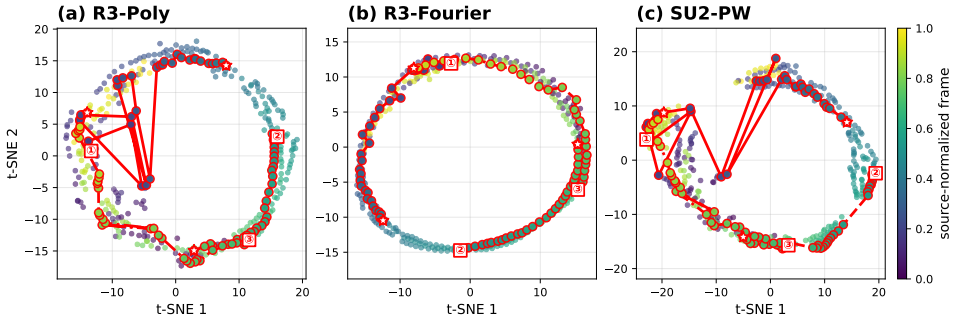

Figure 3: CMU Run D-t-SNE latent organization; paths 1–3 trace selected temporal trajectories.

probes bounded internal molecular dynamics: Benzene is a small, highly symmetric aromatic ring, whereas Aspirin is larger, less symmetric, and structurally more heterogeneous.

Benchmark results. As summarized in Table 2, R3-Poly yields $H = 0$ in the highly symmetric Benzene system despite its generally non-invariant coupling, while N-body and Aspirin show only modest differences in D-PR and G-PR across the three realizations. Thus, non-invariance need not be exposed on the latent support, and coupling invariance need not be consistently reflected in temporal latent organization.

In contrast to N-body and Aspirin, CMU Run and Benzene show clearer geometry-specific organization. In CMU Run, R3-Fourier yields the lowest D-PR and G-PR, with panel (b) of Figure 3 showing the clearest qualitative temporal organization. This organization is consistent with recurrently revisited configurations through the gait cycle being well matched to the periodic parameterization of the Fourier field. In Benzene, SU2-PW yields the lowest D-PR and G-PR, with its stronger organization consistent with the bounded dynamics and high molecular symmetry being well matched to the compact group geometry. See Appendix B.3 for D- and D_G-based t-SNE visualizations.

Together, these results indicate that the correspondence between invariant cross-state coupling and latent organization is system-dependent, with different latent geometries being compatible with particular physical dynamics.

Forecasting and Irregular Temporal Sampling. Tables 3 and 4 report F-MSE and A-MSE under regular and random irregular (**RI**) temporal sampling; MD17 A-MSE results are provided in Appendix B.4. Across the benchmarks, the three realizations remain competitive with established forecasting baselines, while forecasting performance does not consistently track H, D-PR, or G-PR. F-MSE and A-MSE remain comparable between regular and RI sampling, indicating that the continuous formulation accommodates random irregular sampling without systematic degradation.

Ablation. On N-body with RI sampling, removing the equivariant spatial pathway consistently increases F-MSE across all three latent realizations and rollout resolutions, whereas changing number of rollout steps does not produce a consistent performance trend. See Appendix B.5 for details.

5 Limitations and Discussion

The exact symmetry characterization applies specifically to the shared latent field and its coefficient kernel in Eq. 16. Accordingly, invariance of the shared-field coupling does not imply invariance of

Table 3: MSE ($\times 10^{-2}$) on N-body and CMU Motion.

	N-body		CMU #Walk		CMU #Run	
Method	F-MSE ↓	A-MSE ↓	F-MSE ↓	A-MSE ↓	F-MSE ↓	A-MSE ↓
MPNN (Gilmer et al., 2017)	1.07	–	36.1±1.5	–	66.4±2.2	–
TFN (Thomas et al., 2018)	1.55	–	32.0±1.8	–	56.6±1.7	–
RF (Köhler et al., 2019)	1.04	–	188.0±1.9	–	521.3±2.3	–
SE(3)-Tr. (Fuchs et al., 2020)	2.44	–	31.5±2.1	–	61.2±2.3	–
EGNN (Satorras et al., 2021)	0.71	–	28.7±1.6	–	50.9±0.9	–
EGNO (Xu et al., 2024)	**0.54**	**0.22**	8.1±1.6	3.5±0.5	33.9±1.7	14.9±0.9
Ours R3-Poly	0.61	0.25	7.9±1.4	3.9±0.4	31.9±1.4	14.0±0.7
Ours R3-Fourier	0.61	0.24	**7.5**±1.4	**3.4**±0.5	31.8±1.7	13.9±0.8
Ours SU2-PW	0.57	0.23	9.1±2.0	4.1±0.7	**27.9**±0.9	**12.8**±0.7
Ours R3-Poly + RI	0.58	0.25	8.2±1.8	3.9±0.6	29.2±1.1	14.4±0.6
Ours R3-Fourier + RI	0.62	0.26	7.8±1.7	3.7±0.5	28.2±1.2	13.7±0.6
Ours SU2-PW + RI	0.60	0.25	7.6±1.7	3.6±0.5	30.9±1.5	15.4±0.8

Table 4: F-MSE ($\times 10^{-3}$), MD17.

Method	Aspirin	Benzene	Ethanol	Malonald.	Naphthal.	Salicylic	Toluene	Uracil
TFN (Thomas et al., 2018)	123.7±1.8	584.8±19.8	48.1±0.4	136.2±0.8	4.9±0.1	10.3±0.2	108.9±0.1	8.4±0.2
RF (Köhler et al., 2019)	109.4±0.1	1037.2±12.9	46.4±0.1	139.3±0.3	5.0±0.1	12.3±0.1	109.3±0.4	6.4±0.1
SE(3)-Tr. (Fuchs et al., 2020)	111.2±0.6	681.1±6.7	47.4±1.3	138.9±0.2	5.2±0.1	11.3±0.2	108.8±0.6	7.9±0.2
EGNN (Satorras et al., 2021)	144.1±1.5	624.0±5.3	46.4±0.1	136.4±0.1	4.7±0.2	10.2±0.2	117.8±0.7	6.4±0.1
EGNO (Xu et al., 2024)	**91.8**±0.6	488.5±5.5	**46.2**±0.1	128.0±0.2	3.7±0.1	8.6±0.2	102.1±0.5	**5.2**±0.2
Ours R3-Poly	92.6±0.7	437.4±4.0	46.3±0.1	**127.9**±0.2	**3.6**±0.1	8.9±0.2	103.0±0.6	5.3±0.2
Ours R3-Fourier	92.8±0.6	481.3±6.8	46.3±0.1	128.1±0.1	**3.6**±0.1	9.1±0.2	100.9±0.4	5.5±0.2
Ours SU2-PW	93.7±0.5	500.2±4.9	47.0±0.1	128.3±0.2	**3.6**±0.1	8.8±0.2	101.7±0.5	5.5±0.2
Ours R3-Poly + RI	93.7±0.5	438.6±4.2	46.7±0.1	**127.9**±0.2	3.9±0.1	9.2±0.2	**100.5**±0.5	5.8±0.2
Ours R3-Fourier + RI	92.7±0.7	**427.1**±6.1	47.0±0.1	128.1±0.2	3.9±0.1	**8.5**±0.1	101.8±0.6	5.8±0.2
Ours SU2-PW + RI	92.9±0.4	461.8±5.5	46.6±0.1	**127.9**±0.2	3.9±0.1	9.1±0.2	100.8±0.4	5.7±0.2

the full model's parameter-gradient kernel. Likewise, H measures invariance only over the realized latent support and should not be interpreted as a global invariance guarantee.

The motivation for invariant cross-state coupling is not that latent organization should necessarily translate into lower forecasting error, but to impose a geometric inductive bias on cross-state learning. Our results show that this structure can be reflected in a more coherent latent organization, while forecasting performance does not consistently track that organization. This separation suggests that latent organization should be viewed as a representational property of the learned dynamics rather than a direct proxy for forecasting performance. Determining when such structured representations yield practical forecasting gains remains an important direction for further study.

The relationship between invariant cross-state coupling and latent organization is system-dependent. Relative-configuration invariance may also become restrictive when absolute context is informative or when anisotropic dynamics distinguish orientations that the kernel treats as equivalent. Consistent with this dependence, the benchmark results in Section 4.2 suggest that different latent geometries may be better suited to different physical dynamics.

6 Conclusion

We introduced a modular framework for continuous physical dynamics that preserves end-to-end $E(3)$ equivariance while allowing latent geometry and field parameterization to be varied. A shared continuous vector field induces gradient-mediated cross-state coupling whose symmetry follows from the latent design. Our controlled experiments support the hypothesis that invariance in cross-state coupling can be reflected in the organization of relative latent configuration. Across standard benchmarks, this correspondence is system-dependent, with different latent geometries being compatible with particular physical dynamics. At the same time, the continuous formulation remains competitive in forecasting and accommodates random irregular temporal sampling without systematic degradation. More broadly, the framework positions latent geometry not only as a representation choice, but also as a mechanism for structuring cross-state learning.

AI USE STATEMENT

In this work, we did not use generative AI tools for tasks requiring disclosure under the ICLR 2027 AI Policy. Generative AI tools were used to assist with editing, rephrasing, and polishing portions of the manuscript for clarity and readability. We have reviewed all AI-assisted work and verified the final manuscript. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We will publicly release the source code, experiment configurations, and reproduction scripts upon publication. The model formulation and evaluation procedure are described in Sections 2, 3, and 4. Additional details on the end-to-end $E(3)$-equivariant rollout and random irregular temporal sampling are provided in Appendices A.3 and B.2, respectively. For the standard dynamics benchmarks, we follow the established experimental setups and corresponding public implementations used in prior work, including EGNO (Xu et al., 2024).

REFERENCES

Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *Advances in neural information processing systems*, 32, 2019.

Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. *Advances in neural information processing systems*, 29, 2016.

Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.

Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116 (45):22445–22451, 2019.

Stefan Chmiela, Alexandre Tkatchenko, Huziel E Sauceda, Igor Poltavsky, Kristof T Schütt, and Klaus-Robert Müller. Machine learning of accurate energy-conserving molecular force fields. *Science advances*, 3(5):e1603015, 2017.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.

Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018.

Taco S Cohen, Mario Geiger, and Maurice Weiler. A general theory of equivariant cnns on homogeneous spaces. *Advances in neural information processing systems*, 32, 2019.

Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-translation equivariant attention networks. *Advances in neural information processing systems*, 33:1970–1981, 2020.

Victor Garcia Satorras, Emiel Hoogeboom, Fabian Fuchs, Ingmar Posner, and Max Welling. E (n) equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 34:4181–4192, 2021.

Jan Gerken, Oscar Carlsson, Hampus Linander, Fredrik Ohlsson, Christoffer Petersson, and Daniel Persson. Equivariance versus augmentation for spherical images. In *International Conference on Machine Learning*, pp. 7404–7421. PMLR, 2022.

Jan E Gerken and Pan Kessel. Emergent equivariance in deep ensembles. *arXiv preprint arXiv:2403.03103*, 2024.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pp. 1263–1272. Pmlr, 2017.

Jiaqi Han, Wenbing Huang, Tingyang Xu, and Yu Rong. Equivariant graph hierarchy-based neural networks. *Advances in Neural Information Processing Systems*, 35:9176–9187, 2022.

Daning Huang, Hanyang He, John Harlim, and Yan Li. Learning vector fields of differential equations on manifolds with geometrically constrained operator-valued kernels. In *International Conference on Learning Representations*, volume 2025, pp. 19541–19573, 2025.

Wenbing Huang, Jiaqi Han, Yu Rong, Tingyang Xu, Fuchun Sun, and Junzhou Huang. Equivariant graph mechanics networks with constraints. *arXiv preprint arXiv:2203.06442*, 2022.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *International conference on machine learning*, pp. 2688–2697. Pmlr, 2018.

Anthony W Knapp and Anthony William Knapp. *Lie groups beyond an introduction*, volume 140. Springer, 1996.

David M Knigge, David R Wessels, Riccardo Valperga, Samuele Papa, Jan-Jakob Sonke, Efstratios Gavves, and Erik J Bekkers. Space-time continuous pde forecasting using equivariant neural fields. *Advances in neural information processing systems*, 37:76553–76577, 2024.

Jonas Köhler, Leon Klein, and Frank Noé. Equivariant flows: sampling configurations for multi-body systems with symmetric energies. *arXiv preprint arXiv:1910.00753*, 2019.

Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International conference on machine learning*, pp. 2747–2755. PMLR, 2018.

Risi Kondor, Zhen Lin, and Shubhendu Trivedi. Clebsch–gordan nets: a fully fourier space spherical convolutional neural network. *Advances in neural information processing systems*, 31, 2018.

Leon Lang and Maurice Weiler. A wigner-eckart theorem for group equivariant convolution kernels. *arXiv preprint arXiv:2010.10952*, 2020.

Zhiyuan Li, Ruosong Wang, Dingli Yu, Simon S Du, Wei Hu, Ruslan Salakhutdinov, and Sanjeev Arora. Enhanced convolutional neural tangent kernels. *arXiv preprint arXiv:1911.00809*, 2019.

Yang Liu, Jiashun Cheng, Haihong Zhao, Tingyang Xu, Peilin Zhao, Fugee Tsung, Jia Li, and Yu Rong. Segno: Generalizing equivariant graph neural networks with physical inductive biases. In *International Conference on Learning Representations*, volume 2024, pp. 22764–22787, 2024.

Philipp Misof, Pan Kessel, and Jan E Gerken. Equivariant neural tangent kernels. *arXiv preprint arXiv:2406.06504*, 2024.

Alvaro Sanchez-Gonzalez, Victor Bapst, Kyle Cranmer, and Peter Battaglia. Hamiltonian graph networks with ode integrators. *arXiv preprint arXiv:1909.12790*, 2019.

Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pp. 8459–8468. Pmlr, 2020.

Vıctor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pp. 9323–9332. PMLR, 2021.

Chence Shi, Shitong Luo, Minkai Xu, and Jian Tang. Learning gradient fields for molecular conformation generation. In *International conference on machine learning*, pp. 9558–9568. PMLR, 2021.

Gábor J Székely, Maria L Rizzo, and Nail K Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 2007.

Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018.

Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. *Advances in Neural information processing systems*, 31, 2018.

Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*, 2022a.

Minkai Xu, Jiaqi Han, Aaron Lou, Jean Kossaifi, Arvind Ramanathan, Kamyar Azizzadenesheli, Jure Leskovec, Stefano Ermon, and Anima Anandkumar. Equivariant graph neural operator for modeling 3d dynamics. *arXiv preprint arXiv:2401.11037*, 2024.

Yinshuang Xu, Jiahui Lei, Edgar Dobriban, and Kostas Daniilidis. Unified fourier-based kernel and nonlinearity design for equivariant networks on homogeneous spaces. In *International Conference on Machine Learning*, pp. 24596–24614. PMLR, 2022b.

A Mathematical Details and Proofs

A.1 Real-Valued Shared Field from Complex Wigner Matrix Coefficients

The Wigner D-matrix coefficients are generally complex, whereas the shared latent field represents a left-trivialized velocity and therefore takes values in the real Lie algebra $\mathfrak{su}(2) \simeq \mathbb{R}^3$. For $m, n \in \{-j, -j+1, \ldots, j\}$, the shared field is

$$\xi_{\boldsymbol{C}}(q) = \sum_{j \le J} \sqrt{d_j} \sum_{m,n=-j}^{j} D^{(j)}_{mn}(q)\, \boldsymbol{C}^{(j)}_{mn}, \qquad \boldsymbol{C}^{(j)}_{mn} \in \mathbb{C}^3. \tag{22}$$

We constrain the spectral coefficients by

$$\boldsymbol{C}^{(j)}_{-m,-n} = (-1)^{m-n} \overline{\boldsymbol{C}^{(j)}_{mn}}. \tag{23}$$

For each independent conjugate pair, we parameterize

$$\boldsymbol{C}^{(j)}_{mn} = \frac{\boldsymbol{u}^{(j)}_{mn} + i\, \boldsymbol{v}^{(j)}_{mn}}{\sqrt{2}}, \qquad \boldsymbol{u}^{(j)}_{mn}, \boldsymbol{v}^{(j)}_{mn} \in \mathbb{R}^3, \tag{24}$$

while each self-conjugate coefficient is parameterized directly as a real vector.

Under the standard Wigner D-matrix phase convention, the coefficients satisfy

$$D^{(j)}_{-m,-n}(q) = (-1)^{m-n} \overline{D^{(j)}_{mn}(q)}. \tag{25}$$

Hence each conjugate pair contributes

$$\begin{aligned} D^{(j)}_{mn}(q)\boldsymbol{C}^{(j)}_{mn} + D^{(j)}_{-m,-n}(q)\boldsymbol{C}^{(j)}_{-m,-n} \\ = 2\ \mathrm{Re}\Big[D^{(j)}_{mn}(q)\boldsymbol{C}^{(j)}_{mn}\Big] \in \mathbb{R}^3. \end{aligned} \tag{26}$$

Thus $\xi_{\boldsymbol{C}}(q) \in \mathfrak{su}(2) \simeq \mathbb{R}^3$ is real-valued by construction.

Equivalently, the independent real coefficients in Eq. 24 can be collected into $W \in \mathbb{R}^{3 \times M_J}$ together with a real feature map $\boldsymbol{\phi}_J(q) \in \mathbb{R}^{M_J}$ such that

$$\xi_{\boldsymbol{C}}(q) = W \boldsymbol{\phi}_J(q). \tag{27}$$

For a conjugate pair, the corresponding real features are proportional to $\sqrt{2d_j}\,\mathrm{Re}\, D^{(j)}_{mn}(q)$ and $-\sqrt{2d_j}\,\mathrm{Im}\, D^{(j)}_{mn}(q)$; self-conjugate coefficients, when present, contribute their real feature directly. This realification preserves the kernel inner product. Using the complex feature map from Eq. equation 6 and the unitarity of the irreducible representations,

$$\begin{aligned} \boldsymbol{\phi}_J(q)^\top \boldsymbol{\phi}_J(r) &= \mathrm{Re}\Big[\big(\boldsymbol{\phi}^{\mathbb{C}}_J(q)\big)^\dagger \boldsymbol{\phi}^{\mathbb{C}}_J(r)\Big] \\ &= \mathrm{Re}\left[\sum_{j \le J} d_j\ \mathrm{Tr}\Big(D^{(j)}(q)^\dagger D^{(j)}(r)\Big)\right] \\ &= \mathrm{Re}\left[\sum_{j \le J} d_j\ \mathrm{Tr}\Big(D^{(j)}(q^{-1}r)\Big)\right] \\ &= \sum_{j \le J} d_j\, \chi_j(q^{-1}r), \end{aligned} \tag{28}$$

where $\chi_j(g) = \mathrm{Tr}\, D^{(j)}(g)$ is the character of the spin-j irreducible representation. The final equality uses the fact that irreducible characters of $\mathrm{SU}(2)$ are real-valued. Hence the realification preserves exactly the SU2-PW basis kernel used in Eq. equation 19.

A.2 Proof of Gradient-Induced Cross-State Coupling

Let α index an occurrence at which the shared field $\xi_{\boldsymbol{\theta}}$ is evaluated during the training computation, with latent input ℓ_α. Define the corresponding backpropagated learning signal by

$$\boldsymbol{a}_\alpha = \nabla_{\xi_{\boldsymbol{\theta}}(\ell_\alpha)} \mathcal{L} \in \mathbb{R}^{d_\xi}. \tag{29}$$

Here, $\boldsymbol{a}_\alpha$ is the total learning signal arriving at that field evaluation through the computational graph. The parameter Jacobian is taken with the latent input held fixed:

$$\mathcal{J}_{\boldsymbol{\theta}}(\ell) = \frac{\partial \xi_{\boldsymbol{\theta}}(\ell)}{\partial \boldsymbol{\theta}}. \tag{30}$$

Therefore, by the chain rule, the gradient contribution from all evaluations of the shared field can be written as

$$\nabla_{\boldsymbol{\theta}} \mathcal{L} = \sum_\alpha \mathcal{J}_{\boldsymbol{\theta}}(\ell_\alpha)^\top \boldsymbol{a}_\alpha. \tag{31}$$

Proof of Proposition 1. Assume that the parameter Jacobian is locally Lipschitz in $\boldsymbol{\theta}$. For a fixed query latent state ℓ, a Taylor expansion with respect to the shared-field parameters gives

$$\Delta \xi_{\boldsymbol{\theta}}(\ell) = \mathcal{J}_{\boldsymbol{\theta}}(\ell) \Delta \boldsymbol{\theta} + O\big(\|\Delta \boldsymbol{\theta}\|_2^2\big). \tag{32}$$

The query state ℓ is held fixed in this perturbation; thus, Eq. equation 32 characterizes the parameter-induced change of the shared field itself rather than the additional change that may arise from perturbing the latent trajectory.

Substituting $\Delta \boldsymbol{\theta} = -\eta \nabla_{\boldsymbol{\theta}} \mathcal{L}$ and Eq. equation 31 yields

$$\begin{aligned} \Delta \xi_{\boldsymbol{\theta}}(\ell) &= -\eta \sum_\alpha \mathcal{J}_{\boldsymbol{\theta}}(\ell) \mathcal{J}_{\boldsymbol{\theta}}(\ell_\alpha)^\top \boldsymbol{a}_\alpha + O(\eta^2) \\ &= -\eta \sum_\alpha \boldsymbol{\Theta}_{\boldsymbol{\theta}}(\ell, \ell_\alpha) \boldsymbol{a}_\alpha + O(\eta^2), \end{aligned} \tag{33}$$

where

$$\boldsymbol{\Theta}_{\boldsymbol{\theta}}(\ell, r) = \mathcal{J}_{\boldsymbol{\theta}}(\ell) \mathcal{J}_{\boldsymbol{\theta}}(r)^\top. \tag{34}$$

This establishes Proposition 1. □

For the R3-Poly, R3-Fourier, and SU2-PW shared fields used in this work, $\xi_\beta(\ell)$ is linear in its independent real coefficient parameters. Consequently, at fixed latent states, $\mathcal{J}_{\boldsymbol{\theta}_\beta}(\ell)$ is independent of those coefficient values, and the parameter perturbation is exactly linear:

$$\Delta \xi_\beta(\ell) = \mathcal{J}_{\boldsymbol{\theta}_\beta}(\ell) \Delta \boldsymbol{\theta}_\beta. \tag{35}$$

Hence the Taylor remainder in Eq. equation 33 vanishes for updates of the shared-field coefficients, yielding the exact cross-state update and basis-kernel form stated in Eq. equation 16.

A.3 Latent Rollout and End-to-End E(3) Equivariance

We first establish that the latent states remain invariant under the physical $E(3)$ action throughout rollout and query-time interpolation. Let

$$z_{i,b}^0 := z_{i,b}(t_T), \qquad q_i^0 := q_i(t_T)$$

denote the latent states initialized from the final observed frame. From Eqs. equation 9 and equation 10, the initial latent states satisfy

$$\begin{aligned} z_{i,b}^0(g \cdot X) &= z_{i,b}^0(X), \qquad b \in \{\text{R3-Poly}, \text{R3-Fourier}\}, \\ q_i^0(g \cdot X) &= q_i^0(X), \qquad g \in E(3). \end{aligned} \tag{36}$$

Euclidean latent rollout. For the R3-Poly and R3-Fourier realizations, the latent rollout from Eq. equation 2 is

$$z_{i,b}^{k+1} = z_{i,b}^{k} + \Delta\zeta\, \xi_b(z_{i,b}^{k}), \qquad b \in \{\text{R3-Poly}, \text{R3-Fourier}\}. \tag{37}$$

Assume inductively that $z_{i,b}^{k}(g\cdot X) = z_{i,b}^{k}(X)$. Because the shared field is evaluated at the same latent state,

$$\xi_b\big(z_{i,b}^{k}(g\cdot X)\big) = \xi_b\big(z_{i,b}^{k}(X)\big)\,. \tag{38}$$

Applying Eq. equation 37 therefore gives

$$z_{i,b}^{k+1}(g\cdot X) = z_{i,b}^{k+1}(X). \tag{39}$$

Together with Eq. equation 36, this establishes $E(3)$ invariance of every Euclidean latent rollout state.

SU2-PW latent rollout. For SU2-PW, the latent update from Eq. equation 5 is

$$q_i^{k+1} = q_i^{k}\,\mathrm{Exp}\big(\Delta\zeta\, \xi_{\boldsymbol{C}}(q_i^{k})\big)\,. \tag{40}$$

Since $\xi_{\boldsymbol{C}}(q_i^{k}) \in \mathfrak{su}(2)$,

$$\mathrm{Exp}\big(\Delta\zeta\, \xi_{\boldsymbol{C}}(q_i^{k})\big) \in \mathrm{SU}(2). \tag{41}$$

Hence, if $q_i^{k} \in \mathrm{SU}(2)$, group closure implies $q_i^{k+1} \in \mathrm{SU}(2)$. Because the encoder initializes $Q^0 \in \mathrm{SU}(2)^N$, every state of the SU2-PW rollout remains in $\mathrm{SU}(2)^N$.

To establish physical-frame invariance, assume inductively that $q_i^{k}(g\cdot X) = q_i^{k}(X)$. Then

$$\xi_{\boldsymbol{C}}\big(q_i^{k}(g\cdot X)\big) = \xi_{\boldsymbol{C}}\big(q_i^{k}(X)\big)\,, \tag{42}$$

and Eq. equation 40 gives

$$q_i^{k+1}(g\cdot X) = q_i^{k+1}(X). \tag{43}$$

Thus every SU2-PW rollout state is both group-valued and $E(3)$-invariant.

Query-time interpolation. Query states between adjacent rollout points are obtained by geodesic interpolation on the corresponding latent geometry. For $\zeta \in [\zeta_k, \zeta_{k+1}]$, define

$$\lambda = \frac{\zeta - \zeta_k}{\zeta_{k+1} - \zeta_k} \in [0, 1]. \tag{44}$$

For the Euclidean realizations,

$$z_{i,b}(\zeta) = (1-\lambda) z_{i,b}^{k} + \lambda z_{i,b}^{k+1}, \qquad b \in \{\text{R3-Poly}, \text{R3-Fourier}\}. \tag{45}$$

For SU2-PW, the corresponding group geodesic interpolation is

$$q_i(\zeta) = q_i^{k}\,\mathrm{Exp}\big(\lambda\,\mathrm{Log}\big((q_i^{k})^{-1} q_i^{k+1}\big)\big)\,. \tag{46}$$

Since $\mathrm{Log}((q_i^{k})^{-1} q_i^{k+1}) \in \mathfrak{su}(2)$, Eq. equation 46 remains in $\mathrm{SU}(2)$.

The physical $E(3)$ action does not change observation or query times, and therefore leaves the latent phase ζ and interpolation weight λ unchanged. Since the endpoints in Eqs. equation 45 and equation 46 are invariant, the interpolated query states satisfy

$$\begin{aligned} z_{i,b}(\zeta; g\cdot X) &= z_{i,b}(\zeta; X), \\ q_i(\zeta; g\cdot X) &= q_i(\zeta; X). \end{aligned} \tag{47}$$

Hence, for every realization β, the evolved latent state supplied to the decoder satisfies

$$\ell_{i,\beta}(g\cdot X, \tau) = \ell_{i,\beta}(X, \tau). \tag{48}$$

Equivariant physical reconstruction. From Eq. equation 11, the retained spatial displacement transforms as

$$\boldsymbol{s}_i(g\cdot X, t_T) = \boldsymbol{R}\,\boldsymbol{s}_i(X, t_T). \tag{49}$$

The scalar gate from Eq. equation 12 is

$$\alpha_{i,\beta}(\tau) = f_{\mathrm{sp},\beta}(\boldsymbol{h}_i^{\mathrm{enc}}(t_T), \ell_{i,\beta}(\tau), \tau)\,. \tag{50}$$

Its inputs are $E(3)$-invariant: $\boldsymbol{h}_i^{\mathrm{enc}}$ is invariant by the EGNN encoder, $\ell_{i,\beta}(\tau)$ is invariant by Eq. equation 48, and the physical query offset τ is unchanged by the spatial $E(3)$ action. Therefore

$$\alpha_{i,\beta}(g\cdot X, \tau) = \alpha_{i,\beta}(X, \tau). \tag{51}$$

It follows that

$$\begin{aligned}\boldsymbol{o}_{i,\beta}(g\cdot X, \tau) &= \alpha_{i,\beta}(g\cdot X, \tau)\,\boldsymbol{s}_i(g\cdot X, t_T)\\ &= \boldsymbol{R}\,\boldsymbol{o}_{i,\beta}(X, \tau).\end{aligned} \tag{52}$$

The coordinate supplied to the EGNN decoder therefore transforms as

$$\begin{aligned}&\boldsymbol{R}\boldsymbol{x}_i(t_T) + \boldsymbol{p} + \tau\,\boldsymbol{R}\boldsymbol{v}_i(t_T) + \boldsymbol{R}\boldsymbol{o}_{i,\beta}(\tau)\\ &\quad= \boldsymbol{R}(\boldsymbol{x}_i(t_T) + \tau\,\boldsymbol{v}_i(t_T) + \boldsymbol{o}_{i,\beta}(\tau)) + \boldsymbol{p}.\end{aligned} \tag{53}$$

At the same time, the branch-specific decoder context

$$\boldsymbol{h}_{i,\beta}^{\mathrm{ctx}}(\tau) = \mathrm{MLP}_{\mathrm{dec},\beta}(\ell_{i,\beta}(\tau)) \tag{54}$$

is $E(3)$-invariant:

$$\boldsymbol{h}_{i,\beta}^{\mathrm{ctx}}(g\cdot X, \tau) = \boldsymbol{h}_{i,\beta}^{\mathrm{ctx}}(X, \tau). \tag{55}$$

The EGNN decoder therefore receives equivariant coordinates together with invariant scalar contexts. By its $E(3)$ equivariance,

$$\mathcal{F}_{\Theta,\beta}\big(g\cdot X; \tau\big) = g\cdot\mathcal{F}_{\Theta,\beta}\big(X; \tau\big), \qquad g \in E(3), \tag{56}$$

for

$$\beta \in \{\mathrm{R3\text{-}Poly}, \mathrm{R3\text{-}Fourier}, \mathrm{SU2\text{-}PW}\}.$$

This establishes Eq. equation 14 for all three latent realizations.

B Additional Experimental Details and Results

B.1 Hyperparameters and Implementation Details

We use the same backbone and training configuration across R3-Poly, R3-Fourier, and SU2-PW within each experiment. For the standard physical-dynamics benchmarks, we follow the corresponding data protocols and backbone settings of EGNO (Xu et al., 2024). Table 5 summarizes the experiment-specific hyperparameters.

Table 5: Experiment-specific hyperparameters. P is the number of queried future states and K is the number of latent integration steps. "EGNN layers" denotes the number of message-passing layers in each of the encoder and decoder. For the spatial pathway, the value in parentheses is the hidden dimension of $f_{\mathrm{sp},\beta}$ in Eq. 12.

Experiment	Batch	LR	WD	Max epochs	EGNN layers	Hidden	P	K	Spatial
Controls	64	5×10^{-4}	10^{-10}	2000	3	64	5	32	Off
N-body	100	1×10^{-4}	10^{-8}	5000	4	64	5	16	On (32)
CMU Motion	12	5×10^{-4}	10^{-10}	5000	6	128	5	16	On (32)
MD17	100	1×10^{-4}	10^{-15}	5000	5	64	8	16	On (32)

Training. All models are optimized with Adam and selected according to validation F-MSE. We use an early-stopping patience of 100 validation evaluations for the Controls and 50 validation evaluations for the physical benchmarks. For the Controls, the gradient norm is clipped at 1.0.

Latent dynamics. All three realizations evolve over the common latent phase interval $\zeta \in [0, \pi]$. R3-Poly uses the complete degree-five monomial basis, R3-Fourier uses $M = 27$ fixed frequency vectors, and SU2-PW uses $j \in \{0, \frac{1}{2}, 1, \frac{3}{2}, 2\}$, i.e., $j_{\max} = 2$. Their nearly matched shared-field capacities are detailed in Section 3. We use $K = 32$ for the Controls and $K = 16$ for the physical benchmarks, consistent with the rollout-resolution study in Appendix B.5.

B.2 Random Irregular Temporal Sampling

We implement random irregular temporal sampling by varying the intermediate target times while preserving the original prediction horizon and final target. Let s denote the initial frame index, Δ the prediction horizon in frames, and P the number of queried future states. For each sample, we use ordered offsets

$$0 < o_1 < \cdots < o_{P-1} < o_P = \Delta, \tag{57}$$

giving target frame indices

$$s + o_1, \ldots, s + o_{P-1}, s + \Delta. \tag{58}$$

In the notation of Section 2, the physical query offsets are $\tau_k = o_k$. Their normalized offsets are

$$\bar{\tau}_k = \frac{o_k}{\Delta}, \qquad k = 1, \ldots, P, \tag{59}$$

with $\bar{\tau}_P = 1$. These query times are linearly mapped to the latent phase interval $[0, \pi]$, so that $\zeta_k = \pi\bar{\tau}_k$.

N-body. For the charged N-body benchmark, the prediction interval is from frame $f_0 = 30$ to frame $f_T = 40$, giving $\Delta = 10$, with $P = 5$ queried future states. Under random irregular sampling, the $P - 1$ intermediate target frames are sampled without replacement from

$$\{31, 32, \ldots, 39\}, \tag{60}$$

while frame 40 is always retained as the final target. Equivalently, the intermediate offsets are sampled from $\{1, \ldots, 9\}$, with the final offset fixed to 10. Thus, for a sampled frame f,

$$\bar{\tau} = \frac{f - 30}{10}. \tag{61}$$

Intermediate targets are resampled during training, whereas validation and test use fixed seeded schedules.

CMU Motion Capture. For CMU Motion Capture, each trajectory segment starts at frame s and uses a prediction horizon of $\Delta = 30$ frames with $P = 5$ queried future states. Four intermediate offsets are sampled without replacement from

$$\{1, 2, \ldots, 29\}, \tag{62}$$

after which the final offset 30 is appended and the targets are ordered chronologically. The corresponding normalized query offsets are $\bar{\tau}_k = o_k/30$. Intermediate targets are resampled during training, while validation and test retain fixed seeded schedules.

MD17. For MD17, we use a prediction horizon of $\Delta = 3000$ frames with $P = 8$ queried future states, and define candidate target times on a 16-point grid over this horizon. For a sample beginning at frame s, the candidate offsets are

$$b_j = \mathrm{round}\left(\frac{j\Delta}{16}\right), \qquad j = 1, \ldots, 16, \tag{63}$$

where the resulting offsets are required to be distinct. The $P - 1$ intermediate offsets are sampled without replacement from

$$\{b_1, \ldots, b_{15}\}, \tag{64}$$

and the final offset $b_{16} = \Delta$ is always retained. The resulting normalized query offsets are $\bar{\tau}_k = o_k/\Delta$. As in the other benchmarks, intermediate targets are resampled during training, whereas validation and test use fixed seeded schedules.

B.3 Latent-Trajectory Organization

Figures 4, 5, and 6 visualize latent-trajectory organization using the D and D_G distances defined in the main text. The D-t-SNE panels show complete latent-state organization, whereas the D_G-t-SNE panels show within-state relational geometry. Each point represents a latent state and is colored by temporal position; the red paths (1–3) trace selected trajectories or windows in temporal order, with the square and star marking their start and end states.

The figures compare R3-Poly, R3-Fourier, and SU2-PW for N-body and CMU Run, and separately for Benzene and Aspirin in MD17. These embeddings provide only qualitative views; D-PR and G-PR are computed directly in the original D and D_G metric spaces.

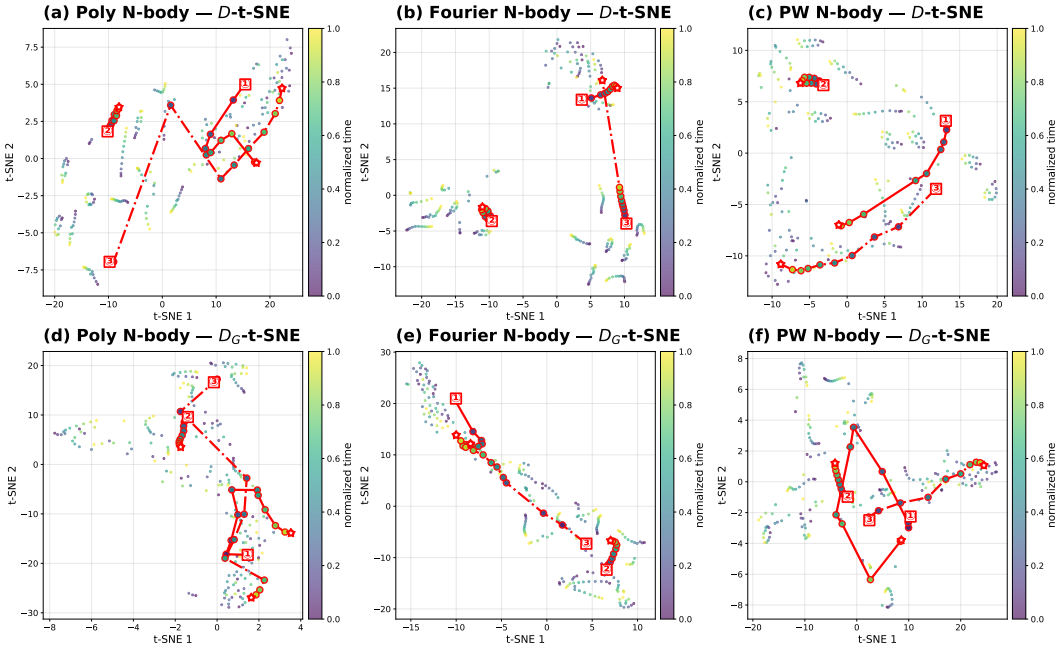

Figure 4: N-body latent-trajectory organization.

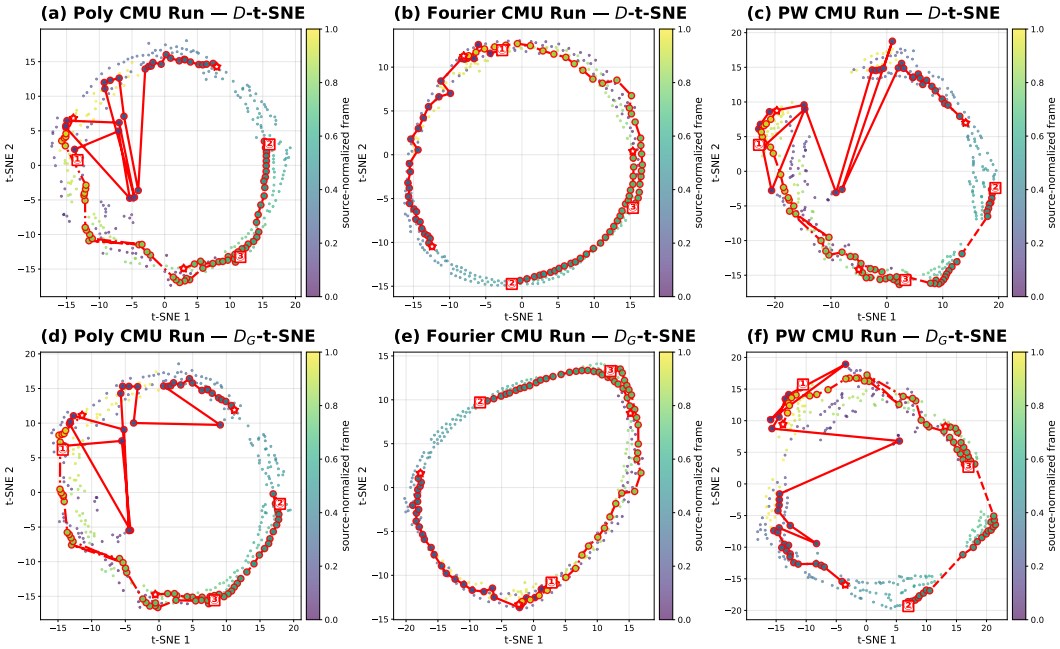

Figure 5: CMU Run latent-trajectory organization.

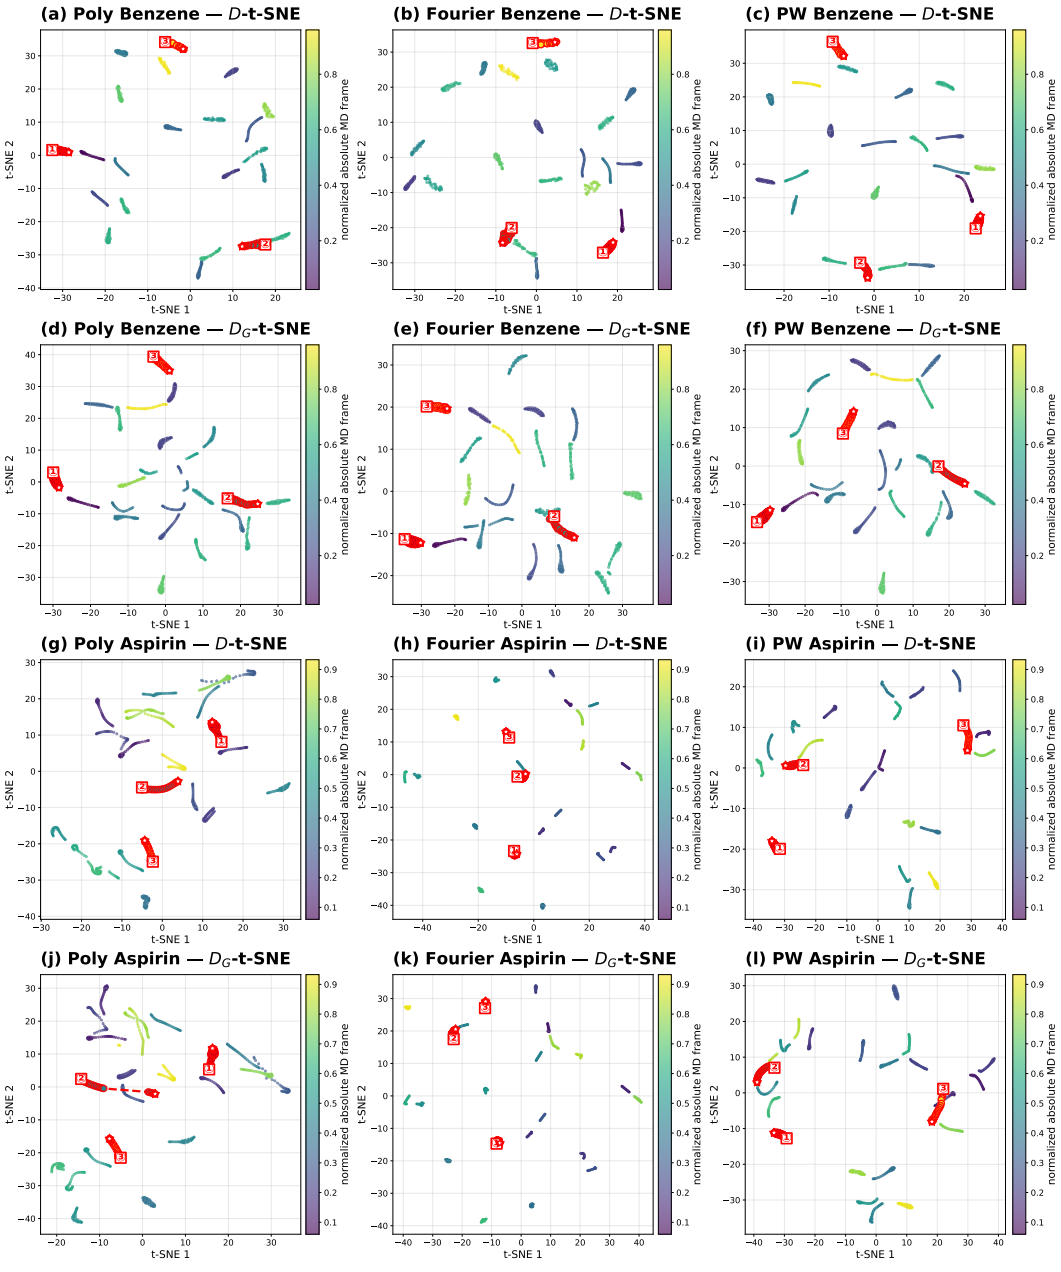

Figure 6: MD17 latent-trajectory organization.

B.4 MD17 A-MSE RESULTS

Table 6 reports the all-state forecasting error on MD17 under regular and random irregular temporal sampling. A-MSE complements the final-state evaluation in Table 4 by measuring prediction accuracy across all evaluated temporal states. The results show that the three latent realizations remain comparable under regular and RI sampling across the molecular systems.

Table 6: A-MSE ($\times 10^{-3}$), MD17.

Method	Aspirin	Benzene	Ethanol	Malonald.	Naphthal.	Salicylic	Toluene	Uracil
EGNO (Xu et al., 2024)	73.7±0.7	224.1±3.1	32.8±0.2	106.7±0.1	3.2±0.1	**7.7**±0.1	45.8±0.3	4.7±0.1
Ours R3-Poly	70.1±0.8	**174.9**±2.2	32.9±0.3	107.3±0.1	**2.8**±0.1	8.6±0.2	46.2±0.3	**4.6**±0.1
Ours R3-Fourier	**68.1**±0.6	199.1±2.9	32.9±0.3	107.4±0.1	2.9±0.1	8.8±0.2	**44.9**±0.2	4.9±0.1
Ours SU2-PW	70.0±0.8	204.8±3.2	33.2±0.3	107.4±0.2	2.9±0.1	8.0±0.1	45.0±0.4	5.1±0.1
Ours R3-Poly + RI	71.0±0.9	179.3±2.5	32.9±0.3	104.5±0.1	4.1±0.2	9.4±0.3	45.5±0.3	5.7±0.2
Ours R3-Fourier + RI	69.5±0.9	176.7±2.6	33.2±0.2	**104.2**±0.1	4.2±0.2	8.6±0.2	46.1±0.3	5.8±0.2
Ours SU2-PW + RI	69.8±0.9	195.7±2.9	**32.7**±0.1	104.7±0.1	4.2±0.2	10.9±0.2	45.6±0.3	5.6±0.1

B.5 ABLATION STUDIES

We conduct all ablations on N-body with RI sampling.

Rollout resolution. We evaluate 8, 16, and 32 latent rollout steps across all three realizations. With the spatial pathway enabled, performance varies only modestly across rollout resolutions for R3-Poly and R3-Fourier, while SU2-PW improves from 8 to 16 steps and remains unchanged at 32 steps. Overall, increasing the rollout resolution does not yield a consistent performance improvement across the three realizations. We use 32 rollout steps for the controlled systems and 16 rollout steps for the physical benchmarks.

Spatial pathway. We also ablate the equivariant spatial pathway used for physical reconstruction. Because the latent states are $E(3)$-invariant, this pathway carries equivariant spatial information through the displacement $\boldsymbol{s}_i = \boldsymbol{x}_i^{\text{enc}} - \boldsymbol{x}_i$. Removing the spatial pathway consistently increases F-MSE across all three latent realizations and all rollout resolutions. This is consistent with the role of the spatial pathway in retaining equivariant physical information for reconstruction.

Table 7: Ablation of rollout resolution and the spatial pathway. Results are F-MSE ($\times 10^{-2}$). Lower is better.

Rollout steps	Spatial pathway	R3-Poly	R3-Fourier	SU2-PW
8	Off	0.77	0.80	0.92
8	On	0.59	0.63	0.68
16	Off	0.82	0.79	0.92
16	On	0.59	0.65	0.61
32	Off	0.78	0.75	0.88
32	On	0.60	0.65	0.61