GFC: GEOMETRY-FIRST CORRUPTION FOR GENERATIVE IMITATION LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative policies model expressive action distributions for robotic manipulation, but standard Gaussian corruption does not explicitly account for dependencies across the action horizon. To address this issue, we propose GFC: *geometry-first corruption*, which generates noisy trajectories by randomly deforming a trajectory-level Riemannian metric. The metric captures action geometry and temporal coupling, refined by cross-segment predictive responses to structured perturbations. At each forward step, a random self-adjoint field deforms the metric through a matrix exponential, and a metric-gradient flow induces trajectory motion without an additional coordinate-noise draw. This connects demonstration geometry along the action horizon with stochastic corruption along noising time. We derive path-wise bounds on trajectory motion and predictive response, and specify reverse-learning objectives and terminal-source requirements. Experiments in simulation and on real robots demonstrate improved trajectory quality and policy performance. More broadly, GFC establishes stochastic geometry as a design principle for generative modeling, allowing learned data structure to shape the corruption process itself. Videos and supplementary materials are available at **https://anonymous.4open.science/r/GFC/**.

1 INTRODUCTION

Generative policies model multimodal action distributions for robotic manipulation through diffusion or flow-based generation (Chi et al., 2023; Lipman et al., 2023; Zhang et al., 2025; Black et al., 2024). These policies often generate action chunks, making dependencies across the prediction horizon central to coherent execution (Zhao et al., 2023; Black et al., 2025).

Expert trajectories contain coordinated variations across time: changing one segment can alter the continuations supported by demonstrations. The manifold hypothesis suggests that such high-dimensional data may concentrate near lower-dimensional sets (Loaiza-Ganem et al., 2024). This motivates studying trajectory geometry even when individual actions have Euclidean coordinates. Standard diffusion policies use isotropic Gaussian corruption without explicitly adapting it to these dependencies (Ho et al., 2020; Chi et al., 2023). Although the reverse model can learn them, the forward process determines the intermediate distributions and thus the denoising problems encountered during training. How should demonstration geometry shape these distributions?

Structured noise and learned forward processes offer ways to adapt corruption to data structure (Voleti et al., 2022; Chang et al., 2024; Daras et al., 2023; Bartosh et al., 2024). Riemannian geometry provides a complementary perspective by assigning different costs to different directions of variation. It supports generation on manifolds (De Bortoli et al., 2022b; Huang et al., 2022; Chen & Lipman, 2024), while data-induced metrics can shape interpolation paths and local noise (Kapuśniak et al., 2024; Bamberger et al., 2026b). We investigate a further possibility: making the trajectory metric itself a stochastic variable and deriving corruption from its deformations.

We propose *geometry-first corruption* (GFC), a forward process organized along two axes. Along the action horizon, cross-segment predictive responses refine a trajectory metric with temporal coupling. Along noising time, random metric deformations induce trajectory updates through a metric-gradient flow. This makes learned trajectory geometry an explicit part of each forward transition. We establish finite-step bounds on trajectory motion and predictive response, and formulate reverse-learning objectives with the source-matching conditions required for generation.

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.

Contributions. i): We introduce a Riemannian noise kernel that, to the best of our knowledge, is the first to inject stochasticity directly into the Riemannian metric and subsequently use the resulting deformed geometry to induce corrupted trajectories. This mechanism provides a geometry-aware corruption kernel for the forward process. ii): We further develop a two-axis geometry-first generative framework that integrates the proposed kernel with response-based trajectory refinement, and formulates the resulting forward chain together with source matching and reverse-learning objectives. iii): We extensively validate GFC in simulation and on real robots, demonstrating improved trajectory quality and policy performance across diverse geometries and manipulation tasks.

2 Related Work

Our work connects two research directions. The first studies generative policies and the construction of their forward processes, especially in trajectory corruptions. Diffusion and flow-based policies model multimodal action sequences and have been extended toward efficient sampling, vision-language conditioning, and streaming execution (Chi et al., 2023; Prasad et al., 2024; Zhang et al., 2025; Black et al., 2024; Jiang et al., 2025). Beyond standard Gaussian noising, general corruptions, learned forward processes, and stochastic interpolants provide more flexible probability paths (Daras et al., 2023; Bartosh et al., 2024; Albergo et al., 2025; Falorsi et al., 2019; De Bortoli et al., 2022a). Non-isotropic noise, temporal correlation, function-space models, and Gaussian-process priors further encode structure in generated trajectories (Voleti et al., 2022; Chang et al., 2024; Biloš et al., 2023; Kerrigan et al., 2024; Kollovieh et al., 2025; Molu, 2026).

The second direction studies how geometric structure can be encoded and exploited for trajectory generation and generative modeling. Robot learning methods construct task or demonstration metrics for policy composition and trajectory generation (Ratliff et al., 2018; Cheng et al., 2021; Bylard et al., 2021; Beik-Mohammadi et al., 2021; 2023; Rozo et al., 2025; Li et al., 2026), while Riemannian diffusion and flow matching generate manifold and robot trajectories (De Bortoli et al., 2022b; Huang et al., 2022; Chen & Lipman, 2024; Braun et al., 2024; Ding et al., 2025). Learned pullback, score-derived, and data-induced metrics can shape interpolation paths or local noise (Arvanitidis et al., 2018; de Kruiff et al., 2024; Azeglio & Di Bernardo, 2025; Saito & Matsubara, 2026; Kapuśniak et al., 2024; Bamberger et al., 2026b;a). Random and time-dependent geometries provide related mathematical foundations (Arnaudon et al., 2008; LaGatta & Wehr, 2014; Clarke et al., 2015; Arnaudon et al., 2019). Motivated by Sobolev curve metrics and pullback information geometry (Bauer et al., 2023; Arvanitidis et al., 2022), we learn a trajectory metric from cross-segment predictive responses. We randomly deform this metric and induce coordinate motion from the resulting geometry, making the metric itself part of the transition mechanism.

3 Method: Two-Axis Geometry-First Forward Process

The method constructs trajectory geometry along horizon h and randomizes it along time t. Each step follows

$$\tau_{t_k} \longrightarrow g^{\mathcal{M}}_{\psi,s_k,\cdot} \longrightarrow \widetilde{g}^{\omega_k}_{s_k} \longrightarrow \tau_{t_{k+1}}, \qquad s_k \in (t_k, t_{k+1}], \tag{1}$$

where $g^{\mathcal{M}}_{\psi,s_k,\cdot}$ denotes the metric fields evaluated at the current internal-flow state. We first define the base metric and random metric kernel, then learn the horizon refinement and assemble the forward chain. Notations, more backgrounds and theoretical proofs are in the appendix B, C and D.

3.1 Base trajectory geometry

Let $\tau = (x_0, \cdots, x_H) \in \mathcal{M} := SE(3)^{H+1}$ be a trajectory manifold, where $x_h \in \mathcal{P} := SE(3)$ is a pose manifold. The tangent space of τ is $\Xi = (\xi_0, \cdots, \xi_H) \in T_\tau\mathcal{M}$, where $\xi_h := T_{x_h}\mathcal{P} \simeq R^6$. The Riemannian metric on $\mathcal{P}$ is denoted $g^{\mathcal{P}}_x : T_x\mathcal{P} \times T_x\mathcal{P} \to R$ and $g^{\mathcal{P}}_x(\xi, \psi) = \langle \xi, \psi \rangle_W = u_h^\top W v_h$, where $\xi \in \Xi$, $\psi \in \Psi \in T_\tau\mathcal{M}$ and $W \succ 0$. Here $\vee$ maps a Lie-algebra element to its coordinate vector and $u_h = (\xi_h)^\vee, v_h = (\psi_h)^\vee,\ u_h, v_h \in \mathbb{R}^6$. The Riemannian metric on $\mathcal{M}$ ($g^{\mathcal{M}}_\tau$) is then denoted as

$$g^{\mathcal{M}}_\tau := \bigoplus_{h=0}^{H} g^{\mathcal{P}}_{x_h}. \tag{2}$$

Note that $\bigoplus$ is used schematically to indicate the aggregation of pose metrics into a trajectory metric; in practice, the specific aggregation operator can be chosen according to the requirements of the application.

Here in this paper, we want to take into account the smoothness between poses, by combining $\xi_{h+1} - \xi_h$ with Eq. 2. However, ξ_{h+1} and ξ_h belong to different tangent spaces, which cannot be compared by direct subtraction. Left trivialization is employed to identify the tangent spaces at different poses with the common Lie algebra at the identity, following the standard Lie-group formulation in (Solà et al., 2018). Let $L_{x_h^{-1}} : \mathcal{P} \to \mathcal{P},\ L_{x_h^{-1}}(x_h) = \mathfrak{e}$, where $\mathfrak{e} = \left(\begin{smallmatrix} \mathbf{I}_3 & 0 \\ 0 & 1 \end{smallmatrix}\right) \in SE(3)$. Let $\mathrm{d}L_{x_h^{-1}} : T_{x_h}\mathcal{P} \to T_e\mathcal{P}, s.t.\ \mathrm{d}L_{x_h^{-1}}(\xi_h) \in T_e\mathcal{P}$, the left trivialized tangent vector is defined as $\tilde{\xi}_h := \mathrm{d}L_{x_h^{-1}}(\xi_h) \in \mathfrak{se}(3)$. Let $D_h^{LT}\Xi := \frac{\tilde{\xi}_{h+1} - \tilde{\xi}_h}{\Delta t} \in \mathfrak{se}(3)$, where 'LT' stands for left trivialization and $\Delta_h > 0$ is the physical sampling interval. $g_\tau^{\mathcal{M}}$ is then specified as

$$g_\tau^{\mathcal{M}}(\Xi, \Psi) = \sum_{h=0}^{H} g_{x_h}^{\mathcal{P}}(\xi_h, \psi_h) + \lambda \sum_{h=0}^{H-1} \langle D_h^{LT}\Xi, D_h^{LT}\Psi \rangle_{\mathfrak{se}(3)}. \tag{3}$$

3.2 Random metric kernel

The noise kernel plays a central role in the forward corruption process. Existing Riemannian noise mechanisms, such as Exp-map-based perturbations (Falorsi et al., 2019) and heat-kernel diffusion (De Bortoli et al., 2022a), define stochastic transitions over samples under a fixed underlying geometry. In contrast, we propose a geometry-aware noise kernel (Random metric kernel) that first deforms the Riemannian metric and then induces the corresponding stochastic transition in the sample space. *Although the random metric kernel is introduced in the context of $SE(3)$, the construction is not specific to $SE(3)$ and can extend naturally to arbitrary Riemannian manifolds.*

Random deformation of the metric. Let $G^{\mathcal{M}}$ denote the space of Riemannian metrics on $\mathcal{M}$, and let Z_t be a random element taking values in a noise space $\mathcal{Z}$. A general metric deformation operator is defined as $\Phi_t : G^{\mathcal{M}} \times \mathcal{Z} \to G^{\mathcal{M}}$, which maps a base metric $g_\tau^{\mathcal{M}}$ and a random element Z_t to a random metric

$$\tilde{g}_{t,\tau} = \Phi_{\varepsilon(t)}(g_\tau^{\mathcal{M}}, Z_t). \tag{4}$$

The random element Z_t may contain one or multiple stochastic components and exhibit time dependence. Z_0 recovers the base geometry $\Phi_0(g_\tau^{\mathcal{M}}, Z_0) = g_\tau^{\mathcal{M}}$.

Different realizations of the metric-deformation operator are possible, depending on how the stochastic deformation is parameterized and constrained. In this work, a simple instantiation using a time correlated noise process ε_t and a self-adjoint random deformation field S_τ^ω sampled from S_τ, i.e., $Z_t = (\varepsilon_t, S_\tau^\omega)$, so that

$$g_\tau^{\mathcal{M}}(S_\tau^w \Xi, \Psi) = g_\tau^{\mathcal{M}}(\Xi, S_\tau^w \Psi) \tag{5}$$

The resulting operator takes the form

$$\tilde{g}_{t,\tau}^w(\Xi, \Psi) = g_\tau^{\mathcal{M}}(e^{\frac{\varepsilon_t}{2} S_\tau^w}\Xi, e^{\frac{\varepsilon_t}{2} S_\tau^w}\Psi). \tag{6}$$

Metric-induced trajectory update. Metric deformation alone does not produce a new trajectory, so an additional mechanism is required to induce a corrupted trajectory. One option is to introduce a second random field, but this would obscure whether the corruption arises from the deformed geometry or from direct trajectory-space noise. We therefore use a single random field to generate the deformation metric in Eq. 6.

Several single-field realizations are possible. The same field may jointly determine the metric deformation and trajectory update, or the deformed metric may deterministically guide the trajectory, e.g., along a Riemannian gradient direction. We adopt the latter, while the framework is not restricted to gradient-based updates: its key requirement is that a single random field induces trajectory corruption through the geometry.

To quantify the local magnitude of the metric deformation, we define

$$\phi_t^\omega(\tau) := \frac{1}{2} \left\| \varepsilon_t S_\tau^\omega \right\|_{HS, g_\tau^{\mathcal{M}}}^2 = \frac{\varepsilon_t^2}{2} \left\| S_\tau^\omega \right\|_{HS, g_\tau^{\mathcal{M}}}^2. \tag{7}$$

For an endomorphism $B : T_\tau\mathcal{M} \to T_\tau\mathcal{M}$, the corresponding Hilbert–Schmidt norm is

$$\|B\|^2_{HS,g^{\mathcal{M}}_\tau} := \sum_{i=1}^{6(H+1)} g^{\mathcal{M}}_\tau(Be_i, Be_i), \tag{8}$$

where $\{e_i\}_{i=1}^{6(H+1)}$ is any $g^{\mathcal{M}}_\tau$-orthonormal basis of $T_\tau\mathcal{M}$, i.e., $g^{\mathcal{M}}_\tau(e_i, e_j) = \delta_{ij}$.

We then define the deformation-induced vector field using the random metric $\tilde{g}^\omega_{t,\tau}$ from Eq. 6:

$$V^\omega_t(\tau) := \operatorname{grad}_{\tilde{g}^w_{t,\tau}} \phi^\omega_t(\tau). \tag{9}$$

Proposition 1 (Gradient under metric deformation)**.** *Under the metric deformation in Eq. 6, the vector field in Eq. 9 satisfies*

$$V^\omega_t(\tau) = e^{-\varepsilon_t S^\omega_\tau} \operatorname{grad}_{g^{\mathcal{M}}_\tau} \phi^\omega_t(\tau). \tag{10}$$

Finally, the corrupted sample is induced by the integral curve of this vector field. Let $\nu \in [0, 1]$ and $\gamma^\omega_t(\nu) \in \mathcal{M}$ satisfy

$$\frac{\partial \gamma^\omega_t(\nu)}{\partial \nu} = V^\omega_t(\gamma^\omega_t(\nu)), \qquad \gamma^\omega_t(0) = \tau, \qquad F^{\tilde{g}^w_{t,\tau}}_t(\tau, \omega) = \gamma^\omega_t(1). \tag{11}$$

Along this flow, let $\dot{\gamma}^\omega_t := \frac{\partial \gamma^\omega_t(\nu)}{\partial \nu}$,

$$\frac{\mathrm{d}}{\mathrm{d}\nu} \phi^\omega_t(\gamma^\omega_t(\nu)) = \|\dot{\gamma}^\omega_t(\nu)\|^2_{\tilde{g}^\omega_{t,\gamma^\omega_t(\nu)}} \geq 0, \tag{12}$$

For any measurable $\mathcal{A} \subseteq \mathcal{M}$, the kernel is

$$K^{\tilde{g}^w_{t,\tau}}_t(\mathcal{A} \mid \tau) = \Pr_\omega\Big[F^{\tilde{g}^w_{t,\tau}}_t(\tau, \omega) \in \mathcal{A}\Big]. \tag{13}$$

3.3 Horizon refinement through conditional responses

The base metric couples action perturbations through pose geometry and temporal variation. We refine it using cross-segment predictions learned from expert trajectories and their corrupted versions.

Learning cross-segment distributions. Let $\mathcal{I} = \{0, \ldots, H\}$, and let $\mathcal{B}$ be a finite set of nonempty proper blocks $B \subset \mathcal{I}$, with complement $\bar{B}$. Choose the blocks so that every action index occurs in both visible and held-out blocks. For a fixed reference measure μ_B on $\mathcal{P}^{|B|}$, define the conditional density

$$q^B_{\phi,t}(z_B \mid \tau) := q_\phi(z_B \mid \tau_{\bar{B}}, t, B), \qquad z_B \in \mathcal{P}^{|B|}. \tag{14}$$

The predictor never receives τ_B. Generate training pairs (τ_0, τ_t, t) using the base kernel, with $\tilde{g}^w_{t,\tau} = g^{\mathcal{M}}_\tau$, and denote this law by $\mathcal{D}_{\mathrm{ref}}$. Fit the predictor with

$$\mathcal{L}_{\mathrm{resp}}(\phi) = \mathbb{E}_{\mathcal{D}_{\mathrm{ref}}} \mathbb{E}_{B \sim \mathrm{Unif}(\mathcal{B})}[-\log q^B_{\phi,t}(\tau_{0,B} \mid \tau_t)]. \tag{15}$$

We then freeze ϕ. This predictor is fitted before the refinement changes the forward process. It need not equal the true conditional law of the refined chain.

Structured horizon probes. For smooth temporal basis functions b_m, construct

$$\widetilde{\xi}^{(j)}_h = \sum_{m=1}^{M_b} c_{j,m} b_m(h), \qquad c_{j,m} \in \mathfrak{se}(3), \qquad \xi^{(j)}_h = \mathrm{d}L_{x_h} \widetilde{\xi}^{(j)}_h. \tag{16}$$

Normalize each nonzero $\Xi^{(j)}_\tau = (\xi^{(j)}_h)_{h=0}^H$ under $g^{\mathcal{M}}_\tau$. Let $\mathrm{Exp}^{\mathcal{M}}$ be the Riemannian exponential of the full base metric, including its temporal term. For sufficiently small $\delta > 0$, set

$$\begin{aligned} \tau^{\pm,j} &= \mathrm{Exp}^{\mathcal{M}}_\tau(\pm\delta \Xi^{(j)}_\tau), \\ r^\delta_j(t, \tau) &= \frac{1}{|\mathcal{B}|\delta^2} \sum_{B \in \mathcal{B}} \sum_{\sigma \in \{-,+\}} D_{\mathrm{KL}}\big(q^B_{\phi,t}(\cdot \mid \tau) \,\big\|\, q^B_{\phi,t}(\cdot \mid \tau^{\sigma,j})\big). \end{aligned} \tag{17}$$

Only visible components enter each prediction. A large response means that the probe changes predictions of other segments. It measures predictive sensitivity, not distance from expert support or physical feasibility.

Local response geometry. Assume positive densities on a common support, twice differentiable dependence on τ, and the integrability conditions for differentiating under the integral. Define

$$\begin{aligned}\mathcal{H}_{t,\tau}(\Xi,\Psi) &= \frac{1}{|\mathcal{B}|}\sum_{B\in\mathcal{B}} \mathbb{E}_{z_B\sim q^B_{\phi,t}(\cdot|\tau)}[\mathrm{d}_\tau \ell^B_{t,z_B}[\Xi]\,\mathrm{d}_\tau \ell^B_{t,z_B}[\Psi]],\\ \ell^B_{t,z_B}(\rho) &= \log q^B_{\phi,t}(z_B \mid \rho).\end{aligned} \tag{18}$$

This is a positive-semidefinite average of pullback Fisher tensors (Arvanitidis et al., 2022). The two KL terms in Eq. 17 each contribute half the quadratic term, so

$$r^\delta_j(t,\tau) = \mathcal{H}_{t,\tau}(\Xi^{(j)}_\tau,\Xi^{(j)}_\tau) + o(1), \qquad \delta\to 0. \tag{19}$$

Joint probes can reveal cross-action terms. The tensor transforms intrinsically when the predictive laws, reference measures, and tangents are transformed consistently.

A smooth positive-definite refinement. To avoid evaluating predictive responses inside every kernel update, fit

$$R_\psi(t,\tau) = L_\psi(t,\tau)L_\psi(t,\tau)^\star, \qquad L_\psi(t,\tau):\mathbb{R}^r\to T_\tau\mathcal{M}, \tag{20}$$

where $\star$ uses $g^{\mathcal{M}}_\tau$ and the Euclidean inner product. Thus R_ψ is self-adjoint and positive semidefinite. Minimize

$$\begin{aligned}\mathcal{L}_{\mathrm{met}}(\psi) = \mathbb{E}_{\mathcal{D}_{\mathrm{ref}}}\Bigg[\frac{1}{J}\sum_{j=1}^{J}\Big(g^{\mathcal{M}}_{\tau_t}(R_\psi(t,\tau_t)\Xi^{(j)}_{\tau_t},\Xi^{(j)}_{\tau_t}) - r^\delta_j(t,\tau_t)\Big)^2\\ + \lambda_R\|R_\psi(t,\tau_t)\|^2_{\mathrm{HS},g^{\mathcal{M}}_{\tau_t}}\Bigg].\end{aligned} \tag{21}$$

Finite probes and a rank constraint do not identify the full tensor. The fit supervises the probed quadratic forms. For $\alpha\ge 0$, define

$$\begin{aligned}g^{\mathcal{M}}_{\psi,t,\tau}(\Xi,\Psi) &= g^{\mathcal{M}}_\tau\Big((I+\alpha R_\psi)^{1/2}\Xi,(I+\alpha R_\psi)^{1/2}\Psi\Big)\\ &= g^{\mathcal{M}}_\tau(\Xi,\Psi) + \alpha g^{\mathcal{M}}_\tau(R_\psi(t,\tau)\Xi,\Psi).\end{aligned} \tag{22}$$

This is positive definite and recovers the base metric at $\alpha=0$. It adds cost to directions with large fitted predictive response. Set $\tilde g^w_{t,\tau} = g^{\mathcal{M}}_{\psi,t,\tau}$ in the kernel. Self-adjointness, norms, energies, and gradients all use this same metric field. The parameters ϕ,ψ are frozen after fitting, while metric values still vary with (t,τ).

3.4 Forward transitions and reverse learning

Stochastic metric evolution. One finite expansion of the random deformation field is

$$S^\omega_t(\tau) = \sum_{m=1}^{M_s} z_m\rho_m(\tau)\mathsf{B}_m(t,\tau), \qquad z_m \overset{\mathrm{iid}}{\sim}\mathcal{N}(0,1), \tag{23}$$

where ρ_m are smooth scalar functions and B_m are smooth $g^{\mathcal{M}}_{\psi,t,\tau}$-self-adjoint operators. A single coefficient vector is held fixed during one internal flow. The field law is part of the kernel specification; smoothness alone does not specify a unique law.

Let $0=t_0<\cdots<t_K=1$ and select $s_k\in(t_k,t_{k+1}]$ as the evaluation time of each transition. This avoids forcing the first transition to be the identity when $\varepsilon(0)=0$. With independent field draws,

$$\begin{aligned}\tau_{t_{k+1}} &= F^\psi_{s_k}(\tau_{t_k},\omega_k), \qquad K^\psi_k = K^{g_\psi}_{s_k},\\ q_\psi(\mathrm{d}\tau_{t_0:t_K}\mid c) &= p_{\mathrm{data}}(\mathrm{d}\tau_{t_0}\mid c)\prod_{k=0}^{K-1}K^\psi_k(\mathrm{d}\tau_{t_{k+1}}\mid\tau_{t_k}).\end{aligned} \tag{24}$$

Here F^ψ_t denotes Eq. 11 with the refined metric. The step convention also defines the reference chain used for $\mathcal{D}_{\mathrm{ref}}$. That dataset records the noising label supplied to the predictor at

each sampled state. The horizon module supplies deterministic response geometry; the sampled operator field supplies transition randomness. For correlated metric draws, we use $z_{k+1,m} = \varrho z_{k,m} + \sqrt{1-\varrho^2}\,\eta_{k+1,m}$, $|\varrho| < 1$, $\eta_{k+1,m} \overset{\text{iid}}{\sim} \mathcal{N}(0,1)$. Starting with standard normal z_0 gives stationary coefficients. The Markov state is then (τ_{t_k}, z_k), and source and reverse models must account for this latent state. The remaining discussion uses independent draws.

Relation to correlated coordinate noise. Coordinate covariances can also encode horizon dependence and vary with state and time. Our distinction is the construction of the transition law: predictive responses define a metric, a random field deforms that metric, and the resulting energy and gradient induce motion.

Terminal source. Let $p_T^{\psi}(\cdot \mid c)$ be the terminal marginal of the refined chain. It is not automatically a known Gaussian. To use a sampleable reference $p_{\mathrm{ref}}(\cdot \mid c)$, calibrate the noise schedule or field-law parameters through $\mathcal{L}_{\mathrm{src}} = \mathbb{E}_c\Big[D\big(p_T^{\psi}(\cdot \mid c), p_{\mathrm{ref}}(\cdot \mid c)\big)\Big]$, where D can be a sample-based discrepancy such as MMD (Gretton et al., 2012). Alternatively, fit a conditional source sampler to terminal samples. Matching only a pooled marginal is insufficient for a shared unconditional source. Finalize any calibration before generating the reverse-training data; residual source mismatch remains an approximation error.

Reverse objectives. Train on pairs from the final frozen chain, for example using a stochastic reverse transition model $p_\theta^{\psi}(\mathrm{d}\tau_{t_k} \mid \tau_{t_{k+1}}, c)$. The measure notation also permits singular kernels. A reconstruction loss can train a denoiser, but does not by itself specify an exact reverse sampler. For flow regression, retain each internal path γ_k and set

$$\tau_s = \gamma_k\left(\frac{s-t_k}{\Delta t_k}\right), \qquad \dot\tau_s = \frac{1}{\Delta t_k}\dot\gamma_k\left(\frac{s-t_k}{\Delta t_k}\right), \qquad \Delta t_k = t_{k+1} - t_k. \tag{25}$$

Regressing these velocities yields the marginal forward field when the required conditional expectations and transport regularity hold. Generation integrates the learned field from $s = 1$ to $s = 0$. A score formulation needs the relevant densities and available or estimated scores. Stochastic interpolant objectives require compatible path and endpoint assumptions (Albergo et al., 2025). The predictor q_ϕ learns response geometry and is separate from the reverse model. A detailed analysis of the response stability under metric corruption is provided in Appendix D.3.

4 Experiments

We evaluate our method from three perspectives. First, we study the intrinsic behavior of the proposed geometry-first corruption process through synthetic experiments (Q1–Q2). Q1 investigates whether randomizing trajectory geometry preserves geometric structure better than coordinate-based corruption under matched noise strength. Q2 examines whether this structured forward process improves the learnability of the reverse generation process. Second, we evaluate the practical impact on generative imitation learning through robot manipulation and navigation benchmarks with different action dimensions (Q3–Q5). These experiments test whether the proposed corruption mechanism improves action generation beyond synthetic geometric advantages. Finally, we provide additional analyses in the terminal-source convergence diagnostic (Q6), which verifies the forward process's source-matching property. Furthermore, we also conduct several real-world manipulation experiments to validate our method.

4.1 Synthetic Geometry and Reverse Learning

Forward corruption (Q1). We compare corruption of branching trajectories on $\mathbb{R}^2$, S^2, and $SO(3)$. All methods are matched to the same intrinsic RMS displacement, $r_t = 0.30t$, using rescaling only for this controlled diagnostic. Figure 1 shows representative trajectories in all three spaces. Table 1 compares our kernel with correlated Gaussian noise. Random metric corruption has lower local-geometry distortion, roughness, and path-length distortion in all three spaces. The local-geometry gap is larger on S^2 and $SO(3)$ than on $\mathbb{R}^2$. Appendix E.1 contains the complete baseline table and results across corruption strengths.

Table 1: Forward corruption at $t = 1$, with intrinsic RMS displacement 0.30 for every trajectory. Mean±std over 120 corrupted trajectories. We show the best-performing baseline and our kernel; Table 7 gives all five methods.

Space	Method	Local geom. ↓	Roughness ↓	Path-length dist. ↓
$\mathbb{R}^2$	Correlated Gaussian	0.414 ± 0.097	0.0116 ± 0.0051	0.175 ± 0.105
	Random Metric (Ours)	$\mathbf{0.329 \pm 0.099}$	$\mathbf{0.0015 \pm 0.0010}$	$\mathbf{0.120 \pm 0.077}$
S^2	Correlated Gaussian	0.795 ± 0.324	0.1955 ± 0.5566	0.616 ± 0.385
	Random Metric (Ours)	$\mathbf{0.351 \pm 0.107}$	$\mathbf{0.0013 \pm 0.0010}$	$\mathbf{0.308 \pm 0.144}$
$SO(3)$	Correlated Gaussian	1.051 ± 0.477	0.4087 ± 0.8123	0.993 ± 0.538
	Random Metric (Ours)	$\mathbf{0.459 \pm 0.116}$	$\mathbf{0.0015 \pm 0.0013}$	$\mathbf{0.459 \pm 0.116}$

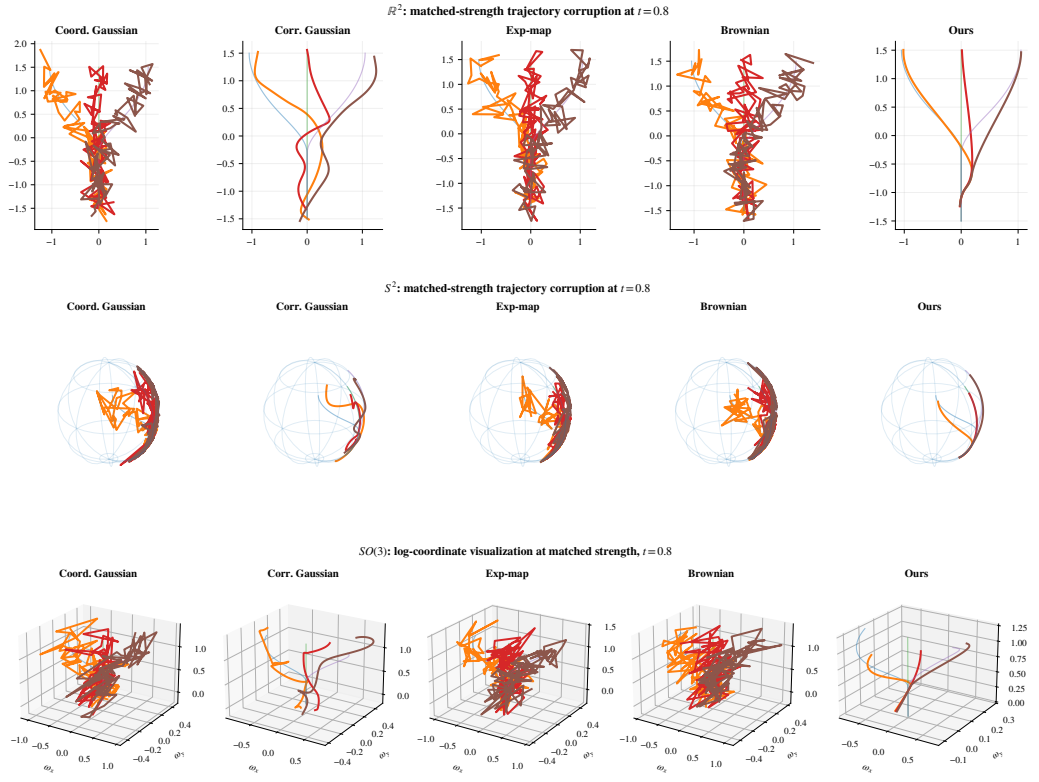

Figure 1: Matched-strength trajectory corruption at $t = 0.8$. Rows show $\mathbb{R}^2$, S^2, and $SO(3)$, respectively. Thin curves are clean trajectories and thick curves are corrupted ones. All methods have the same intrinsic RMS displacement. Rotations are displayed in logarithmic coordinates; perturbations and geometry metrics are evaluated intrinsically.

Learned reverse generation (Q2). We train the same time-conditioned residual denoiser on continuous three-mode trajectory distributions in $\mathbb{R}^2$ and $SO(3)$, using five forward steps with matched one-step RMS displacement. Recovery starts from held-out terminal corruptions. Generation starts from a fitted terminal source, using the same PCA–Gaussian-mixture family for every kernel (Appendix E.2). Our kernel gives the lowest recovery RMS error and local-geometry distortion in both spaces (Table 2). It also yields the lowest SWD and MMD on $\mathbb{R}^2$. On $SO(3)$, IID Gaussian corruption has lower distributional distances, while our kernel gives the lowest generated geometry distortion. Thus, the consistent benefit in Q2 is reverse recovery and geometric fidelity.

4.2 Robot Imitation

Table 2: Learned reverse recovery and source-fitted generation on $\mathbb{R}^2$ and $SO(3)$. All kernels use the same reverse-model architecture and terminal-prior family. Lower is better for every metric.

Space	Forward kernel	Recovery RMS	Recovery geom.	SWD	MMD	Gen. geom.
$\mathbb{R}^2$	IID Gaussian	0.084	0.140	1.917	0.1818	0.093
	Correlated Gaussian	0.086	0.080	3.885	0.2298	0.092
	Random Metric (Ours)	**0.072**	**0.063**	**1.466**	**0.1475**	**0.087**
$SO(3)$	IID Gaussian	0.083	0.389	0.380	**0.0265**	0.153
	Correlated Gaussian	0.083	0.254	0.375	0.0308	0.108
	Random Metric (Ours)	**0.063**	**0.082**	**0.370**	0.0374	**0.087**

Table 3: Controlled Push-T comparison for representative generative policies. Only the forward corruption mechanism is changed. Results are mean±std over multiple random seeds.

Forward corruption	Avg. score ↑	Max. score ↑	Jerk ↓	Chunk inc. ↓
IID Gaussian	90.7 ± 1.2	92.8 ± 1.0	1.000 ± 0.043	1.000 ± 0.039
Correlated Gaussian / GP	91.9 ± 1.1	94.1 ± 0.8	0.816 ± 0.035	0.842 ± 0.033
Smooth-field coordinate noise	93.0 ± 0.9	95.3 ± 0.7	0.651 ± 0.028	0.713 ± 0.027
Random Metric (t only)	94.6 ± 0.8	97.0 ± 0.5	0.684 ± 0.030	0.637 ± 0.025
Full ($h + t$)	$\mathbf{96.2 \pm 0.6}$	$\mathbf{98.6 \pm 0.3}$	$\mathbf{0.512 \pm 0.024}$	$\mathbf{0.486 \pm 0.021}$

Push-T (Q3). Full achieves the highest average and maximum task scores, 96.2 and 98.6, and the lowest jerk and chunk inconsistency (Table 3). Smooth-field coordinate noise has slightly lower jerk than t only (0.651 versus 0.684), but lower task scores and greater chunk inconsistency. This comparison separates action smoothness from the observed control benefit. Also, we compare GFC against Diffusion Policy (DP) (Chi et al., 2023), Flow Matching Policy (FMP) (Zhang & Gienger, 2025), and Streaming Flow Policy (SFP) (Jiang et al., 2025). All baselines are implemented using their official public codebases. Table 4 gives the state-input Push-T results, together with our scores. Our implementation requires 30.1 ms per action-generation call. Additional detail are provided in Appendix E.4.

Table 4: Push-T task scores.

Policy	Avg. score ↑	Max. score ↑
Diffusion Policy (100 DDPM)	90.7	92.8
Diffusion Policy (10 DDIM)	81.4	85.3
Flow Matching Policy	80.6	82.6
Streaming Diffusion Policy	84.2	87.0
Ours	**96.2**	**98.6**

RoboMimic, Adroit (Q4) and Long-Horizon Navigation (Q5). We choose three tasks from RoboMimic (Lift, Can, and Square) (Mandlekar et al., 2021) and four tasks from Adroit (Door, Pen, Hammer, and Relocate) (Rajeswaran et al., 2018). Full has the highest mean success on five of seven tasks. The smooth-field coordinate noise performs best on Pen (Table 5). Full also gives the lowest benchmark-averaged jerk and chunk inconsistency on both RoboMimic and Adroit (Table 9 in appendix). Further, we evaluate the same five corruption mechanisms on Maze2D U-Maze, Medium, and Large. We report goal success, path inefficiency on successful trajectories, invalid-trajectory rate, and action jerk. Full has the highest Medium success rate, the lowest path ratio, and the lowest invalid rate (Table 6). The t-only variant has the highest Large success rate, while smooth-field coordinate noise has the highest U-Maze success. Metric definitions are given in Appendix E.5.

4.3 Terminal-Source Compatibility (Q6)

Preserving intermediate trajectory structure must remain compatible with a sampleable terminal source. We compare each forward marginal p_t with a common reference $p_{\text{ref}} = \mathcal{N}(0, I)$ using RBF-MMD, $D_{\text{src}}(t) = \text{MMD}_k(p_t, p_{\text{ref}})$ (Gretton et al., 2012). Unlike Q2, the reference is fixed across methods. Figure 2 shows that all methods approach the reference. Random Metric corruption reaches a comparable terminal discrepancy while remaining closer at most intermediate times. Together with Q1, this indicates that its geometry preservation is not explained by weaker corruption.

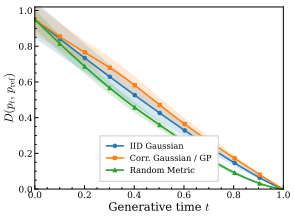

Figure 2: RBF-MMD.

Table 5: Success rates (%) on RoboMimic and Adroit. Results are mean±std.

	RoboMimic			Adroit			
Forward corruption	Lift ↑	Can ↑	Square ↑	Door ↑	Pen ↑	Hammer ↑	Relocate ↑
IID Gaussian	96.0 ± 1.2	89.7 ± 2.1	71.3 ± 3.2	78.7 ± 3.0	72.0 ± 3.5	65.3 ± 3.8	58.7 ± 4.0
Correlated Gaussian / GP	98.1 ± 0.9	91.5 ± 1.9	76.2 ± 2.9	83.0 ± 2.6	75.1 ± 3.0	70.8 ± 3.4	64.2 ± 3.6
Smooth-field coordinate noise	97.8 ± 0.9	93.6 ± 1.6	78.4 ± 2.7	81.9 ± 2.8	$\mathbf{84.1 \pm 2.3}$	72.4 ± 3.1	68.7 ± 3.3
Random Metric (t only)	98.5 ± 0.7	94.8 ± 1.4	85.1 ± 2.1	88.7 ± 2.0	81.9 ± 2.5	80.1 ± 2.7	76.8 ± 2.8
Full ($h + t$)	$\mathbf{98.8 \pm 0.6}$	$\mathbf{96.2 \pm 1.1}$	$\mathbf{86.4 \pm 2.2}$	$\mathbf{91.0 \pm 1.7}$	83.5 ± 2.2	$\mathbf{82.8 \pm 2.3}$	$\mathbf{77.9 \pm 2.7}$

Table 6: Maze2D success and trajectory quality across three layouts. Results are mean±std.

Forward corruption	U-Maze ↑	Medium ↑	Large ↑	Path ratio ↓	Invalid rate ↓	Jerk ↓
IID Gaussian	84.3 ± 2.8	67.1 ± 3.5	44.8 ± 4.0	1.427 ± 0.052	21.6 ± 2.4	1.000 ± 0.046
Correlated Gaussian / GP	92.0 ± 2.1	76.8 ± 3.1	62.7 ± 3.5	1.251 ± 0.043	12.7 ± 1.9	0.758 ± 0.036
Smooth-field coordinate noise	$\mathbf{96.2 \pm 1.1}$	82.5 ± 2.6	69.4 ± 3.4	1.184 ± 0.034	8.4 ± 1.3	0.438 ± 0.023
Random Metric (t only)	94.8 ± 1.5	89.1 ± 2.0	$\mathbf{85.6 \pm 2.1}$	1.112 ± 0.026	3.4 ± 0.8	0.611 ± 0.029
Full ($h + t$)	95.8 ± 1.2	$\mathbf{92.9 \pm 1.6}$	83.8 ± 2.3	$\mathbf{1.081 \pm 0.020}$	$\mathbf{2.7 \pm 0.7}$	$\mathbf{0.417 \pm 0.024}$

4.4 Real-world Experiments

We evaluate GFC and Diffusion Policy (Chi et al., 2023) on a Kinova Gen3 arm equipped with an Intel RealSense D435 camera on two tabletop tasks (Figure 3). **T-block pushing** requires moving a T-shaped block to a marked target while avoiding obstacles, with varying initial block and obstacle positions. Its difficulty increases across three configurations: an unobstructed route, a bowl obstacle, and a bowl plus an additional object. **Box uprighting** requires turning a flat cardboard box upright through out-of-plane object rotation. Both tasks require more than locally smooth actions: successive action chunks must maintain coherent trajectory structure across repeated replanning, supporting a consistent route around obstacles or coordinated motion throughout object reorientation. For each task, we collect 200 demonstrations and run 100 trials per policy. GFC succeeds in 20/20, 35/40, and 32/40 Push-T trials across the three configurations (87/100 overall); Diffusion Policy succeeds in 2/20 trials in the easiest setting and none in the other two. On box uprighting, GFC succeeds in 75/100 trials, compared with 0/100 for Diffusion Policy. Additional details are provided in Appendix E.6



Figure 3: Real-robot task setups.T-block pushing (left) and box uprighting (right).

5 Conclusion, Limitation and Broader Impact

We presented GFC, a geometry-first forward process for generative imitation learning. It learns a trajectory metric from cross-segment predictive responses and corrupts trajectories through random metric deformations, without an additional coordinate-noise draw. We derived pathwise bounds on trajectory motion and predictive response, and stated the conditions needed for reverse generation. Experiments across synthetic geometries, manipulation, navigation, and real robots show improved trajectory fidelity and policy performance. The two-axis ablation further highlights the complementary effects of response-based refinement and stochastic metric corruption.

Limitation. GFC requires predictor and response operator. This staged pipeline adds training complexity and may be sensitive to errors in the learned response geometry. Simplifying the pipeline and improving source matching are directions for future work.

Broader Impact. On the robotics side, improved trajectory fidelity and lower jerk/chunk inconsistency are practically meaningful for control stability and hardware safety. Furthermore, the concept that "learned geometric structures should shape the degradation process" holds value far beyond the scope of robotics; fields such as manifold-valued generative modeling, analysis of strongly coupled time series, and structured motion priors all stand to benefit from it. We will continue to study on this.

REFERENCES

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26:1–80, 2025.

Alexis Arnaudon, Darryl D. Holm, and Stefan Sommer. A geometric framework for stochastic shape analysis. *Foundations of Computational Mathematics*, 19:653–701, 2019. doi: 10.1007/s10208-018-9394-z. URL https://link.springer.com/article/10.1007/s10208-018-9394-z.

Marc Arnaudon, Kolehe Abdoulaye Coulibaly, and Anton Thalmaier. Brownian motion with respect to a metric depending on time; definition, existence and applications to Ricci flow. *Comptes Rendus Mathématique*, 346:773–778, 2008. doi: 10.1016/j.crma.2008.05.004. URL https://www.numdam.org/item/10.1016/j.crma.2008.05.004.pdf.

Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: On the curvature of deep generative models. In *International Conference on Learning Representations*, 2018. URL https://arxiv.org/abs/1710.11379.

Georgios Arvanitidis, Miguel González-Duque, Alison Pouplin, Dimitris Kalatzis, and Søren Hauberg. Pulling back information geometry, 2022. URL https://arxiv.org/abs/2106.05367.

Simone Azeglio and Arianna Di Bernardo. What's inside your diffusion model? a score-based Riemannian metric to explore the data manifold. *arXiv preprint arXiv:2505.11128*, 2025. URL https://arxiv.org/abs/2505.11128.

Jacob Bamberger, Adam Gosztolai, Pierre Vandergheynst, Michael Bronstein, and Iolo Jones. Riemannian metric matching for scalable geometric modeling of distributions. *arXiv preprint arXiv:2606.14334*, 2026a.

Jacob Bamberger, Iolo Jones, Dennis Duncan, Michael Bronstein, Pierre Vandergheynst, and Adam Gosztolai. Carré du champ flow matching: Better quality-generalisation tradeoff in generative models. In *International Conference on Learning Representations*, 2026b.

Grigory Bartosh, Dmitry Vetrov, and Christian A. Naesseth. Neural flow diffusion models: Learnable forward process for improved diffusion modelling. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2352. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/871a8ccb9232487366feb5e2d9069915-Abstract-Conference.html.

Martin Bauer, Cy Maor, and Peter W. Michor. Sobolev metrics on spaces of manifold-valued curves. *Annali della Scuola Normale Superiore di Pisa, Classe di Scienze*, 24(4):1895–1948, 2023. doi: 10.2422/2036-2145.202010_016.

Hadi Beik-Mohammadi, Søren Hauberg, Georgios Arvanitidis, Gerhard Neumann, and Leonel Rozo. Learning riemannian manifolds for geodesic motion skills. In *Robotics: Science and Systems*, 2021. doi: 10.15607/RSS.2021.XVII.082.

Hadi Beik-Mohammadi, Søren Hauberg, Georgios Arvanitidis, Gerhard Neumann, and Leonel Rozo. Reactive motion generation on learned riemannian manifolds. *The International Journal of Robotics Research*, 42(10):729–754, 2023. doi: 10.1177/02783649231193046.

Marin Biloš, Kashif Rasul, Anderson Schneider, Yuriy Nevmyvaka, and Stephan Günnemann. Modeling temporal data as continuous functions with stochastic process diffusion. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 2452–2470. PMLR, 2023. URL https://proceedings.mlr.press/v202/bilos23a.html.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

Kevin Black, Manuel Y. Galliker, and Sergey Levine. Real-time execution of action chunking flow policies, 2025. URL https://arxiv.org/abs/2506.07339.

Max Braun, Noémie Jaquier, Leonel Rozo, and Tamim Asfour. Riemannian flow matching policy for robot motion learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5144–5151. IEEE, 2024. doi: 10.1109/IROS58592.2024.10801521.

Andrew Bylard, Riccardo Bonalli, and Marco Pavone. Composable geometric motion policies using multi-task pullback bundle dynamical systems. In *IEEE International Conference on Robotics and Automation*. IEEE, 2021. doi: 10.1109/ICRA48506.2021.9561320.

Pascal Chang, Jingwei Tang, Markus Gross, and Vinicius C. Azevedo. How i warped your noise: A temporally-correlated noise prior for diffusion models. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=pzElnMrgSD.

Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations*, 2024.

Ching-An Cheng, Mustafa Mukadam, Jan Issac, Stan Birchfield, Dieter Fox, Byron Boots, and Nathan D. Ratliff. Rmpflow: A geometric framework for generation of multitask motion policies. *IEEE Transactions on Automation Science and Engineering*, 18(3):968–987, 2021. doi: 10.1109/TASE.2021.3053422.

Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin C. M. Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Robotics: Science and Systems*, 2023. doi: 10.15607/RSS.2023.XIX.026.

Brian Clarke, Dmitry Jakobson, Niky Kamran, Lior Silberman, and Jonathan Taylor. Gaussian measures on the of space of Riemannian metrics. *Annales Mathématiques du Québec*, 39: 129–145, 2015. doi: 10.1007/s40316-015-0037-3. URL https://link.springer.com/article/10.1007/s40316-015-0037-3.

Giannis Daras, Mauricio Delbracio, Hossein Talebi, Alexandros G. Dimakis, and Peyman Milanfar. Soft diffusion: Score matching with general corruptions. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=W98rebBxlQ.

Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. *Advances in neural information processing systems*, 35:2406–2422, 2022a.

Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian score-based generative modelling. In *Advances in Neural Information Processing Systems*, volume 35, 2022b.

Friso de Kruiff, Erik Bekkers, Ozan Öktem, Carola-Bibiane Schönlieb, and Willem Diepeveen. Pullback flow matching on data manifolds. *arXiv preprint arXiv:2410.04543*, 2024. URL https://arxiv.org/abs/2410.04543.

Haoran Ding, Noémie Jaquier, Jan Peters, and Leonel Rozo. Fast and robust visuomotor riemannian flow matching policy. *IEEE Transactions on Robotics*, 41:5327–5343, 2025. doi: 10.1109/TRO.2025.3601293.

Luca Falorsi, Pim De Haan, Tim R Davidson, and Patrick Forré. Reparameterizing distributions on lie groups. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 3244–3253. PMLR, 2019.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.

Chin-Wei Huang, Milad Aghajohari, Avishek Joey Bose, Prakash Panangaden, and Aaron Courville. Riemannian diffusion models. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Sunshine Jiang, Xiaolin Fang, Nicholas Roy, Tomás Lozano-Pérez, Leslie Pack Kaelbling, and Siddharth Ancha. Streaming flow policy: Simplifying diffusion/flow-matching policies by treating action trajectories as flow trajectories. In *Proceedings of the 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 238–257. PMLR, 2025.

Kacper Kapuśniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Avishek Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-4291.

Gavin Kerrigan, Giosue Migliorini, and Padhraic Smyth. Functional flow matching. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 3934–3942. PMLR, 2024. URL https://proceedings.mlr.press/v238/kerrigan24a.html.

Marcel Kollovieh, Marten Lienen, David Lüdke, Leo Schwinn, and Stephan Günnemann. Flow matching with Gaussian process priors for probabilistic time series forecasting. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=uxVBbSlKQ4.

Tom LaGatta and Jan Wehr. Geodesics of random riemannian metrics. *Communications in Mathematical Physics*, 327(1):181–241, 2014. doi: 10.1007/s00220-014-1901-8.

Yiming Li, Nael Darwiche, Amirreza Razmjoo, Sichao Liu, Yilun Du, Auke Ijspeert, and Sylvain Calinon. Geometry-aware policy imitation. In *International Conference on Learning Representations*, 2026.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Gabriel Loaiza-Ganem, Brendan Leigh Ross, Rasa Hosseinzadeh, Anthony L. Caterini, and Jesse C. Cresswell. Deep generative models through the lens of the manifold hypothesis: A survey and new connections. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=a90WpmSi0I.

Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation, 2021. URL https://arxiv.org/abs/2108.03298.

Lekan Molu. Kolmogorov regression for robust diffusion policies. *arXiv preprint arXiv:2606.18186*, 2026. URL https://arxiv.org/abs/2606.18186.

Xavier Pennec. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006. doi: 10.1007/s10851-006-6228-4.

Aaditya Prasad, Kevin Lin, Jimmy Wu, Linqi Zhou, and Jeannette Bohg. Consistency policy: Accelerated visuomotor policies via consistency distillation. In *Robotics: Science and Systems*, 2024. doi: 10.15607/RSS.2024.XX.071.

Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations, 2018. URL https://arxiv.org/abs/1709.10087.

Nathan D. Ratliff, Jan Issac, Daniel Kappler, Stan Birchfield, and Dieter Fox. Riemannian motion policies. *arXiv preprint arXiv:1801.02854*, 2018.

Leonel Rozo, Miguel González-Duque, Noémie Jaquier, and Søren Hauberg. Riemann^2: Learning riemannian submanifolds from riemannian data. In *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 3097–3105. PMLR, 2025.

Shinnosuke Saito and Takashi Matsubara. Be tangential to manifold: Discovering Riemannian metric for diffusion models. In *European Conference on Computer Vision*, 2026. URL https://arxiv.org/abs/2510.05509.

Joan Solà, Jérémie Deray, and Dinesh Atchuthan. A micro lie theory for state estimation in robotics. *arXiv preprint arXiv:1812.01537*, 2018.

Vikram Voleti, Christopher Pal, and Adam Oberman. Score-based denoising diffusion with non-isotropic Gaussian noise models. In *NeurIPS Workshop on Score-Based Methods*, 2022. URL https://arxiv.org/abs/2210.12254.

Fan Zhang and Michael Gienger. Affordance-based robot manipulation with flow matching, 2025. URL https://arxiv.org/abs/2409.01083.

Qinglun Zhang, Zhen Liu, Haoqiang Fan, Guanghui Liu, Bing Zeng, and Shuaicheng Liu. Flowpolicy: Enabling fast and robust 3d flow-based policy via consistency flow matching for robot manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 14754–14762, 2025. doi: 10.1609/aaai.v39i14.33617.

Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware, 2023. URL https://arxiv.org/abs/2304.13705.

A USE OF AI ASSISTANCE

Generative AI tools were used to assist with manuscript editing and polishing, the retrieval and discovery of related literature, and coding support. All AI-assisted code was reviewed, edited, and validated by the authors. Generative AI tools were not used to develop the core methodology or algorithms, formulate theoretical results, conduct experiments, or generate experimental results. The authors reviewed all AI-assisted content and take full responsibility for the final content of this work.

B NOTATION

Action trajectories and two time axes. Let $\mathcal{P}$ be a d-dimensional action manifold and $\mathcal{M} = \mathcal{P}^{H+1}$ the ambient trajectory manifold. A trajectory contains $H+1$ actions indexed from 0 to H:

$$A_{t,h} \in \mathcal{P}, \qquad \tau_t = (A_{t,0}, \dots, A_{t,H}) \in \mathcal{M}. \tag{26}$$

Here h indexes the physical action horizon, while $t \in [0,1]$ indexes forward noising time. We denote the forward marginal by $p_t(\cdot \mid c)$, with $p_0(\cdot \mid c) = p_{\mathrm{data}}(\cdot \mid c)$, where c is an optional observation or language condition. Reverse generation starts from a source matched to the terminal law $p_1(\cdot \mid c)$. At fixed t, we write $x_h = A_{t,h}$ and $\tau = (x_0, \dots, x_H)$. The manifold $\mathcal{M}$ specifies the action representation; it is not an estimated low-dimensional support of demonstrations.

Tangent vectors and trajectory metrics. A trajectory tangent is $\Xi = (\xi_0, \dots, \xi_H) \in T_\tau\mathcal{M}$, where $\xi_h \in T_{x_h}\mathcal{P}$. We write $\Psi = (\psi_0, \dots, \psi_H)$ for a second tangent. The base trajectory metric is $g^{\mathcal{M}}_\tau$, and its learned refinement is $g^{\mathcal{M}}_{\psi,t,\tau}$. The parameter symbol ψ is distinct from the tangent components ψ_h. The notation $\widetilde{g}^{\omega}_{t,\tau}$ is reserved for its random deformation, where ω indexes a random-field realization. Thus, $\bar{g}_{t,\tau} = g^{\mathcal{M}}_{\psi,t,\tau}$ in the full method.

Metric-dependent operators. For a metric g, the tangent norm is $\|\Xi\|_g = \sqrt{g(\Xi,\Xi)}$. The metric gradient of a scalar function f is defined by

$$g_\tau\big(\mathrm{grad}_g f(\tau), \Xi\big) = \mathrm{d}f_\tau[\Xi], \qquad \Xi \in T_\tau\mathcal{M}. \tag{27}$$

We use $\|\cdot\|_{\mathrm{HS},g}$ and $\|\cdot\|_{\mathrm{op},g}$ for the Hilbert–Schmidt and operator norms induced by g. The superscript $\star$ denotes the adjoint under the specified inner products. In $R_\psi = L_\psi L_\psi^\star$, the trajectory-space inner product is the base metric $g^{\mathcal{M}}$; self-adjointness of S_t^ω is instead defined under the kernel metric $\bar{g}_t$. We distinguish the manifold exponential $\mathrm{Exp}^{\mathcal{M}}_\tau$ from the operator exponential e^S. Unless otherwise specified, $\mathrm{Exp}^{\mathcal{M}}$ uses the base metric.

Cross-segment response geometry. A block $B \in \mathcal{B}$ is a subset of action indices, and $\bar{B}$ denotes its complement. The predictor $q^B_{\phi,t}$ models the held-out block using the visible trajectory segments at noise level t. Its parameters ϕ are frozen before fitting the response operator R_ψ. We use $\Xi^{(j)}_\tau$, $j = 1, \dots, J$, for structured tangent probes and $\delta > 0$ for their perturbation scale. The weight $\alpha \geq 0$ controls the contribution of the response operator to the refined metric.

Forward steps and internal flow. The discrete forward chain uses $0 = t_0 < \cdots < t_K = 1$. For transition k, the kernel is evaluated at $s_k = t_{k+1}$ and uses a fresh realization ω_k. The internal path $\gamma_k(\nu)$ satisfies $\gamma_k(0) = \tau_{t_k}$ and $\gamma_k(1) = \tau_{t_{k+1}}$, where $\nu \in [0,1]$ parameterizes the flow within one transition. Both s_k and ω_k remain fixed during this flow, while the metric and operator field are evaluated at the evolving state $\gamma_k(\nu)$. The variable ν is distinct from both the action index h and the forward time t.

C PRIMER ON RIEMANNIAN GEOMETRY

This appendix summarizes the geometric concepts used in the main text. It is intended as a self-contained reference rather than a complete introduction to differential geometry. We review smooth manifolds, tangent spaces, Riemannian metrics, geodesics, exponential and logarithmic maps, the space of Riemannian metrics, and product geometry for trajectories. We follow standard conventions in Riemannian geometry; see Pennec (2006).

C.1 SMOOTH MANIFOLDS

A d-dimensional smooth manifold $\mathcal{N}$ is a space that is locally equivalent to $\mathbb{R}^d$. For each point $x \in \mathcal{N}$, there is an open neighborhood $U \subset \mathcal{N}$ and a coordinate chart

$$\varphi : U \to \varphi(U) \subset \mathbb{R}^d. \tag{28}$$

An atlas is a collection of such charts covering $\mathcal{N}$. The manifold is smooth when the transition maps between overlapping charts are smooth. A chart assigns coordinates to points locally, but the manifold and its geometric objects are independent of the chosen coordinates.

Euclidean space $\mathbb{R}^d$ is a smooth manifold. Other common examples are the sphere $\mathbb{S}^d$, the rotation group $\mathrm{SO}(d)$, and the rigid-transformation group $\mathrm{SE}(d)$. These spaces can be represented in ambient Euclidean coordinates, but their intrinsic dimensions and admissible motions are determined by the manifolds themselves.

A smooth manifold specifies the underlying space but does not, by itself, define lengths, angles, or distances. These quantities require an additional geometric structure: a Riemannian metric.

C.2 TANGENT SPACES AND TANGENT VECTORS

Although a manifold may be globally nonlinear, its infinitesimal structure at a point is linear. The tangent space at $x \in \mathcal{N}$, denoted by $T_x\mathcal{N}$, is a d-dimensional vector space containing the possible infinitesimal directions at x. A tangent vector $v \in T_x\mathcal{N}$ can be represented as the velocity of a smooth curve $\gamma : (-\epsilon, \epsilon) \to \mathcal{N}$ with $\gamma(0) = x$:

$$v = \dot{\gamma}(0). \tag{29}$$

Curves that have the same velocity in one, and hence every, local chart represent the same tangent vector.

The tangent bundle is the disjoint union of all tangent spaces,

$$T\mathcal{N} = \bigsqcup_{x \in \mathcal{N}} T_x\mathcal{N}. \tag{30}$$

It is itself a smooth manifold of dimension $2d$. A smooth vector field assigns a tangent vector $V(x) \in T_x\mathcal{N}$ to every point x.

For a smooth map $F : \mathcal{N} \to \mathcal{Q}$, its differential at x is the linear map

$$\mathrm{d}F_x : T_x\mathcal{N} \to T_{F(x)}\mathcal{Q}. \tag{31}$$

It describes the first-order response of F to a tangent perturbation. In Euclidean space, $\mathrm{d}F_x$ is represented by the usual Jacobian.

All tangent spaces of $\mathbb{R}^d$ can be canonically identified with $\mathbb{R}^d$. On a general manifold, however, $T_x\mathcal{N}$ and $T_y\mathcal{N}$ are different vector spaces when $x \neq y$. Their vectors cannot be subtracted or compared directly without an identification such as a coordinate trivialization or parallel transport.

C.3 RIEMANNIAN METRICS

A Riemannian metric g is a smoothly varying inner product on the tangent spaces of $\mathcal{N}$. At each point $x \in \mathcal{N}$, it assigns a symmetric, positive-definite bilinear form

$$g_x : T_x\mathcal{N} \times T_x\mathcal{N} \to \mathbb{R}. \tag{32}$$

For $u, v \in T_x\mathcal{N}$, the induced norm is

$$\|v\|_{g_x} = \sqrt{g_x(v, v)}, \tag{33}$$

and the angle between two nonzero vectors satisfies

$$\cos\theta = \frac{g_x(u, v)}{\|u\|_{g_x}\|v\|_{g_x}}. \tag{34}$$

The pair $(\mathcal{N}, g)$ is called a Riemannian manifold.

The metric determines the local cost of motion. If two directions have different metric norms, equal coordinate displacements along those directions have different geometric magnitudes. Two metrics can therefore define different geometries on the same underlying manifold.

C.4 Local Matrix Representation of a Metric

Let $(x^1, \dots, x^d)$ be local coordinates and let $\{\partial_1, \dots, \partial_d\}$ be the associated coordinate basis of $T_x\mathcal{N}$. The metric coefficients are

$$G_{ij}(x) = g_x(\partial_i, \partial_j). \tag{35}$$

They form a symmetric positive-definite matrix

$$G(x) \in \mathbb{S}^d_{++}. \tag{36}$$

If u and v denote the coordinate vectors of two tangent vectors, then

$$g_x(u, v) = u^\top G(x) v, \qquad \|v\|^2_{g_x} = v^\top G(x) v. \tag{37}$$

The Euclidean metric is the special case $G(x) = I$. A spatially varying matrix field $G(x)$ defines a nonuniform geometry in which the same coordinate increment may have different magnitudes at different points.

The matrix $G(x)$ is a coordinate representation, not the metric itself. Under a change of coordinates, both tangent-vector coordinates and $G(x)$ transform so that the scalar $g_x(u, v)$ remains unchanged. This invariance distinguishes the intrinsic metric from any particular matrix representation.

C.5 Curves, Length, and Geodesic Distance

A smooth curve is a map $\gamma : [0, 1] \to \mathcal{N}$ with velocity $\dot\gamma(s) \in T_{\gamma(s)}\mathcal{N}$. Its Riemannian length is

$$L_g(\gamma) = \int_0^1 \|\dot\gamma(s)\|_{g_{\gamma(s)}} \, \mathrm{d}s. \tag{38}$$

In local coordinates,

$$L_g(\gamma) = \int_0^1 \sqrt{\dot\gamma(s)^\top G(\gamma(s)) \dot\gamma(s)} \, \mathrm{d}s. \tag{39}$$

The Riemannian distance between $x, y \in \mathcal{N}$ is the infimum of the lengths of all piecewise smooth curves joining them:

$$d_g(x, y) = \inf_{\substack{\gamma(0)=x \\ \gamma(1)=y}} L_g(\gamma). \tag{40}$$

The metric also determines a unique torsion-free, metric-compatible connection, called the Levi–Civita connection and denoted by ∇. A geodesic satisfies

$$\nabla_{\dot\gamma} \dot\gamma = 0. \tag{41}$$

It has zero covariant acceleration. Geodesics are locally length minimizing, although a geodesic need not be globally shortest after it passes a cut point. Changing g can change both the length of a fixed curve and the geodesics between fixed endpoints.

C.6 Exponential and Logarithmic Maps

Let γ_v be the geodesic satisfying $\gamma_v(0) = x$ and $\dot\gamma_v(0) = v$. The exponential map at x is defined by

$$\mathrm{Exp}_x(v) = \gamma_v(1) \tag{42}$$

for every v for which γ_v exists on $[0, 1]$. Thus,

$$\mathrm{Exp}_x : \mathcal{D}_x \subseteq T_x\mathcal{N} \to \mathcal{N}. \tag{43}$$

If $(\mathcal{N}, g)$ is geodesically complete, then $\mathcal{D}_x = T_x\mathcal{N}$ for every x.

Within a normal neighborhood of x, Exp_x is invertible. Its local inverse is the logarithmic map

$$\mathrm{Log}_x(y) = \mathrm{Exp}_x^{-1}(y) \in T_x\mathcal{N}. \tag{44}$$

When the minimizing geodesic from x to y is unique,

$$d_g(x, y) = \|\mathrm{Log}_x(y)\|_{g_x}. \tag{45}$$

Outside such a neighborhood, the logarithmic map may be multivalued or undefined because several minimizing geodesics may connect the same points.

For the standard Euclidean metric,

$$\mathrm{Exp}_x(v) = x + v, \qquad \mathrm{Log}_x(y) = y - x. \tag{46}$$

C.7 Riemannian Gradients and Parallel Transport

For a smooth scalar function $f : \mathcal{N} \to \mathbb{R}$, the Riemannian gradient $\operatorname{grad}_g f(x) \in T_x\mathcal{N}$ is defined by

$$g_x\big(\operatorname{grad}_g f(x), v\big) = \mathrm{d}f_x[v], \qquad \forall v \in T_x\mathcal{N}. \tag{47}$$

In local coordinates, if $\nabla f(x)$ is the ordinary coordinate gradient, then

$$\operatorname{grad}_g f(x) = G(x)^{-1}\nabla f(x). \tag{48}$$

The metric therefore converts the differential, which is a covector, into a tangent vector and determines the direction of steepest ascent.

Parallel transport provides an intrinsic way to compare tangent vectors at different points along a curve. A vector field $V(s) \in T_{\gamma(s)}\mathcal{N}$ is parallel along γ when

$$\nabla_{\dot\gamma(s)} V(s) = 0. \tag{49}$$

Parallel transport preserves inner products and norms under g. Its result can depend on the chosen path when the manifold is curved.

C.8 The Space of Riemannian Metrics

For a fixed smooth manifold $\mathcal{N}$, the collection of smooth Riemannian metrics is denoted by

$$\operatorname{Met}(\mathcal{N}) = \left\{ g \in \Gamma\big(S^2T^*\mathcal{N}\big) : g_x \succ 0 \text{ for every } x \in \mathcal{N} \right\}. \tag{50}$$

Here, $T^*\mathcal{N}$ is the cotangent bundle, $S^2T^*\mathcal{N}$ is the bundle of symmetric covariant 2-tensors, and $\Gamma(S^2T^*\mathcal{N})$ denotes its smooth sections. Equivalently, each $g \in \operatorname{Met}(\mathcal{N})$ assigns a smoothly varying positive-definite bilinear form to the tangent spaces of $\mathcal{N}$.

In a coordinate chart, a metric may be represented as a smooth field

$$x \longmapsto G(x) \in \mathbb{S}^d_{++}. \tag{51}$$

Accordingly, $\operatorname{Met}(\mathcal{N})$ can be viewed as a cone inside the vector space of smooth symmetric 2-tensor fields. The notation separates the choice of a metric from the underlying manifold: g and $\widetilde{g}$ may belong to $\operatorname{Met}(\mathcal{N})$ while assigning different lengths, angles, and geodesics to the same points.

C.9 Product Manifolds and Trajectory Geometry

Let $\mathcal{P}$ be the manifold of an individual state or action. A discrete trajectory with horizon H is

$$\tau = (x_0, \dots, x_H), \qquad x_h \in \mathcal{P}. \tag{52}$$

It is a point on the product manifold

$$\mathcal{M} = \mathcal{P}^{H+1} = \underbrace{\mathcal{P} \times \cdots \times \mathcal{P}}_{H+1 \text{ factors}}. \tag{53}$$

Its tangent space is the direct sum

$$T_\tau\mathcal{M} \cong T_{x_0}\mathcal{P} \oplus \cdots \oplus T_{x_H}\mathcal{P}. \tag{54}$$

Thus, a tangent perturbation of the full trajectory has the form

$$\Xi = (\xi_0, \dots, \xi_H), \qquad \xi_h \in T_{x_h}\mathcal{P}. \tag{55}$$

Given pointwise metrics $g^{\mathcal{P}}_{x_h}$, the canonical product metric is

$$g^{\mathrm{prod}}_\tau(\Xi, \Psi) = \sum_{h=0}^{H} g^{\mathcal{P}}_{x_h}(\xi_h, \psi_h). \tag{56}$$

This metric is block diagonal and treats horizon positions independently. A more general metric on $\mathcal{M}$ may also include cross-position terms. In a local chart or a chosen trivialization, its matrix representation can contain off-diagonal blocks that couple perturbations at different horizon positions. Such a metric measures the geometry of an entire trajectory rather than a set of unrelated points.

C.10 Riemannian Metrics and Covariance Matrices

A Riemannian metric matrix and a covariance matrix may both be symmetric and positive definite, but they play different roles. A covariance matrix Σ describes the second-order variation of a random vector; for example,

$$\epsilon \sim \mathcal{N}(0, \Sigma). \tag{57}$$

By contrast, a metric matrix $G(x)$ represents a bilinear form used to measure tangent vectors:

$$\|v\|_{g_x}^2 = v^\top G(x) v. \tag{58}$$

Changing Σ changes a probability law, whereas changing G changes the notion of geometric magnitude, length, and distance.

The two objects are related in some constructions but should not be identified. For example, the inverse covariance Σ^{-1} defines the Mahalanobis quadratic form

$$\|v\|_{\Sigma^{-1}}^2 = v^\top \Sigma^{-1} v, \tag{59}$$

which can be interpreted as a constant metric in a fixed coordinate system. More generally, covariance matrices and metric matrices obey different coordinate-transformation rules: covariance acts on vectors, while a metric is a covariant bilinear form. Their meaning is therefore determined by the role they play, not by positive definiteness alone.

D Theoretical Guarantees, Proofs, and Scope

D.1 Properties of the random metric kernel

Proof of Proposition 1. Using the self-adjointness of S_τ^ω with respect to $g_\tau^{\mathcal{M}}$, we have

$$\tilde{g}_{t,\tau}^{\omega}(\Xi, \Psi) = g_\tau^{\mathcal{M}}\left(e^{\varepsilon_t S_\tau^\omega}\Xi, \Psi\right). \tag{60}$$

By the definition of the Riemannian gradient, for any $\Xi \in T_\tau\mathcal{M}$,

$$d\phi_t^\omega[\Xi] = \tilde{g}_{t,\tau}^{\omega}(V_t^\omega, \Xi) = g_\tau^{\mathcal{M}}\left(e^{\varepsilon_t S_\tau^\omega} V_t^\omega, \Xi\right) = g_\tau^{\mathcal{M}}\left(\mathrm{grad}_{g_\tau^{\mathcal{M}}}\phi_t^\omega, \Xi\right). \tag{61}$$

Since $g_\tau^{\mathcal{M}}$ is nondegenerate,

$$V_t^\omega = e^{-\varepsilon_t S_\tau^\omega}\,\mathrm{grad}_{g_\tau^{\mathcal{M}}}\phi_t^\omega. \tag{62}$$

□

Smoothness of the vector field gives local existence and uniqueness. A sufficient condition for existence on the full unit interval is bounded speed in a complete reference Riemannian metric. More general linear-growth conditions can also prevent escape in finite time. The kernel definition assumes such existence; local Lipschitz continuity alone is insufficient on a noncompact manifold.

Sampling self-adjoint fields. On $SE(3)^{H+1}$ or $\mathbb{R}^{d(H+1)}$, a global trivialized frame is available. Let $G(t,\tau)$ be the matrix of $\tilde{g}_{t,\tau}^{w}$ in that frame. For any smooth symmetric matrix field $E_m(t,\tau)$, the operator represented by

$$B_m(t,\tau) = G(t,\tau)^{-1/2} E_m(t,\tau) G(t,\tau)^{1/2} \tag{63}$$

satisfies $GB_m = B_m^\top G$. This gives a concrete way to satisfy Eq. 5 for either the base or refined metric. The scalar fields ρ_m, symmetric fields E_m, and amplitude schedule remain design choices. Spatially varying coefficients should be checked for nonconstant deformation energy and non-explosion.

D.2 Local KL response

Fix B, t, τ and a tangent Ξ. Let $f(s) = D_{\mathrm{KL}}(q_{\phi,t}^B(\cdot \mid \tau) \| q_{\phi,t}^B(\cdot \mid \mathrm{Exp}_\tau^{\mathcal{M}}(s\Xi)))$. Under the regularity conditions in the main text, normalization gives $f(0) = f'(0) = 0$ and

$$f''(0) = \mathbb{E}_{q_{\phi,t}^B(\cdot|\tau)}\left[(\mathrm{d}_\tau \log q_{\phi,t}^B[\Xi])^2\right].$$

Thus $f(\delta) + f(-\delta) = \delta^2 f''(0) + o(\delta^2)$. Averaging over B and dividing by δ^2 proves Eq. 19. The zero first derivative makes the second-order term independent of the acceleration of the chosen curve.

D.3 Response stability under metric corruption

We analyze one continuous internal update at a fixed noising time t. Using the same predictive laws and reference measures as above, define

$$\mathrm{Hel}_{t,\mathcal{B}}^2(\tau,\rho) = \frac{1}{2|\mathcal{B}|}\sum_{B\in\mathcal{B}}\int\left(\sqrt{q_{\phi,t}^B(z_B \mid \tau)} - \sqrt{q_{\phi,t}^B(z_B \mid \rho)}\right)^2 \mathrm{d}\mu_B(z_B). \tag{64}$$

This is the averaged squared Hellinger distance between the predicted distributions. The following subsections provide the proofs and discuss extensions and limitations of these guarantees.

Proposition 2 (Finite-step motion–response budget)**.** *Assume the square-root density maps are continuously differentiable in $L^2(\mu_B)$ and $\mathcal{H}_{t,\rho}$ is finite along the path. For $\alpha > 0$, let*

$$g_{\alpha,t,\rho} = g_\rho^{\mathcal{M}} + \alpha\mathcal{H}_{t,\rho}. \tag{65}$$

Assume this is a smooth metric on a neighborhood of the path. Use it as $\bar{g}_t$ in Eqs. 6–11, with a smooth field self-adjoint under this same metric. Suppose the flow γ exists on $[0,1]$, with $\gamma(0)=\tau$ and $\gamma(1)=\tau^+$. Define

$$m_{t,\omega}(\gamma)=\sup_{\nu\in[0,1]}\|\varepsilon(t)S_t^{\omega}(\gamma(\nu))\|_{\mathrm{op},g_{\alpha,t,\gamma(\nu)}}<\infty,$$
$$\Delta\varphi_{t,\omega}=\varphi_t^{\omega}(\tau^+)-\varphi_t^{\omega}(\tau),\qquad \mathcal{E}_{\mathcal{M}}(\gamma)=\int_0^1 g^{\mathcal{M}}_{\gamma(\nu)}(\dot\gamma,\dot\gamma)\,\mathrm{d}\nu. \tag{66}$$

Then

$$\boxed{\mathcal{E}_{\mathcal{M}}(\gamma)+8\alpha\,\mathrm{Hel}^2_{t,\mathcal{B}}(\tau,\tau^+)\le e^{m_{t,\omega}(\gamma)}\Delta\varphi_{t,\omega}.} \tag{67}$$

In particular, the base-metric distance satisfies $d^2_{\mathcal{M}}(\tau,\tau^+)\le\mathcal{E}_{\mathcal{M}}(\gamma)$, and the response is bounded by $e^{m_{t,\omega}(\gamma)}\Delta\varphi_{t,\omega}/(8\alpha)$.

The bound links a complete metric-induced update to its predictive response. Both the distortion and energy increment can depend on α, so it does not imply monotone improvement as α increases.

Corollary 1 (Learned response geometry)**.** *For $\alpha>0$, set $\widehat{\mathcal{H}}_{t,\rho}(\Xi,\Psi)=g^{\mathcal{M}}_{\rho}(R_{\psi}(t,\rho)\Xi,\Psi)$ and $\widehat{g}_{\alpha,t,\rho}=g^{\mathcal{M}}_{\rho}+\alpha\widehat{\mathcal{H}}_{t,\rho}$. Construct $\widehat{\bar{g}}^{\omega}_t$, $\widehat{\varphi}^{\omega}_t$, and the flow $\widehat{\gamma}$ using this metric in Eqs. 6–11. The sampled field must be self-adjoint under $\widehat{g}_{\alpha}$. Define $\widehat{m}_{t,\omega}$ and $\Delta\widehat{\varphi}_{t,\omega}$ by Eq. 66 for this flow and metric. If, along the path,*

$$\mathcal{H}_{t,\rho}(\Xi,\Xi)\le\widehat{\mathcal{H}}_{t,\rho}(\Xi,\Xi)+\eta g^{\mathcal{M}}_{\rho}(\Xi,\Xi),\qquad \eta\ge0, \tag{68}$$

for all tangent directions, then

$$\mathrm{Hel}^2_{t,\mathcal{B}}(\widehat{\gamma}(0),\widehat{\gamma}(1))\le\frac{e^{\widehat{m}_{t,\omega}}}{8}\max\{\alpha^{-1},\eta\}\,\Delta\widehat{\varphi}_{t,\omega}. \tag{69}$$

A small probe-fitting loss does not establish the one-sided spectral assumption in Eq. 68.

Proposition 3 (A case where updating geometry matters)**.** *Consider one predictive law $q(z\mid x,y)=\mathcal{N}(z;x^2,1)$ on $\mathbb{R}^2$ with base metric I. Its response metric is $g_\alpha=\mathrm{diag}(1+4\alpha x^2,1)$. There exists a smooth Gaussian-coefficient field of the form 23, self-adjoint under all these metrics, such that the update from $(0,0)$ satisfies, almost surely,*

$$\mathrm{Hel}^2_q\big((0,0),(x_\alpha(1),y(1))\big)=\Theta(\alpha^{-4/3}),\qquad \alpha\to\infty, \tag{70}$$

while $y(1)>0$ is independent of α. Here $\mathrm{Hel}^2_q(a,b)=\frac12\int(\sqrt{q(z\mid a)}-\sqrt{q(z\mid b)})^2\mathrm{d}z$. Freezing the response tensor at the initial point instead gives a strictly positive response change independent of α.

This example shows that updating the response geometry can matter; it is not a universal comparison with fixed metrics. These guarantees concern the specified predictive laws and continuous internal flow. They do not establish source matching, unbiased noise, or downstream task success.

D.4 Finite-step motion–response budget

For fixed t, define the Hilbert space $\mathcal{K}=\bigoplus_{B\in\mathcal{B}}L^2(\mu_B)$ and the square-root map

$$U_t(\tau)=\frac{1}{\sqrt{|\mathcal{B}|}}\big(\sqrt{q^B_{\phi,t}(\cdot\mid\tau)}\big)_{B\in\mathcal{B}}. \tag{71}$$

Then $\|U_t(\tau)\|_{\mathcal{K}}=1$ and

$$\|\mathrm{D}U_t(\tau)[\Xi]\|^2_{\mathcal{K}}=\frac14\mathcal{H}_{t,\tau}(\Xi,\Xi),\qquad \mathrm{Hel}^2_{t,\mathcal{B}}(\tau,\rho)=\frac12\|U_t(\tau)-U_t(\rho)\|^2_{\mathcal{K}}. \tag{72}$$

The first identity follows by differentiating each square-root density. The assumed L^2 differentiability justifies the Hilbert-space fundamental theorem of calculus.

Proof of Proposition 2. For the path γ, Cauchy–Schwarz on the unit interval gives

$$\begin{aligned}\mathrm{Hel}^2_{t,\mathcal{B}}(\gamma(0),\gamma(1)) &= \frac{1}{2}\left\|\int_0^1 \mathrm{D}U_t(\gamma(\nu))[\dot\gamma(\nu)]\,\mathrm{d}\nu\right\|^2_{\mathcal{K}} \\ &\le \frac{1}{8}\int_0^1 \mathcal{H}_{t,\gamma(\nu)}(\dot\gamma,\dot\gamma)\,\mathrm{d}\nu.\end{aligned} \tag{73}$$

Write $m = m_{t,\omega}(\gamma)$. Self-adjointness and the operator-norm bound imply the quadratic-form comparison

$$e^{-m} g_{\alpha,t} \preceq \widetilde{g}^{\omega}_t \preceq e^{m} g_{\alpha,t} \quad \text{along } \gamma. \tag{74}$$

No commutativity between the response tensor and random operator is needed. Combining the path bound and metric comparison yields

$$\begin{aligned}\mathcal{E}_{\mathcal{M}}(\gamma) + 8\alpha\,\mathrm{Hel}^2_{t,\mathcal{B}}(\tau,\tau^+) &\le \int_0^1 g_{\alpha,t,\gamma(\nu)}(\dot\gamma,\dot\gamma)\,\mathrm{d}\nu \\ &\le e^{m}\int_0^1 \widetilde{g}^{\omega}_{t,\gamma(\nu)}(\dot\gamma,\dot\gamma)\,\mathrm{d}\nu = e^{m}\Delta\varphi_{t,\omega}.\end{aligned} \tag{75}$$

The energy increment is nonnegative. The base distance is at most the path length, whose square is at most $\mathcal{E}_{\mathcal{M}}(\gamma)$ on a unit interval. This proves both endpoint bounds. □

D.5 Approximate response geometry

Proof of Corollary 1. Let $a = g^{\mathcal{M}}_{\rho}(\Xi,\Xi) \ge 0$ and $b = \widehat{\mathcal{H}}_{t,\rho}(\Xi,\Xi) \ge 0$. For $\kappa = \max\{\alpha^{-1},\eta\}$,

$$\mathcal{H}_{t,\rho}(\Xi,\Xi) \le b + \eta a \le \kappa(a+\alpha b) = \kappa \widehat{g}_{\alpha,t,\rho}(\Xi,\Xi). \tag{76}$$

Apply Eq. 73 to $\widehat{\gamma}$, compare $\widehat{g}_\alpha$ with its random deformation, and use the energy identity for $\widehat{\varphi}^{\omega}_t$. This gives Eq. 69. □

The same argument retains a motion term:

$$(1-\alpha\eta)\mathcal{E}_{\mathcal{M}}(\widehat{\gamma}) + 8\alpha\,\mathrm{Hel}^2_{t,\mathcal{B}}(\widehat{\gamma}(0),\widehat{\gamma}(1)) \le e^{\widehat{m}_{t,\omega}}\Delta\widehat{\varphi}_{t,\omega}. \tag{77}$$

The motion coefficient is nonnegative when $\alpha\eta \le 1$. The error assumption is a one-sided spectral bound. A small empirical loss on a finite probe set does not imply it. The proof only needs the bound on visited points and velocity directions, but relating these to training probes requires an additional generalization argument.

D.6 Adaptive and frozen response geometry

Proof of Proposition 3. For $q(z \mid x,y) = \mathcal{N}(z; x^2, 1)$, $\partial_x \log q = 2x(z-x^2)$ and $\partial_y \log q = 0$. Hence

$$\mathcal{H}(x,y) = \mathrm{diag}(4x^2, 0), \qquad G_\alpha(x,y) = \mathrm{diag}(1+4\alpha x^2, 1). \tag{78}$$

Fix a noising time with $\varepsilon = \varepsilon(t) > 0$. Draw independent $z_x, z_y \sim \mathcal{N}(0,1)$ and set

$$S^{\omega}(x,y) = \mathrm{diag}\left(z_x \tanh(1+x), z_y \tanh(1+y)\right). \tag{79}$$

This is a two-term field of the form 23. It is smooth and self-adjoint under G_α for every α. Its distortion is bounded by $\varepsilon\max\{|z_x|,|z_y|\}$ on the whole plane. The deformation energy is

$$\varphi^{\omega}(x,y) = \frac{\varepsilon^2}{2}\left[z_x^2 \tanh^2(1+x) + z_y^2 \tanh^2(1+y)\right]. \tag{80}$$

The Hilbert–Schmidt norm is the sum of squared diagonal entries in the metric-orthonormal frame, so no state-dependent energy term is omitted. Define

$$f_z(u) = \varepsilon^2 z^2 \tanh(1+u)\,\mathrm{sech}^2(1+u) e^{-\varepsilon z \tanh(1+u)}. \tag{81}$$

The update from $(0,0)$ is

$$\dot x_\alpha = \frac{f_{z_x}(x_\alpha)}{1+4\alpha x_\alpha^2}, \qquad \dot y = f_{z_y}(y), \qquad x_\alpha(0) = y(0) = 0. \tag{82}$$

These smooth vector fields are bounded, so the solutions exist on $[0,1]$. Almost surely $z_x, z_y \neq 0$; their velocities are positive on the nonnegative axes. In particular, $y(1) > 0$ and is independent of α.

Write $b_\alpha = x_\alpha(1)$. Separation of variables gives

$$1 = \int_0^{b_\alpha} \frac{1 + 4\alpha u^2}{f_{z_x}(u)} \, \mathrm{d}u. \tag{83}$$

We have $0 < b_\alpha \leq b_0$. On $[0, b_0]$, the function f_{z_x} is positive and bounded above. The integral identity first gives $b_\alpha = O(\alpha^{-1/3})$. Continuity at zero then implies

$$b_\alpha \sim \left(\frac{3 f_{z_x}(0)}{4\alpha} \right)^{1/3}. \tag{84}$$

For equal-variance Gaussians, the squared Hellinger response is

$$1 - e^{-b_\alpha^4/8} \sim \frac{1}{8} \left(\frac{3 f_{z_x}(0)}{4} \right)^{4/3} \alpha^{-4/3}. \tag{85}$$

The constants are positive and may depend on the field realization, but not on α. At the initial point, $\mathcal{H}(0,0) = 0$. Freezing this tensor gives $G_{\mathrm{frozen}} = I$ for all α. Its endpoint is $b_0 > 0$, so the response remains the positive constant $1 - e^{-b_0^4/8}$. This proves the claim. □

The example uses a single response head and isolates two deformation directions. It does not establish equal displacement budgets, centered increments, or universal superiority over a fixed metric.

D.7 Prediction error and numerical integration

Let $P_t(\tau) = (p_t^B(\cdot \mid \tau))_{B \in \mathcal{B}}$ be the actual conditional laws of interest and $Q_t(\tau) = (q_{\phi,t}^B(\cdot \mid \tau))_{B \in \mathcal{B}}$. For families P, Q, write

$$\mathrm{Hel}^2_{\mathcal{B}}(P, Q) = \frac{1}{2|\mathcal{B}|} \sum_B \int (\sqrt{p^B} - \sqrt{q^B})^2 \mathrm{d}\mu_B.$$

Suppose $\mathrm{Hel}_{\mathcal{B}}(P_t(x), Q_t(x)) \leq \zeta(x)$ at both endpoints. The triangle inequality gives

$$\mathrm{Hel}_{\mathcal{B}}(P_t(\tau), P_t(\tau^+)) \leq \zeta(\tau) + \zeta(\tau^+) + \sqrt{\frac{\kappa e^m}{8} \Delta\varphi}, \tag{86}$$

where $m, \Delta\varphi$ are computed for the actual flow and $\kappa = \alpha^{-1}$ for the ideal metric or $\max\{\alpha^{-1}, \eta\}$ for the learned metric. Reference-chain likelihood training alone does not guarantee the endpoint error assumption on the refined chain.

For a piecewise smooth numerical interpolant γ^{num}, set

$$r(\nu) = \dot{\gamma}^{\mathrm{num}}(\nu) - \mathrm{grad}_{\widetilde{g}} \, \varphi(\gamma^{\mathrm{num}}(\nu)), \qquad \mathcal{R} = \int_0^1 \|r(\nu)\|_{\widetilde{g}}^2 \, \mathrm{d}\nu. \tag{87}$$

Let $E = \int_0^1 \|\dot{\gamma}^{\mathrm{num}}\|_{\widetilde{g}}^2 \mathrm{d}\nu$ and $\Delta\varphi^{\mathrm{num}}$ be the endpoint energy change. The chain rule and Cauchy–Schwarz imply

$$E = \Delta\varphi^{\mathrm{num}} + \int_0^1 \widetilde{g}(r, \dot{\gamma}^{\mathrm{num}}) \, \mathrm{d}\nu \leq [\Delta\varphi^{\mathrm{num}}]_+ + \sqrt{\mathcal{R} E}.$$

Thus $\sqrt{E} \leq \sqrt{[\Delta\varphi^{\mathrm{num}}]_+} + \sqrt{\mathcal{R}}$. Under the same response domination and metric comparison along this path,

$$\mathrm{Hel}^2_{t,\mathcal{B}}(\gamma^{\mathrm{num}}(0), \gamma^{\mathrm{num}}(1)) \leq \frac{\kappa e^m}{8} \left(\sqrt{[\Delta\varphi^{\mathrm{num}}]_+} + \sqrt{\mathcal{R}} \right)^2, \tag{88}$$

where $[a]_+ = \max\{a, 0\}$. A fixed integration step count does not provide a bound on $\mathcal{R}$.

D.8 Random draws and changes in noising time

The distortion constant depends on the field realization. If the right-hand side is integrable and response domination holds with deterministic κ, conditioning on the current trajectory gives

$$\mathbb{E}_\omega[\mathrm{Hel}^2_{t,\mathcal{B}}(\tau,\tau^+) \mid \tau] \le \frac{\kappa}{8}\,\mathbb{E}_\omega[e^m \Delta\varphi \mid \tau]. \tag{89}$$

Gaussian coefficients need not admit a uniform deterministic distortion bound.

The one-step theorem compares predictions at the same noising label. For the chain evaluated at s_k, write $a_k = \sqrt{\kappa_k e^{m_k}\Delta\varphi_k/8}$ and

$$b_k = \mathrm{Hel}_{\mathcal{B}}\left(Q_{s_k}(\tau_{t_{k+1}}), Q_{s_{k+1}}(\tau_{t_{k+1}})\right), \quad k = 0, \dots, K-2.$$

Repeated use of the triangle inequality yields

$$\mathrm{Hel}_{\mathcal{B}}\left(Q_{s_0}(\tau_{t_0}), Q_{s_{K-1}}(\tau_{t_K})\right) \le \sum_{k=0}^{K-1} a_k + \sum_{k=0}^{K-2} b_k. \tag{90}$$

The b_k terms account for explicit changes of the predictive family with noising time. None of these response bounds proves convergence to a prescribed source distribution.

Algorithm 1 Two-Axis Geometry-First Forward Process

Require: Expert data $\mathcal{D} = \{(\tau_0^{(i)}, c^{(i)})\}_{i=1}^N$; base metric $g^{\mathcal{M}}$; blocks $\mathcal{B}$; number of probes J and scale δ; refinement weight α; time grid $0 = t_0 < \cdots < t_K = 1$; random-field construction and amplitude $\varepsilon(t)$.

Ensure: Refined metric field $g^{\mathcal{M}}_{\psi,t,\rho}$ and forward training data $\mathcal{D}_{\mathrm{fwd}}$.

Learn horizon-response geometry (h axis)

Generate reference chains with the base random metric kernel, using $s_k = t_{k+1}$.
Store reference tuples (c, τ_0, t, τ_t) in $\mathcal{D}_{\mathrm{ref}}$.
Fit $q^B_{\phi,t}$ using $\mathcal{L}_{\mathrm{resp}}$ (Eq. 15); freeze ϕ.
$\mathcal{D}_{\mathrm{met}} \leftarrow \emptyset$.
for all $(c, \tau_0, t, \tau_t) \in \mathcal{D}_{\mathrm{ref}}$ **do**
    Construct structured probes $\{\Xi^{(j)}_{\tau_t}\}_{j=1}^J$ using Eq. 16, normalized under $g^{\mathcal{M}}_{\tau_t}$.
    for $j = 1, \ldots, J$ **do**
        $\tau^{\pm,j} \leftarrow \mathrm{Exp}^{\mathcal{M}}_{\tau_t}(\pm\delta\Xi^{(j)}_{\tau_t})$.
        Compute $r^{\delta}_j(t, \tau_t)$ using Eq. 17.
        Add $(t, \tau_t, \Xi^{(j)}_{\tau_t}, r^{\delta}_j(t, \tau_t))$ to $\mathcal{D}_{\mathrm{met}}$.
    end for
end for
Fit $R_\psi = L_\psi L^\star_\psi$ on $\mathcal{D}_{\mathrm{met}}$ using $\mathcal{L}_{\mathrm{met}}$ (Eq. 21); freeze ψ.
Define $\bar{g}_{t,\rho} \leftarrow g^{\mathcal{M}}_{\psi,t,\rho}$ using Eq. 22.

Generate random metric corruption (t axis)

If source calibration is used, calibrate with pilot chains under the refined metric.
Freeze the forward configuration.
$\mathcal{D}_{\mathrm{fwd}} \leftarrow \emptyset$.
for all $(\tau_0, c) \in \mathcal{D}$ **do**
    $\tau_{t_0} \leftarrow \tau_0$.
    for $k = 0, \ldots, K-1$ **do**
        $s_k \leftarrow t_{k+1}$.
        Draw a fresh ω_k and construct a smooth field $S^{\omega_k}_{s_k}$ that is self-adjoint under $\bar{g}_{s_k}$.
        Define the deformed metric field

$$\widetilde{g}^{\omega_k}_{s_k} \leftarrow \Phi_{\varepsilon(s_k)}(\bar{g}_{s_k}, S^{\omega_k}_{s_k}).$$

        Define the deformation energy

$$\varphi^{\omega_k}_{s_k}(\rho) \leftarrow \frac{\varepsilon(s_k)^2}{2} \left\| S^{\omega_k}_{s_k}(\rho) \right\|^2_{\mathrm{HS},\bar{g}_{s_k},\rho}.$$

        Integrate, with s_k and ω_k fixed,

$$\dot{\gamma}_k(\nu) = \mathrm{grad}_{\widetilde{g}^{\omega_k}_{s_k}} \varphi^{\omega_k}_{s_k}(\gamma_k(\nu)), \qquad \gamma_k(0) = \tau_{t_k}, \quad \nu \in [0, 1].$$

        $\tau_{t_{k+1}} \leftarrow \gamma_k(1)$.
    end for
    Add $(c, \tau_{t_0:t_K}, \gamma_{0:K-1})$ to $\mathcal{D}_{\mathrm{fwd}}$.
end for
return $g^{\mathcal{M}}_{\psi,\cdot,\cdot}, \mathcal{D}_{\mathrm{fwd}}$.

Freezing ψ fixes the learned parameters, not the metric values. Within each transition, the metric and operator field are evaluated at the evolving state $\gamma_k(\nu)$, and the energy differential includes all state dependence. Paths may be stored as numerical integration samples for the reverse objectives in Sec. 3.4.

Source calibration does not by itself guarantee terminal matching. If calibration changes the reference corruption configuration, the reference data and response geometry should be updated accordingly. Alternatively, a conditional terminal sampler is fitted to $\{(c, \tau_{t_K})\}$ from $\mathcal{D}_{\mathrm{fwd}}$ after the forward configuration is frozen.

E Additional Experimental Details

E.1 Forward Corruption Across Geometries (Q1)

Trajectory families. Each space contains three smooth trajectories with a shared initial segment and different later motions. Each trajectory has $H = 80$ actions. The planar trajectories branch in $\mathbb{R}^2$; the spherical trajectories share an arc before separating in longitude; the rotation trajectories share an initial rotation before diverging into three orientation modes. Distances are intrinsic to the corresponding space. On $SO(3)$, increments use the left-trivialized representation of the method. For each method, space, and corruption level, we use 40 random draws per mode, giving 120 corrupted trajectories.

Random metric field. The field is instantiated through an intrinsic radial-basis expansion,

$$S^{\omega}(x) = \sum_{m=1}^{M} \sum_{j=1}^{J} a_{mj}^{\omega} k_{\ell}(d(x, c_j))\, B_m, \qquad a_{mj}^{\omega} \sim \mathcal{N}(0, \sigma_S^2), \tag{91}$$

where d is the intrinsic distance, $\{B_m\}$ is a Frobenius-orthonormal basis of symmetric operators in the tangent representation, and $k_{\ell}(r) = \exp\!\big(-r^2/(2\ell^2)\big)$. We use $J = 12$, $\sigma_S = 0.18$, and eight integration steps for the metric-induced flow. The length scales are 1.2, 1.5, and 2.2 for $\mathbb{R}^2$, S^2, and $SO(3)$, respectively, and are fixed across corruption levels. Coordinates are induced by the deformation energy and metric-gradient flow of Sec. 3.2.

Baselines. Coordinate Gaussian adds independent noise in the chosen coordinate representation and projects the result onto S^2 or $SO(3)$ when needed. Correlated Gaussian samples a squared-exponential Gaussian process over the action horizon and maps tangent perturbations through the exponential map; the horizon is normalized to $[0, 1]$ and the correlation length is 0.13. Exp-map Gaussian uses independent isotropic tangent increments. Fixed-metric Brownian takes ten Brownian steps under the canonical metric, following standard manifold diffusion constructions (De Bortoli et al., 2022b; Huang et al., 2022).

Matched corruption strength. After a method proposes a noisy trajectory, we radially rescale its intrinsic tangent displacements from the clean trajectory to satisfy

$$\left(\frac{1}{H} \sum_{h=1}^{H} d^2(A_{0,h}, A_{t,h}) \right)^{1/2} = r_t, \qquad r_t = 0.30t, \tag{92}$$

for $t \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$. This rescaling is used only for controlled evaluation and is not part of the proposed transition kernel. RMS displacement is recorded to verify matching.

Geometry metrics. We measure distortion of distances between nearby actions:

$$D_{\text{local}} = \frac{\displaystyle\sum_{k=1}^{K} \sum_{h=1}^{H-k} |d(A_{t,h}, A_{t,h+k}) - d(A_{0,h}, A_{0,h+k})|}{\displaystyle\sum_{k=1}^{K} \sum_{h=1}^{H-k} d(A_{0,h}, A_{0,h+k})}, \qquad K = 3. \tag{93}$$

Trajectory roughness measures second-order variation of intrinsic increments, using parallel transport on S^2 and left-trivialized increments on $SO(3)$. Path-length distortion is

$$D_L = \frac{|L(\tau_t) - L(\tau_0)|}{L(\tau_0)}, \qquad L(\tau) = \sum_{h=1}^{H-1} d(A_h, A_{h+1}). \tag{94}$$

Lower values are better for all three metrics.

Complete results. Table 7 contains all five methods at $t = 1$. Figure 1 in the main text shows the qualitative comparison across all three spaces. Q1 evaluates the base kernel without response-based metric refinement.

Table 7: Complete matched-strength comparison at $t = 1$. Intrinsic RMS displacement is 0.30 for every trajectory. Values are mean±std over 120 corrupted trajectories; lower is better.

Space	Method	Local geom. ↓	Roughness ↓	Path-length dist. ↓
$\mathbb{R}^2$	Coordinate Gaussian	3.784 ± 0.270	2.9625 ± 0.1073	8.321 ± 0.612
	Correlated Gaussian	0.414 ± 0.097	0.0116 ± 0.0051	0.175 ± 0.105
	Exp-map Gaussian	3.784 ± 0.242	2.9507 ± 0.1226	8.284 ± 0.594
	Fixed-metric Brownian	3.762 ± 0.257	2.9442 ± 0.1232	8.207 ± 0.604
	Random Metric (Ours)	$\mathbf{0.329 \pm 0.099}$	$\mathbf{0.0015 \pm 0.0010}$	$\mathbf{0.120 \pm 0.077}$
S^2	Coordinate Gaussian	7.265 ± 0.907	2.9621 ± 0.1215	15.393 ± 1.943
	Correlated Gaussian	0.795 ± 0.324	0.1955 ± 0.5566	0.616 ± 0.385
	Exp-map Gaussian	7.386 ± 0.958	2.9762 ± 0.1258	15.700 ± 2.073
	Fixed-metric Brownian	7.406 ± 0.940	2.9586 ± 0.1180	15.629 ± 1.884
	Random Metric (Ours)	$\mathbf{0.351 \pm 0.107}$	$\mathbf{0.0013 \pm 0.0010}$	$\mathbf{0.308 \pm 0.144}$
$SO(3)$	Coordinate Gaussian	9.359 ± 1.285	2.9840 ± 0.0834	19.559 ± 2.556
	Correlated Gaussian	1.051 ± 0.477	0.4087 ± 0.8123	0.993 ± 0.538
	Exp-map Gaussian	9.617 ± 1.326	2.9789 ± 0.1046	20.071 ± 2.682
	Fixed-metric Brownian	9.651 ± 1.369	2.9807 ± 0.0880	20.152 ± 2.657
	Random Metric (Ours)	$\mathbf{0.459 \pm 0.116}$	$\mathbf{0.0015 \pm 0.0013}$	$\mathbf{0.459 \pm 0.116}$

E.2 Learned Reverse Recovery and Generation (Q2)

Data and forward chains. We construct continuous three-mode expert distributions on $\mathbb{R}^2$ and $SO(3)$, with $H = 24$ actions per trajectory. The training set contains 500 expert trajectories and one independent forward chain per trajectory. Each chain has five steps, with matched one-step intrinsic RMS displacement increasing linearly from 0.055 to 0.090. The compared kernels are IID Gaussian, correlated Gaussian, and Random Metric. On $SO(3)$, Gaussian increments are mapped from the Lie algebra with the exponential map.

Reverse model and evaluation. All kernels use the same time-conditioned residual denoiser, optimizer, architecture, and training budget. Given τ_k, the model predicts the chart-space residual to τ_{k-1}. Recovery begins at a held-out terminal corruption with a known clean trajectory. Generation fits an intrinsic-rank PCA representation followed by a three-component full-covariance Gaussian mixture to the terminal law, then applies the learned reverse chain to samples from that source. The prior family is shared across methods; its fitted parameters may differ. PCA rank is selected from the supported terminal variance to exclude numerical-null directions.

Recovery is evaluated using intrinsic RMS error and local-geometry distortion. Generation uses sliced Wasserstein distance (SWD), RBF MMD (Gretton et al., 2012), and local-geometry distortion relative to the nearest expert trajectory. SWD and MMD are evaluated in a data-supported PCA subspace, excluding numerical-null directions instead of amplifying them through whitening. For $SO(3)$, distributional metrics use a common logarithmic chart; trajectory geometry metrics remain intrinsic. The fitted terminal source provides a synthetic generation test rather than a universal source-matching construction.

E.3 Two-Axis Ablation

We further analyze the complementary roles of the response-based h-axis refinement and the stochastic t-axis metric corruption. We ablate response-based metric refinement along h and random metric corruption along t on $\mathbb{R}^2$. The h-only variant uses tangent Gaussian noise under the learned response metric; t only uses the base random metric kernel. Full uses the learned response metric inside that kernel. The metric is trained from frozen cross-segment predictions and structured KL probes, as in Sec. 3.3. We use an independent split and training budget, so its absolute values are compared within this ablation.

Table 8 shows distinct effects. The h-only variant best preserves predictive responses, while t only has the lowest forward geometry distortion, SWD, and generated geometry distortion. Combining the axes reduces recovery RMS from 0.150 for t only to 0.067, the lowest among all variants, and

reduces response change from 0.406 to 0.245. The full model therefore improves recoverability while retaining lower forward geometry distortion than coordinate corruption.

Table 8: Response-based 2×2 ablation on $\mathbb{R}^2$ Response change is defined in Eq. 95; lower is better for every column.

Variant	Resp. change	Fwd. geom.	Recovery RMS	SWD	Gen. geom.
Neither	0.108	0.461	0.119	3.194	0.142
h only	**0.078**	0.458	0.112	2.241	0.132
t only	0.406	**0.074**	0.150	**1.245**	**0.064**
$h + t$ (Full)	0.245	0.205	**0.067**	2.050	0.102

Metric learning. We generate a reference corruption set with the base random metric kernel and train the cross-segment conditional predictor $q^B_{\phi,t}$. After freezing ϕ, we evaluate symmetric structured probes $\tau^{\pm,j} = \mathrm{Exp}_\tau(\pm\delta\Xi^{(j)})$ using the symmetric KL response in Sec. 3.3. These responses supervise the rank-4 positive-semidefinite operator $R_\psi = L_\psi L_\psi^\star$ in the refined trajectory metric $g^{\mathcal{M}}_{\psi,t,\tau}$. We use four contiguous held-out blocks, eight smooth low-frequency horizon probes, $\delta = 0.06$, and $\alpha = 8$.

Variants and training. Neither uses coordinate Gaussian noise. The h-only variant samples tangent Gaussian perturbations under the learned response metric without randomizing that metric. The t-only variant uses the random metric kernel with the base trajectory metric. Full uses the learned response metric as the base geometry of the random metric kernel. All variants use the same five-step RMS schedule, reverse architecture, and optimization budget. The independent Q3 split contains 280 training, 60 validation, and 120 test trajectories; the denoiser is trained for 60 epochs.

Predictive response change. We compare clean and terminal trajectories using the same frozen predictor at noise label 1:

$$D_{\mathrm{resp}} = \frac{1}{|\mathcal{B}|} \sum_{B \in \mathcal{B}} \mathrm{Hel}^2\big(q^B_{\phi,1}(\cdot \mid \tau_0), q^B_{\phi,1}(\cdot \mid \tau_T)\big)\,. \tag{95}$$

Conditioning on a trajectory here denotes conditioning on its unmasked segments, as in Sec. 3.3. Lower values indicate smaller changes in the predictive relations encoded by the response model.

E.4 Robot Evaluation Details (Q3–Q4)

Benchmarks and task metrics. Push-T is a planar pushing task with a low-dimensional Euclidean action space. We report its average and maximum continuous rollout scores. RoboMimic contains the robot-arm tasks Lift, Can, and Square, while Adroit contains the dexterous manipulation tasks Door, Pen, Hammer, and Relocate. Task success rate is the primary metric for both RoboMimic and Adroit.

Robomimic and Adroit Tasks. Figures 4a and 4b show representative execution snapshots on these two benchmark suites, respectively. Each row shows the temporal progression of one manipulation episode.

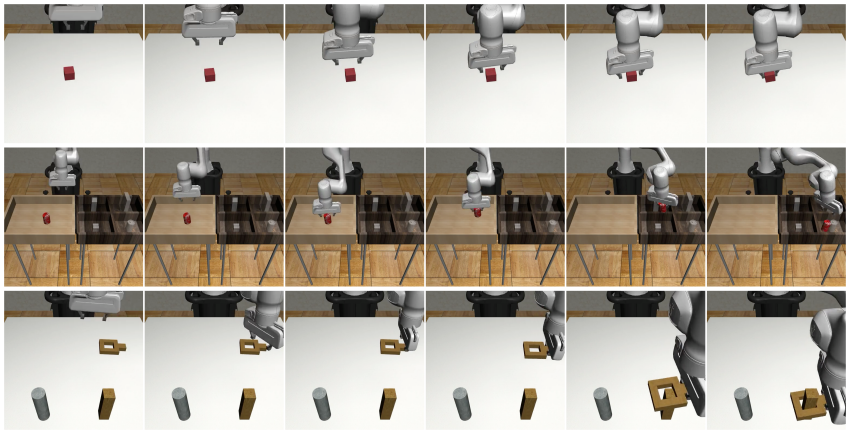

(a) Execution snapshots on the RoboMimic benchmark. From top to bottom: Lift, Can, and Square.

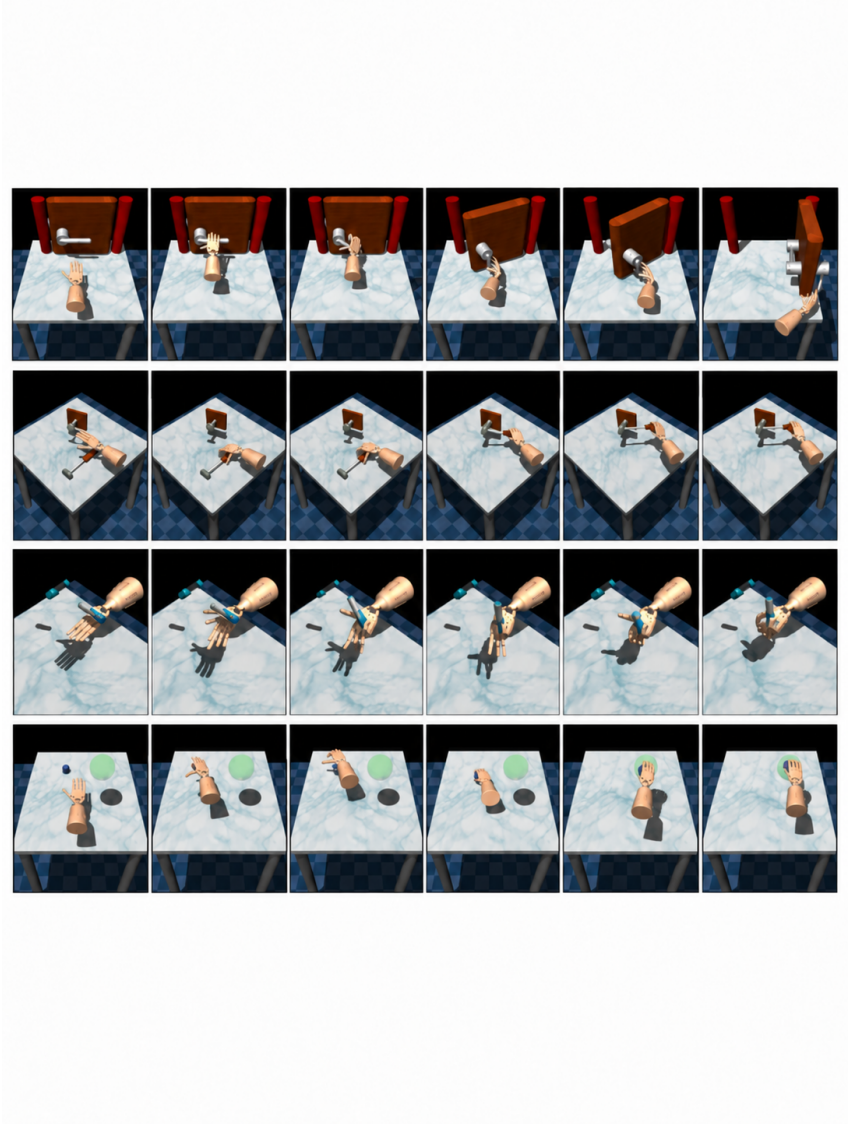

(b) Execution snapshots on the Adroit hand benchmark. From top to bottom: Door, Hammer, Pen, and Relocate.

Figure 4: RoboMimic and Adroit tasks.

Table 9: Trajectory quality averaged within each benchmark. Values are normalized by the corresponding IID Gaussian baseline and reported as mean±standard deviation. Lower is better.

	RoboMimic		Adroit	
Forward corruption	Jerk ↓	Chunk inc. ↓	Jerk ↓	Chunk inc. ↓
IID Gaussian	1.000 ± 0.041	1.000 ± 0.037	1.000 ± 0.048	1.000 ± 0.043
Correlated Gaussian / GP	0.811 ± 0.034	0.842 ± 0.031	0.836 ± 0.038	0.858 ± 0.035
Smooth-field coordinate noise	0.624 ± 0.026	0.719 ± 0.027	0.658 ± 0.031	0.744 ± 0.030
Random Metric (t only)	0.671 ± 0.028	0.641 ± 0.024	0.704 ± 0.033	0.666 ± 0.027
Full ($h + t$)	$\mathbf{0.482 \pm 0.021}$	$\mathbf{0.493 \pm 0.020}$	$\mathbf{0.521 \pm 0.025}$	$\mathbf{0.536 \pm 0.023}$

Controlled comparison. Within each task, all variants use identical demonstrations, observation histories, policy architectures, action horizons, optimizers, numbers of training updates, sampling procedures, and evaluation seeds. Only the forward corruption mechanism is changed. IID Gaussian applies independent coordinate noise. Correlated Gaussian/GP noise introduces temporal correlation across the action horizon. Smooth-field coordinate noise applies a smooth random field directly to action coordinates without modifying their metric. Random Metric (t only) applies stochastic metric corruption to the base trajectory metric, whereas Full ($h + t$) uses the response-refined trajectory metric within the same random metric kernel. Because Push-T has Euclidean actions, exponential-map Gaussian noise is identical to coordinate Gaussian noise: $\mathrm{Exp}_x(v) = x + v$.

Action jerk and chunk inconsistency. Let $\{a_k\}_{k=1}^{N}$ be the executed action sequence. We measure discrete action jerk by

$$J_{\text{jerk}} = \frac{1}{N-3} \sum_{k=1}^{N-3} \|a_{k+3} - 3a_{k+2} + 3a_{k+1} - a_k\|_2^2 . \tag{96}$$

This metric penalizes rapid changes in the local action profile.

At replanning step r, let $A_h^{(r)}$ denote the action predicted for horizon position h. If the policy executes s actions before replanning, predictions for the same physical time are compared using

$$C_{\text{chunk}} = \frac{1}{|\mathcal{I}|} \sum_{(r,h)\in\mathcal{I}} \left\| A_{h+s}^{(r)} - A_h^{(r+1)} \right\|_2^2 , \tag{97}$$

where $\mathcal{I}$ contains all valid overlapping indices. We use $s = 1$ for unit-stride replanning. Lower jerk indicates smoother execution, while lower chunk inconsistency indicates greater agreement between consecutive action chunks. Each quantity is normalized by the IID Gaussian result for the same task before aggregation across tasks.

Rotational components. When an action contains an orientation, we compute angular increments using the vector-valued logarithmic map

$$\omega_h = \mathrm{Log}\big(R_h^\top R_{h+1}\big) , \tag{98}$$

and measure their variation by

$$J_{\text{rot}} = \frac{1}{H-2} \sum_{h=2}^{H-1} \|\omega_h - \omega_{h-1}\|_2^2 . \tag{99}$$

This left-trivialized representation avoids treating orientation coordinates as ordinary Euclidean vectors.

Trajectory quality. Table 9 reports jerk and chunk inconsistency averaged within RoboMimic and Adroit.

E.5 Maze2D Protocol and Metrics (Q5)

Maze2D evaluates whether the forward corruption mechanism remains effective when trajectory structure must be maintained over longer horizons. We consider U-Maze, Medium, and Large,

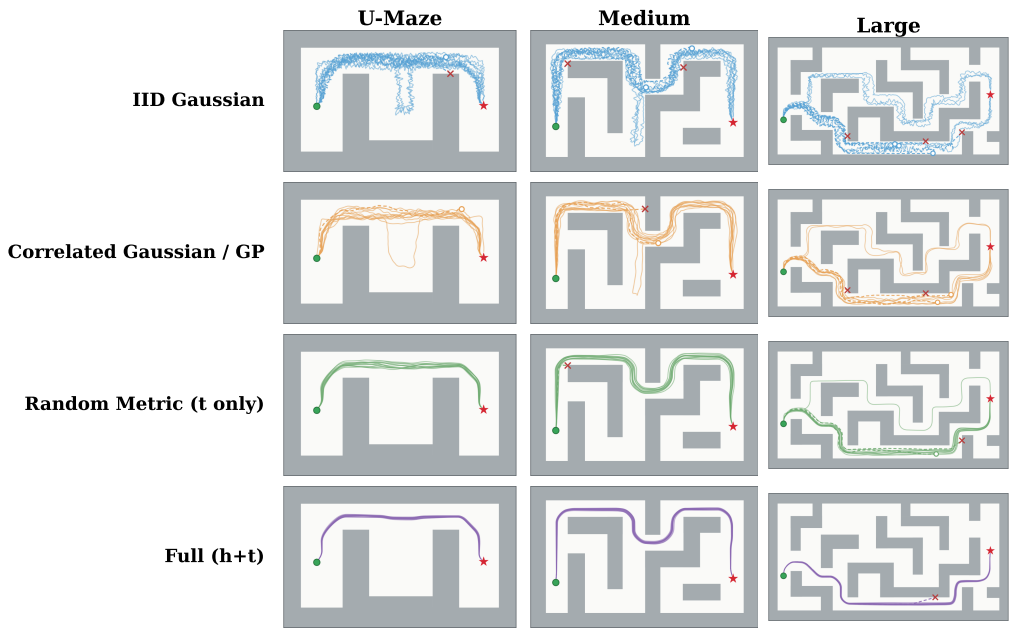

Figure 5: Representative Maze2D trajectories for matched start–goal pairs. The examples show route geometry, local oscillation, and crossings into invalid regions.

which provide increasing layout complexity. Within each environment, all variants use the same trajectory dataset, policy architecture, prediction horizon, training budget, and evaluation initial states.

We report goal success together with three trajectory diagnostics. The invalid-trajectory rate measures the fraction of generated trajectories that intersect the designated invalid set:

$$R_{\text{invalid}} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\left[\tau^{(i)} \cap \mathcal{X}_{\text{invalid}} \neq \emptyset\right]. \tag{100}$$

For successful trajectories, path inefficiency is defined as

$$R_{\text{path}} = \frac{L(\tau)}{L(\tau^\star)}, \tag{101}$$

where $\tau^\star$ is a shortest valid path for the same start–goal pair. A value closer to 1 indicates a more efficient route. Local oscillation is measured by the action jerk in Eq. (96).

In Table 6, path ratio and invalid rate are averaged over the benchmark, while jerk is normalized by the corresponding IID Gaussian baseline. Figure 5 provides qualitative results for matched start–goal pairs.

E.6 Real-Robot Experimental Details

Hardware and tasks. The experimental platform consists of a Kinova Gen3 arm and an Intel RealSense D435 camera (Fig. 6). In *T-block pushing*, the target is indicated by a T-shaped outline on the table. In *Box uprighting*, the robot turns a cardboard box from a flat configuration to an upright configuration.

Demonstrations and observations. We collected 200 demonstrations for *T-block pushing* task and 100 demonstrations for *Box uprighting* task, using kinesthetic teaching. Demonstrations record synchronized camera observations, robot states, and actions. The policy uses RGB images and robot proprioception. The camera records images at an 848×480 resolution with a 30 Hz frequency, and then crops them to a 320×240 resolution for policy training.

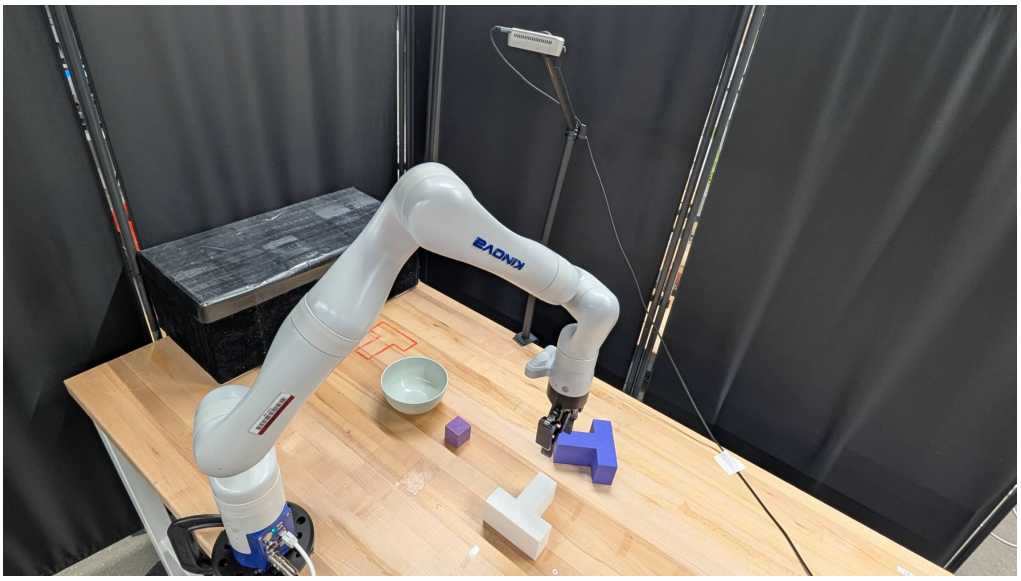

Figure 6: Experimental platform.

Action chunks and execution. The policy predicts an action chunk containing 16 actions. At each replanning step, 8 actions are executed before a new chunk is generated. The control frequency is 50 Hz. Actions are represented as Cartesian end-effector increments and normalized using the training-set preprocessing.

More detailed experimental results For T-block pushing, we vary the difficulty across three settings: pushing along an unobstructed route, pushing around a bowl, and pushing around both a bowl and an additional object. We conduct 20, 40, and 40 trials in these settings, respectively. GFC succeeds in all 20 unobstructed trials, 35 of 40 trials with the bowl, and 32 of 40 trials with both obstacles, yielding 87 successes overall. Diffusion Policy succeeds in only two trials, both in the unobstructed setting.

We also evaluate box uprighting, which requires rotating a flat cardboard box into an upright configuration. GFC succeeds in 75 of 100 trials, whereas Diffusion Policy does not succeed in any trial. Figure 7 shows representative successful and failed executions. Table 10 summarizes the results across task settings.

Table 10: Real-robot success counts for GFC and Diffusion Policy.

Task setting	GFC	Diffusion Policy
T-block: unobstructed route	$20/20$	$2/20$
T-block: bowl obstacle	$35/40$	$0/40$
T-block: bowl and object obstacles	$32/40$	$0/40$
T-block pushing (overall)	$87/100$	$2/100$
Box uprighting	$75/100$	$0/100$

GFC training. Training follows three stages. First, the cross-segment predictor is trained on reference corruptions. Its parameters are then frozen while the response metric is fitted using structured probes. Finally, the refined forward process is fixed for training the reverse model and fitting the conditional terminal source.

The training configuration uses the Adam optimizer with a learning rate of 10^{-4} and a batch size of 64. The response predictor, response metric, and reverse policy are trained for 50, 50, and 500 epochs, respectively. The forward chain uses five corruption steps, with four internal integration steps per random metric transition. The response metric has rank four, and its fitting procedure uses eight structured probes with perturbation scale $\delta = 0.06$ and refinement weight $\alpha = 8$.

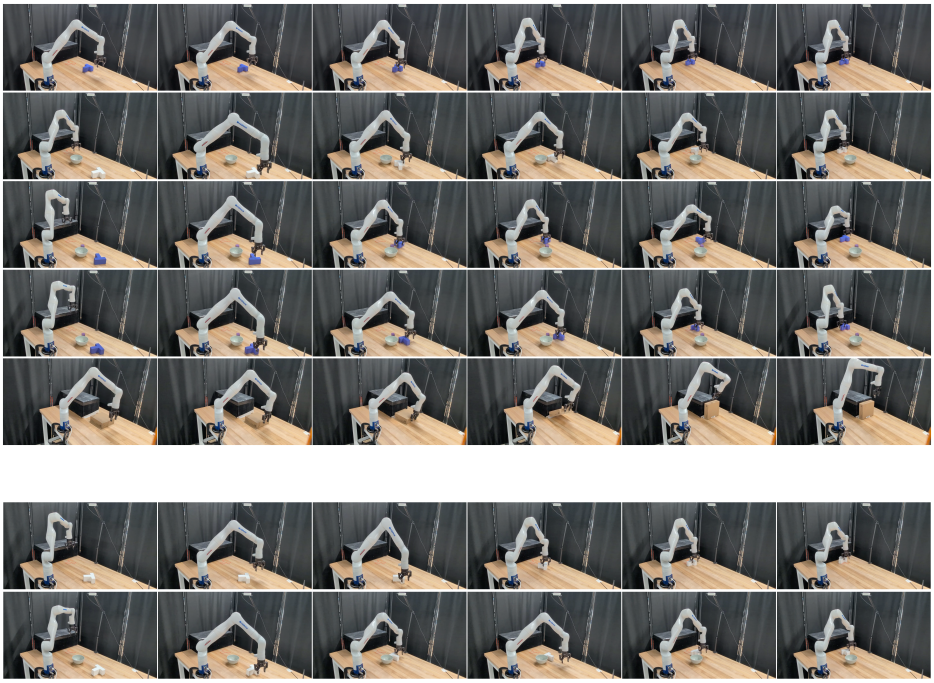

Figure 7: Real-robot outcomes for GFC. The top image shows successful trials, and the bottom image shows failed trials across varied object poses and obstacle arrangements.

Compute and runtime. Training is performed on one NVIDIA RTX 4090 GPU. Total training time is 4 hours for *T-block pushing* and 3 hours for *Box uprighting*, including the response-predictor and metric-fitting stages. Action generation requires 35 ms per chunk at batch size one. This latency includes observation encoding and reverse sampling.