MANIFOLD-STABLE FLOW MATCHING

Anonymous authors
Paper under double-blind review

ABSTRACT

Flow matching (FM) learns generative dynamics through velocity regression. A common assumption in geometric FM variants is that the prior distribution for the data is, in fact, a distribution on the data manifold. However, this requires knowledge of the data manifold, which in many applications is not available. In the absence of such prior knowledge, a low regression error alone does not guarantee that the FM conforms to the true data manifold. Manifold adherence keeps generated samples within valid configurations and is empirically associated with better task performance. In this paper, we introduce manifold-stable flow matching (MSFM), which can start from an arbitrary ambient prior (i.e., a prior distribution that is not necessarily on the manifold). The method leverages tools from nonlinear dynamics, namely contraction theory, to learn tangential transport with prescribed normal contraction to the manifold. Contraction theory provides a theoretical guarantee of convergence to the data manifold within a desired time window (e.g., one second). The construction uses analytical projectors when the data manifold is known, and local affine proxies estimated by principal component analysis of data samples when the data manifold is not known. We derive a family of probability paths compatible with our contraction guarantees and separate the training loss into a learnable manifold tangential term and a manifold normal residual. We establish manifold invariance and transverse convergence under explicit projection and proxy-selection assumptions. An ellipse experiment attains a mean terminal off-manifold error of order 10^{-6} (controlled by our set numerical tolerance). In our Push-T robotic experiments, MSFM raises the success rate from 74% to 82%. In the Robomimic Square task, the success rate increases from 60% to 72%, while the rotation error order decreases from 10^{-2} to 10^{-7} (controlled by our set numerical tolerance). These results show stronger geometric adherence and higher observed task performance, produced by learning the data manifold in generative AI tasks.

1 INTRODUCTION

Flow matching (FM) learns a time-dependent vector field whose flow transports a simple source distribution to a target data distribution through a regression objective rather than repeated ODE simulation during training (Lipman et al., 2023; Albergo et al., 2025). This framework is especially attractive for generative policies because inference reduces to integrating a learned dynamical system. However, many target distributions in robotics and structured generation are supported on, or concentrated near, a lower-dimensional manifold $\mathcal{M} \subset \mathbb{R}^n$. Examples include rotation matrix parameterizations, rigid-body poses, constrained actions, and data-driven task manifolds.

Adherence to this manifold matters. First, samples off the manifold may not be valid outputs: a generated rotation that is not in $\mathrm{SO}(3)$ cannot be executed and must be projected after the fact, and projecting a sample that has drifted far from the manifold can yield an unintended action. Second, when the manifold is unknown, as for expert actions in imitation learning, samples off the data geometry correspond to actions the demonstrations do not support. In closed-loop control, such errors can compound, and adherence to the expert action geometry can matter more for performance than low validation error (Pan et al., 2026).

Conventional FM, however, controls only the expected velocity error on sampled training paths. It matches distributions through training but provides no mechanism that pulls generated trajectories back to the manifold, so a model with low regression loss can still drift away from it. Existing

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.

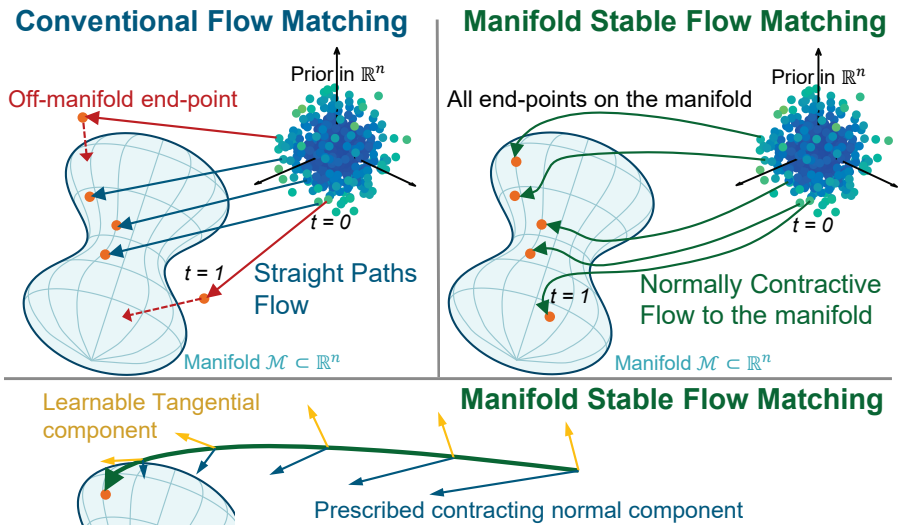


Figure 1: Schematic comparison. Linear conditional FM paths do not necessarily stay near a curved target manifold (top-left panel). MSFM (top-right panel) learns tangential transport and prescribes normal contraction, which guarantees convergence to the manifold. The learned MSFM sampling trajectories can curve because MSFM learns tangential motion, while it prescribes normal contraction (bottom panel). The normal component of the flow uniquely minimizes kinetic action.

remedies are limited: post-hoc projection requires a known manifold and corrects only the final sample, while geometric FM methods require both a known manifold and a prior supported on it.

We propose *manifold-stable flow matching* (MSFM), which separates learned transport along the data geometry from prescribed convergence toward it. For a known manifold, analytical tangent and normal projectors define this separation, while local principal component analysis (PCA) is used for unknown manifolds. Using contraction theory from nonlinear dynamics, we form a normal contractive flow to the data manifold, such that the flow (starting from a prior distribution outside of the manifold) is guaranteed to converge to the manifold within a desired time window (e.g., 1 second). Figure 1 provides a schematic comparison between the conventional FM and our novel MSFM.

Our main contribution is MSFM, a flow-matching framework that provably converges to the data manifold from an arbitrary ambient prior within a finite time window. This guarantee holds regardless of regression accuracy since the network learns only tangential transport while normal contraction is prescribed. We use analytical projectors when the geometry is known and local PCA affine proxies otherwise. We support this framework theoretically by deriving a family of compatible tangent–normal probability paths and an orthogonal loss decomposition that separates what the network learns from what is prescribed. We prove manifold invariance and transverse convergence. We support it empirically on an ellipse and two robotic manipulation tasks: under matched network capacity, MSFM reduces off-manifold error by multiple orders of magnitude on Push-T and Robomimic Square, and raises observed success.

2 Related Work

Intrinsic flow matching. Flow matching learns transport through velocity regression (Lipman et al., 2023). Riemannian flow matching defines paths and tangent vector fields on a known manifold (Chen & Lipman, 2024), with applications to robot policies and pose estimation (Braun et al., 2024; Ding et al., 2025; Ouyang et al., 2025). Extensions modify the regression objective (Zaghen et al., 2025; 2026) or target one-step generation (Zhong et al., 2026). These intrinsic formulations start from manifold-supported priors, often by mapping noise onto the manifold. Pullback flow matching learns geometry for intrinsic transport but also uses a manifold-adapted source (de Kruiff et al., 2025).

Learned geometry and ambient-space transport. Metric flow matching learns ambient metrics and interpolants that favor proximity to observations (Kapuśniak et al., 2024). Meta Flow Matching instead studies dynamics on the Wasserstein space of distributions, rather than imposing a geometric constraint on individual samples (Atanackovic et al., 2025). Energy matching transports ambient noise using a learned potential (Balcerak et al., 2026). Recent works also provided flow matching variations that respect data geometry (Bamberger et al., 2026; Jiang et al., 2026; Kumar et al., 2026; Habashy & Eliasmith, 2026; Cai et al., 2026; Wu et al., 2026). Flow matching has also been investigated in constrained generation, stability, and domain topology contexts (Li et al., 2026; Guan et al., 2026; Yang et al., 2026; Wyrwal et al., 2026; Baldan et al., 2026; Tauberschmidt et al., 2026). For data on a linear subspace, Pi et al. (2026) use ambient Gaussian noise, retain an analytical normal velocity, and derive statistical convergence bounds. Such bounds complement geometric stability: finite distributional error does not ensure exact manifold membership or a prescribed trajectory-wise contraction rate, although vanishing Wasserstein error implies concentration near a manifold-supported target.

Stable learned dynamics. Contraction theory provides tools for incremental and partial stability (Lohmiller & Slotine, 1998; Wang & Slotine, 2005; Tsukamoto et al., 2021; Aminzare & Sontag, 2014; Bullo, 2026), with developments in nonlinear control, modular systems, and differential growth analysis (Lohmiller & Slotine, 2000; Slotine & Lohmiller, 2001; Slotine, 2003; Aminzare & Sontag, 2013; Nazerian et al., 2024). Structured neural dynamics impose contraction (Beik Mohammadi et al., 2024; Jaffe et al., 2024); related stability constructions appear in flow matching (Sprague et al., 2024) and LaSalle-based Riemannian policies (Ding et al., 2025).

Generative robotic policies. Diffusion Policy and Robomimic provide established settings for generative policies and offline imitation learning (Chi et al., 2025; Mandlekar et al., 2022). Generative robotic policies have used manifold-valued poses or states (Ryu et al., 2024; Chatzipantazis et al., 2025), spatial equivariance (Tie et al., 2025; Wang et al., 2025; Yang et al., 2025; Zhu et al., 2025), and projection onto feasible trajectories (Bouvier et al., 2025), while empirical work has examined how accurately diffusion policies learn kinematic constraint manifolds (Foland et al., 2025). Flow-matching policies have also studied point-cloud-conditioned action generation, spatial equivariance, and consistency training for manipulation (Chisari et al., 2025; Zhang et al., 2026; Yan et al., 2025).

MSFM learns tangential transport while prescribing normal contraction independently of tangential regression accuracy. Crucially, its prior need not lie on, or first be mapped onto, the target manifold. This permits analytical projectors for known manifolds or local PCA affine proxies for unknown expert-action geometry, without learning a global intrinsic parameterization. Under the stated projection and proxy-selection assumptions, contraction theory guarantees convergence from admissible ambient initializations to the known manifold or the modeled proxy geometry, respectively.

3 Background and Problem Formulation

We first review the conventional FM. Let p_0 be a source distribution on $\mathbb{R}^n$ and p_{data} the target distribution. Conditional flow matching samples $x_0 \sim p_0$, $x_1 \sim p_{\text{data}}$, and $t \sim \mathcal{U}[0,1]$, constructs $x_t = \gamma(t; x_0, x_1)$ with velocity $u_t = \partial_t \gamma(t; x_0, x_1)$, and minimizes

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}\left[\|v_\theta(t, x_t) - u_t\|^2\right]. \tag{1}$$

At inference, one draws $x(0) \sim p_0$ and integrates $\dot{x} = v_\theta(t, x)$. The usual Euclidean path is $x_t = (1-t)x_0 + tx_1$, with $\dot{x}_t = x_1 - x_0$. More general scalar interpolations change all ambient directions, including components locally tangent and normal to the data manifold.

3.1 Embedded-Manifold Geometry

Let $\mathcal{M} \subset \mathbb{R}^n$ be a smooth embedded manifold. Close enough to $\mathcal{M}$, in a tubular neighborhood $\mathcal{U}$, every point x has a unique nearest point on the manifold, $y = \Pi(x)$, which depends smoothly on x. We define an offset $r(x) = x - \Pi(x) = x - y$, which means we can split x into this base point and the offset, i.e., $x = y + r(x)$. At y, let P_y project onto the tangent space $T_y\mathcal{M}$ and $Q_y = I - P_y$ onto the normal space. The offset r points straight away from the manifold, so $Q_y r = r$, and its length is the distance to $\mathcal{M}$. We measure this distance with

$$V_{\mathcal{M}}(x) = \tfrac{1}{2}\operatorname{dist}(x, \mathcal{M})^2 = \tfrac{1}{2}\|r(x)\|^2, \qquad \nabla V_{\mathcal{M}}(x) = r(x). \tag{2}$$

Why conventional flow matching does not imply stability. Along a generated trajectory $\dot{x} = v_\theta(t, x)$, the distance changes as $\dot{V}_\mathcal{M} = r(x)^\top v_\theta(t, x)$: only the normal component of the velocity matters. Samples are attracted to $\mathcal{M}$ if this rate is negative everywhere off the manifold, for instance

$$\dot{V}_\mathcal{M} = r(x)^\top v_\theta(t, x) \le -\alpha(t)\|r(x)\|^2, \qquad \alpha(t) > 0, \tag{3}$$

which makes the distance shrink at rate at least $\alpha(t)$ at every point of every trajectory. The regression loss (1) only makes the velocity error small on average over sampled training paths, so it cannot enforce this pointwise condition: a model with low loss can still drift away from $\mathcal{M}$ at inference. MSFM enforces (3) by construction while learning tangential transport.

4 Manifold-Stable Flow Matching

MSFM has two ingredients: a *sampling field*, used at inference, that learns motion along the manifold and prescribes contraction toward it; and a *training path* that supplies the regression targets. The two must agree in the normal direction. We construct both for known manifolds (Section 4.1) and for unknown manifolds via local affine proxies (Section 4.2).

4.1 Known-Manifold Construction

For an unconstrained network w_θ and $y = \Pi(x)$, define

$$v_\theta(t, x) = P_y w_\theta(t, x) - \alpha(t) r(x). \tag{4}$$

The learned component is the tangent $P_y w_\theta(t, x)$; the prescribed normal feedback $-\alpha(t)r(x)$ vanishes on $\mathcal{M}$. The positive, locally integrable rate $\alpha(t)$ is specified in Section 4.4.

For source $x_0 \in \mathcal{U}$ and target $x_1 \in \mathcal{M}$, construct

$$x_t = y_t + r_t, \qquad y_t \in \mathcal{M}, \quad r_t \in N_{y_t}\mathcal{M}. \tag{5}$$

Choose a smooth base curve from $y_0 = \Pi(x_0)$ to $y_1 = x_1$, for example the geodesic $y_t = \mathrm{Exp}_{y_0}(t \,\mathrm{Log}_{y_0}(x_1))$, using the Riemannian exponential and a suitable logarithm branch. Set $r_0 = x_0 - y_0$ and $Q_t = Q_{y_t}$, and define

$$\dot{r}_t = \dot{Q}_t r_t - \alpha(t) r_t, \qquad r(0) = r_0. \tag{6}$$

Here $\dot{Q}_t = \mathrm{d}Q_{y_t}/\mathrm{d}t$ accounts for the changing normal space, while

$$\|r_t\| = q_\perp(t)\|r_0\|, \qquad q_\perp(t) = \exp\!\left(-\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right).$$

For each sampled time t, first obtain r_t by evaluating the solution of (6), either in closed form when available or by integrating this prescribed linear ODE from 0 to t. The training pair is then

$$x_t = y_t + r_t, \qquad u_t = \dot{y}_t + \dot{Q}_t r_t - \alpha(t) r_t.$$

For the $\mathrm{SO}(3)$ geometry used in our experiments, r_t has the explicit expression given in (15) (Appendix A). Assuming $\Pi(x_t) = y_t$ throughout, sampling $x_0 \sim p_0$ defines $p_t(\cdot \mid x_1)$ as the law of x_t. Lemma 1 in Appendix B proves that r_t remains normal at y_t, with $\|r_t\| = q_\perp(t)\|r_0\|$, and that the target velocity satisfies $Q_t u_t = -\alpha(t) r_t$. If $y_t \to x_1$ and $\int_0^1 \alpha(t)\,\mathrm{d}t = \infty$, the resulting path converges to x_1 as $t \to 1^-$.

4.2 Unknown Manifold: Local Affine Proxies

For samples $\mathcal{D} = \{x_i^*\}_{i=1}^N$, local PCA supplies orthonormal tangent bases T_i and affine proxies

$$P_i = T_i T_i^\top, \quad Q_i = I - P_i, \quad \widehat{\mathcal{M}}_i = x_i^* + \mathrm{range}(T_i).$$

Define $r_i(x) = Q_i(x - x_i^*)$. Select a proxy by nearest anchor or minimum normal residual, $i(x) \in \arg\min_i \frac{1}{2}\|r_i(x)\|^2$, and use

$$v_\theta(t, x) = P_{i(x)} w_\theta(t, x) - \alpha(t) r_{i(x)}(x). \tag{7}$$

This data-derived geometry is related to off-manifold residual metrics (Pan et al., 2026). Algorithm 1 gives its construction.

For a fixed proxy containing x_1, set $s_0 = P_i(x_0 - x_1)$ and $r_0 = Q_i(x_0 - x_1)$. Since $\dot{Q}_i = 0$, a compatible path is

$$x_t = x_1 + (1 - t)s_0 + q_\perp(t) r_0, \qquad u_t = -s_0 - \alpha(t) q_\perp(t) r_0. \tag{8}$$

4.3 ALGORITHMS

Training regresses the tangential velocity and records the normal residual separately. Sampling integrates the learned field from p_0, using current-state geometry.

Algorithm 1 is needed only when the manifold is unknown. Algorithms 2 and 3 apply to both known and unknown manifolds, through the following shared notation:

$$(P(x), r(x)) = \begin{cases} \big(P_{\Pi(x)},\ x - \Pi(x)\big), & \text{known manifold,} \\ \big(P_{i(x)},\ Q_{i(x)}(x - x^*_{i(x)})\big), & \text{unknown manifold,} \end{cases} \qquad Q(x) = I - P(x). \quad (9)$$

Algorithm 1 Local Proxy Preprocessing (Only when the manifold is unknown)

Require: Dataset $\mathcal{D}$, neighbor count k, dimension rule
1: **for** each anchor $x_i^* \in \mathcal{D}$ **do**
2: Form Y_i from its k nearest-neighbor differences.
3: Let T_i contain the leading d_i left singular vectors of Y_i.
4: Set $P_i = T_i T_i^\top$ and $Q_i = I - P_i$.
5: **end for**
6: **return** $\{P_i, Q_i\}_{i=1}^N$.

Algorithm 2 MSFM Training

Require: p_0, data, geometry, rate α, cutoff δ, network w_θ
1: **while** not converged **do**
2: Sample (x_0, x_1, t) and construct (x_t, u_t) using the chosen path.
3: Evaluate P, Q, and r at x_t using the sampling geometry.
4: Update θ to minimize $\|P(w_\theta(t, x_t) - u_t)\|^2$.
5: Record the normal residual $\|Qu_t + \alpha(t)r\|^2$.
6: **end while**

Algorithm 3 MSFM Sampling

Require: p_0, trained w_θ, geometry, rate α, cutoff δ
1: Draw $x(0) \sim p_0$; hold any observed condition fixed.
2: **for** ODE steps up to $t = 1 - \delta$ **do**
3: Evaluate the current-state field (4) or (7).
4: Advance the numerical solver.
5: **end for**
6: **return** $x(1 - \delta)$.

4.4 CONTRACTION SCHEDULE

We use constant plus singular contraction, with $\alpha_0 \geq 0$ and $\beta > 0$:

$$\alpha(t) = \alpha_0 + \frac{\beta}{1 - t}, \qquad 0 \leq t < 1, \quad (10)$$

and $q_\perp(t) = e^{-\alpha_0 t}(1 - t)^\beta$. Since $q_\perp(t) \to 0$, normal displacement vanishes at the terminal limit. Numerically, integrate to $1 - \delta$ in log time $\tau = -\log(1 - t)$, which makes $(1 - t)\alpha(t) = \alpha_0 e^{-\tau} + \beta$ bounded; details are in Appendix D.

Linear-path special case. For a fixed affine proxy containing x_1, $\alpha_0 = 0$ and $\beta = 1$ reduce (8) to $x_t = (1 - t)x_0 + tx_1$ and $u_t = x_1 - x_0$. This does not extend to curved-manifold paths.

Physical interpretation. The normal potential $V = \frac{1}{2}\|r_t\|^2$ satisfies $\dot{V} = -2\alpha(t)V$: tangential transport accompanies normal relaxation. For a fixed base curve and $\alpha_0 = 0$, the normal path also uniquely minimizes a time-weighted normal kinetic action. When $\beta = 1$, the weight is constant; the

full ambient path can still curve as the base point and normal space move. This is normal-motion optimality, not an optimal-transport claim for the full trajectory. Appendix B.1 defines the action and proves the result.

Since all contracting dynamical systems are dissipative (i.e., the volume of the ball of initial conditions tends to zero eventually), our MSFM is also normally dissipative: the volume of all perturbations to the initial condition in the normal direction to the manifold will tend to zero.

5 Theoretical Guarantees

Proposition 1 (Orthogonal MSFM loss decomposition). *At a fixed training pair (x_t, u_t), let P be either geometry's tangent projector in* (9)*, $Q = I - P$, and r its normal residual. For $v_\theta = Pw_\theta - \alpha r$,*

$$\|v_\theta - u_t\|^2 = \|P(w_\theta - u_t)\|^2 + \|Qu_t + \alpha(t)r\|^2. \tag{11}$$

The normal term is parameter-independent and vanishes exactly when $Qu_t = -\alpha(t)r$.

Assume $\alpha(t)$ is positive and locally integrable for $t \in [0, 1)$, with $\int_0^1 \alpha(t)\,\mathrm{d}t = \infty$ (e.g., as in (10)). Solutions must exist up to each $t < 1$; for the known-manifold field, assume local uniqueness and that the trajectory remains in the smooth projection neighborhood. We first show that, for a known manifold, the prescribed normal feedback alone keeps samples on $\mathcal{M}$ and drives off-manifold samples toward it.

Theorem 1 (Manifold invariance and convergence). *Suppose the solution of* (4) *remains in $\mathcal{U}$. If $x(0) \in \mathcal{M}$, then $x(t) \in \mathcal{M}$ for all times of existence. For arbitrary $x(0) \in \mathcal{U}$,*

$$\operatorname{dist}(x(t), \mathcal{M}) \le \exp\left(-\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right) \operatorname{dist}(x(0), \mathcal{M}). \tag{12}$$

Then $\operatorname{dist}(x(t), \mathcal{M}) \to 0$ as $t \to 1^-$.

The same mechanism applies when the manifold is unknown: using a single local PCA proxy.

Theorem 2 (Fixed-proxy convergence). *Fix i and consider $\dot{x} = P_i w_\theta(t, x) - \alpha(t) Q_i(x - x_i^*)$. Then*

$$\|Q_i(x(t) - x_i^*)\| \le \exp\left(-\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right) \|Q_i(x(0) - x_i^*)\|. \tag{13}$$

Consequently, $\|Q_i(x(t) - x_i^*)\| \to 0$ as $t \to 1^-$.

During sampling, the active proxy may switch as the state moves. Define $V_{\mathrm{proxy}}(x) = \min_i \frac{1}{2}\|Q_i(x - x_i^*)\|^2$ and $\widehat{\mathcal{M}}_{\mathrm{proxy}} = \bigcup_i \widehat{\mathcal{M}}_i$. The next result shows that such switching does not break convergence.

Theorem 3 (Convergence to the proxy set). *Under the minimum-residual selector $i(x) \in \arg\min_i \frac{1}{2}\|r_i(x)\|^2$, any absolutely continuous solution for which* (7) *is well defined satisfies*

$$V_{\mathrm{proxy}}(x(t)) \le \exp\left(-2\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right) V_{\mathrm{proxy}}(x(0)). \tag{14}$$

Consequently, $\operatorname{dist}(x(t), \widehat{\mathcal{M}}_{\mathrm{proxy}})$ obeys the corresponding square-root bound and converges to zero.

Finally, for $\alpha(t)$ in (10), these bounds take an explicit form that quantifies convergence by $t = 1$.

Corollary 1 (Exponential–polynomial terminal bound). *For* (10)*, let $\mathcal{S}$ denote either $\mathcal{M}$, or a fixed proxy $\widehat{\mathcal{M}}_i$, or the proxy union $\widehat{\mathcal{M}}_{\mathrm{proxy}}$, as appropriate. Then*

$$\operatorname{dist}(x(t), \mathcal{S}) \le e^{-\alpha_0 t}(1 - t)^\beta \operatorname{dist}(x(0), \mathcal{S}).$$

At the numerical terminal time $1 - \delta$, the bound is $\operatorname{dist}(x(1 - \delta), \mathcal{S}) \le e^{-\alpha_0(1-\delta)}\delta^\beta \operatorname{dist}(x(0), \mathcal{S})$.

Proofs are given in Appendix C. Theorems 1–3 guarantee geometric stability, which should not be confused with exact matching of the tangential data distribution. Under normal path compatibility and exact population regression, the conventional flow-matching marginalization argument applies to the resulting compatible conditional field (Lipman et al., 2023). Without compatibility, the normal residual in (11) must be treated as a modeling tradeoff rather than hidden inside the training loss.

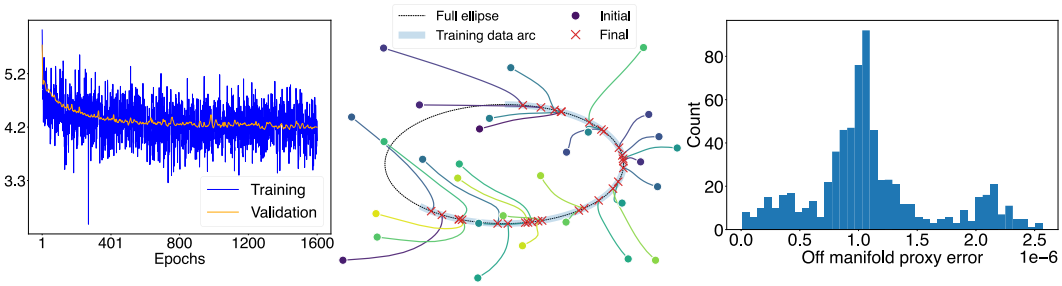


Figure 2: Ellipse: training/validation losses over 1,600 epochs, sampling trajectories, and terminal off-proxy errors for 700 samples. Here, we approximate the manifold (the ellipse) by local affine proxies using PCA.

6 Experiments

We choose three experiments that together cover both settings of MSFM: an ellipse, an unknown curved manifold whose analytic form lets us verify convergence; disturbed Push-T control, an unknown manifold of expert actions approximated by local PCA proxies; and nominal Robomimic Square manipulation, a known $\mathrm{SO}(3)$ manifold with analytical projectors. In each case, we report geometric validity separately from task success.

6.1 Ellipse With Local Linear Proxies

The target manifold is the ellipse arc $(2\cos\varphi, \sin\varphi)$, $\varphi \in [-3\pi/4, \pi/2)$, represented by one-dimensional local PCA proxies; its analytic geometry is not supplied to the model. MSFM uses a uniform ambient prior, linear training paths, and $\alpha(t) = 1/(1-t)$. Across 700 generated samples, mean terminal off-proxy error is 1.10×10^{-6} and mean analytic ellipse residual is 9.74×10^{-5} (Figure 2). These results show adherence to the learned geometry, not distributional equivalence or superiority over FM. Appendix E gives data, training, and diagnostic details.

Shared robotic-policy protocol. Within the following Push-T and Robomimic Square tasks, FM and MSFM share state observations, action representations, temporal U-Net capacity, Gaussian priors, and AdamW optimization. Both use validation-selected checkpoints, receding-horizon execution, and matched RK4 field-evaluation budgets with method-specific time grids. Our U-net has channel widths $(16, 32, 64)$ for Push-T, and $(32, 60, 124)$ for Robomimic Square.

6.2 Push-T

Push-T requires a planar pusher to align a T-shaped block with a target pose (Chi et al., 2025; Hugging Face, 2024). The target geometry is an unknown manifold of expert action horizons in $\mathbb{R}^{16}$: 8 successive 2D pusher positions. MSFM approximates this geometry by local PCA affine proxies and contracts toward the active proxy while learning tangential motion. It does not impose a manifold on the block's physical pose. Conventional FM uses unconstrained linear-path velocity regression. Both policies condition on the pusher/block state, execute the first target, and replan.

We report task performance only under repeated block-position disturbances. MSFM raises success from 74% to 82% and mean maximum reward from 0.86 to 0.92 (Figure 3). This improvement is consistent with transverse stabilization toward expert-supported action geometry, as shown in Fig. 3e as an off-manifold proxy error. Figure 3f shows a sample matched episode run in which the MSFM successfully performs the Push-T task, while the conventional FM fails. Appendix F provides additional details

6.3 Robomimic Square

A Panda robot must place a square nut over its matching peg (Mandlekar et al., 2022; Zhu et al., 2020). The known action manifold is $(\mathbb{R}^3 \times \mathrm{SO}(3) \times \mathbb{R})^{64}$: 64-step horizons of translation, rotation,

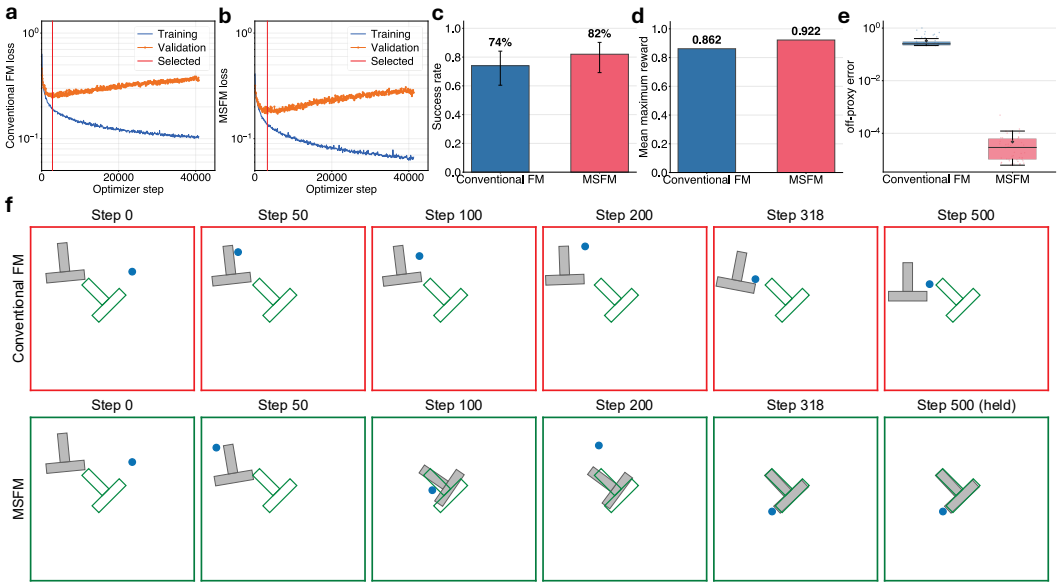


Figure 3: Push-T Task. (a,b) Total FM and parallel-only MSFM training/validation losses; red lines retain the original checkpoint selection. (c,d) Disturbed success and mean maximum reward. (e) Shows that off-manifold-proxy error is 4 orders of magnitude larger for the conventional FM. (f) A matched episode run where MSFM is successful in performing the task, while the conventional FM fails. Here, we approximate the manifold (expert actions) by local affine proxies using PCA.

and gripper commands, represented in $\mathbb{R}^{64\times 13}$. MSFM constrains the rotation blocks to approach $\mathrm{SO}(3)$ using the tubular path and analytical projectors; translation and gripper components remain unconstrained. Both methods encode absolute pose goals relative to the current end-effector frame, decode with the same final rotation projection, and execute four actions before replanning. FM uses ambient linear paths.

Using the best validation checkpoints, MSFM achieves 72% success versus 60% for FM (Figure 4). Panel d also shows that, given a budget for the length of an episode, MSFM achieves a higher success rate. Mean rotation deviation before final projection falls from 10^{-2} to 10^{-7} in panel e. Thus, stronger $\mathrm{SO}(3)$ adherence accompanies higher observed success. Appendix G provides additional information.

7 Conclusion

We introduced manifold-stable flow matching (MSFM), which separates learned tangential transport from prescribed normal contraction. The formulation accommodates both known manifolds, through analytical projectors, and unknown geometry, through local affine proxies. We derived compatible probability paths and decomposed the regression loss into a learnable tangential term and a normal compatibility residual. Under the stated geometric and solution assumptions, we established manifold invariance and explicit transverse convergence bounds for known manifolds and fixed proxies, together with a proxy-union guarantee under minimum-residual selection. With a suitable contraction schedule, these bounds ensure vanishing normal deviation as flow time approaches its terminal value, without requiring exact learning of the tangential dynamics.

Our work complements the notion of manifold adherence discussed by Pan et al. (2026): producing actions consistent with expert-supported geometry under unfamiliar observations can be more relevant to closed-loop performance than reducing validation error alone. MSFM makes attraction toward the chosen geometry an explicit property of the generative dynamics. In Push-T, this geometry is estimated from demonstrated action horizons; in Square, it includes the analytical rotation manifold $\mathrm{SO}(3)$. The stronger geometric adherence and higher observed task performance of MSFM are consistent with the importance of manifold adherence.

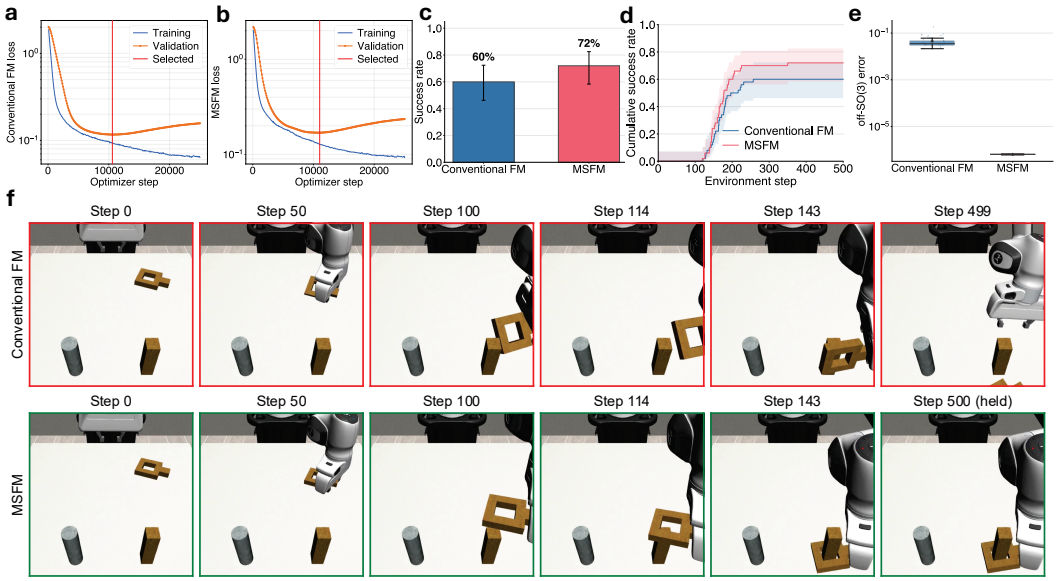


Figure 4: Robomimic Square Task. (a,b) Training/velocity-validation losses; red lines mark the checkpoint selection. (c) Success on paired nominal rollouts using best integrated-policy checkpoints. (d) Cumulative success rate vs. Robomimic environment time step. (e) Pre-projection rotation deviation, nearly five orders of magnitude lower with MSFM. (f) shows a completed episode from MSFM, while a conventional FM failed. Here, the manifold for rotation poses is SO(3).

The absolute success rates in Push-T and Square are below those reported for state-of-the-art generative policies (e.g., Chi et al., 2023). This gap reflects our experimental design rather than a limitation of MSFM. Our goal is a controlled comparison: FM and MSFM share identical temporal U-Net backbones, conditioning, training budgets, and field-evaluation counts, so that observed differences can be attributed to the manifold-stable construction rather than to architecture or tuning. To this end, we use deliberately compact networks (156K parameters for Push-T and 859K for Square), substantially smaller than the backbones typically used in diffusion-policy benchmarks. In addition, Push-T is evaluated under repeated block disturbances, a protocol that differs from standard benchmarks, so absolute numbers are not directly comparable. Because MSFM modifies only the output field through projection and prescribed normal contraction, it is agnostic to the backbone and can be combined with larger architectures and longer training. Evaluating whether its geometric advantages persist at state-of-the-art scale is left to future work.

The experiments show geometric adherence in ellipse sampling, higher observed success and reward under Push-T disturbances, and improved Square success with substantially smaller rotation deviations. Together, these results support transverse stabilization as a useful complement to learned generative transport.

AI USE STATEMENT

Generative AI assisted with language editing and final code polishing, optimizing readability and architectural consistency. We take responsibility for the final content of this work, including text, claims, and all numerical results.

REPRODUCIBILITY STATEMENT

Appendix A provides a special rotation-matrix formulation. Appendix B discusses path derivations. Appendix C gives proofs, and Appendix D provides more information on the algorithms. Appendices E, F, and G document data, optimization, checkpoint selection, and solvers for our numerical experiments.

REFERENCES

Michael Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025.

Zahra Aminzare and Eduardo D Sontag. Logarithmic lipschitz norms and diffusion-induced instability. *Nonlinear Analysis: Theory, Methods & Applications*, 83:31–49, 2013.

Zahra Aminzare and Eduardo D Sontag. Contraction methods for nonlinear systems: A brief introduction and some open problems. In *53rd IEEE Conference on Decision and Control*, pp. 3835–3847. IEEE, 2014.

Lazar Atanackovic, Xi Nicole Zhang, Brandon Amos, Mathieu Blanchette, Leo J Lee, Yoshua Bengio, Alexander Tong, and Kirill Neklyudov. Meta flow matching: Integrating vector fields on the wasserstein manifold. In *International Conference on Learning Representations*, volume 2025, pp. 94586–94610, 2025.

Michal Balcerak, Tamaz Amiranashvili, Antonio Terpin, Suprosanna Shit, Lea Bogensperger, Sebastian Kaltenbach, Petros Koumoutsakos, and Bjoern Menze. Energy matching: Unifying flow matching and energy-based models for generative modeling. *Advances in Neural Information Processing Systems*, 38:8583–8609, 2026.

Giacomo Baldan, Qiang Liu, Alberto Guardone, and Nils Thuerey. Physics vs distributions: Pareto optimal flow matching with physics constraints. In *International Conference on Learning Representations*, volume 2026, pp. 101215–101243, 2026.

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*, volume 2026, pp. 135754–135779, 2026.

Hadi Beik Mohammadi, Sø ren Hauberg, Georgios Arvanitidis, Nadia Figueroa, Gerhard Neumann, and Leonel Rozo. Neural contractive dynamical systems. In *International Conference on Learning Representations*, volume 2024, pp. 49097–49120, 2024.

Jean-Baptiste Bouvier, Kanghyun Ryu, Qiayuan Liao, Koushil Sreenath, and Negar Mehr. DDAT: Diffusion Policies Enforcing Dynamically Admissible Robot Trajectories. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025.

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

F. Bullo. *Contraction Theory for Dynamical Systems*. Kindle Direct Publishing, 1.3 edition, 2026. ISBN 979-8836646806. URL `https://fbullo.github.io/ctds`.

Jian-Feng Cai, Haixia Liu, Zhengyi Su, and Chao Wang. Improving classifier-free guidance of flow matching via manifold projection. In *Forty-third International Conference on Machine Learning*, 2026. URL `https://openreview.net/forum?id=IPp3LD6u16`.

Evangelos Chatzipantazis, Nishanth Rao, and Kostas Daniilidis. Stride: State-space riemannian diffusion for equivariant planning. In *Proceedings of the 7th Annual Learning for Dynamics\& Control Conference*, volume 283 of *Proceedings of Machine Learning Research*, pp. 1338–1352. PMLR, 04–06 Jun 2025.

Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations*, volume 2024, pp. 47922–47945, 2024.

Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10-11):1684–1704, 2025.

Eugenio Chisari, Nick Heppert, Max Argus, Tim Welschehold, Thomas Brox, and Abhinav Valada. Learning robotic manipulation policies from point clouds with conditional flow matching. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 982–993. PMLR, 06–09 Nov 2025.

Friso de Kruiff, Erik J Bekkers, Ozan Öktem, Carola-Bibiane Schönlieb, and Willem Diepeveen. Pullback flow matching on data manifolds. In *ICML 2025 Generative AI and Biology (GenBio) Workshop*, 2025.

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.

Lexi Foland, Thomas Cohn, Adam Wei, Nicholas Pfaff, Boyuan Chen, and Russ Tedrake. How well do diffusion policies learn kinematic constraint manifolds?, 2025. URL https://arxiv.org/abs/2510.01404.

Yunrui Guan, Krishna Balasubramanian, and Shiqian Ma. Mirror flow matching with heavy-tailed priors for generative modeling on convex domains. In *International Conference on Learning Representations*, volume 2026, pp. 130098–130124, 2026.

Karim Habashy and Chris Eliasmith. Geodesic flow matching for denoising high-dimensional structured representations. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=a9CuW1f0CT.

Hugging Face. gym-pusht: A gymnasium environment for push-t. https://github.com/huggingface/gym-pusht, 2024. Accessed 2026-07-08.

Sean Jaffe, Alexander Davydov, Deniz Lapsekili, Ambuj K. Singh, and Francesco Bullo. Learning neural contracting dynamics: Extended linearization and global guarantees. In *Advances in Neural Information Processing Systems*, volume 37, pp. 66204–66225. Curran Associates, Inc., 2024.

Keyue Jiang, Jiahao Cui, Xiaowen Dong, and Laura Toni. Bures-wasserstein flow matching for graph generation. In *International Conference on Learning Representations*, volume 2026, pp. 140522–140560, 2026.

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, pp. 135011–135042. Curran Associates, Inc., 2024.

Shivam Kumar, Yixin Wang, and Lizhen Lin. Flow matching is adaptive to manifold structures, 2026. URL https://arxiv.org/abs/2602.22486.

Xinpeng Li, Enming Liang, and Minghua Chen. Gauge flow matching: Efficient constrained generative modeling over general convex set and beyond. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=vxq1OnaAMq.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=PqvMRDCJT9t.

Winfried Lohmiller and Jean-Jacques E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998. doi: 10.1016/S0005-1098(98)00019-3.

Winfried Lohmiller and Jean-Jacques E Slotine. Nonlinear process control using contraction theory. *AIChE journal*, 46(3):588–596, 2000.

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. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pp. 1678–1690. PMLR, 08–11 Nov 2022.

Amirhossein Nazerian, Francesco Sorrentino, and Zahra Aminzare. Bridging the gap between reactivity, contraction, and finite-time lyapunov exponents. *arXiv preprint arXiv:2410.23435*, 2024.

Wenzhe Ouyang, Qi Ye, Jinghua Wang, Zenglin Xu, and Jiming Chen. Rfmpose: Generative category-level object pose estimation via riemannian flow matching. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 67591–67610. Curran Associates, Inc., 2025.

Chaoyi Pan, Giridharan Anantharaman, Nai-Chieh Huang, Claire Jin, Daniel Pfrommer, Chenyang Yuan, Frank Permenter, Guannan Qu, Nicholas Boffi, Guanya Shi, and Max Simchowitz. Much ado about noising: Dispelling the myths of generative robotic control. In *International Conference on Learning Representations*, volume 2026, pp. 90575–90614, 2026.

Sophia Pi, Mingcheng Lu, Jerry Yao-Chieh Hu, Maojiang Su, Weimin Wu, and Han Liu. Learning manifold data with flow matching. In *ICML 2026 Workshop on Foundations of Deep Generative Models: Understanding Memorization, Generalization, and Reasoning*, 2026.

Hyunwoo Ryu, Jiwoo Kim, Hyunseok An, Junwoo Chang, Joohwan Seo, Taehan Kim, Yubin Kim, Chaewon Hwang, Jongeun Choi, and Roberto Horowitz. Diffusion-edfs: Bi-equivariant denoising generative modeling on se(3) for visual robotic manipulation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18007–18018, 2024.

J.-J.E. Slotine and W. Lohmiller. Modularity, evolution, and the binding problem: a view from stability theory. *Neural Networks*, 14(2):137–145, 2001. ISSN 0893-6080.

Jean-Jacques E Slotine. Modular stability tools for distributed computation and control. *International Journal of Adaptive Control and Signal Processing*, 17(6):397–416, 2003.

Christopher Iliffe Sprague, Arne Elofsson, and Hossein Azizpour. Incorporating stability into flow matching. In *ICML 2024 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2024.

Jan Tauberschmidt, Sophie Fellenz, Sebastian Vollmer, and Andrew Duncan. Physics-constrained fine-tuning of flow-matching models for generation and inverse problems. In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 143992–144037, 2026.

Chenrui Tie, Yue Chen, Ruihai Wu, Boxuan Dong, Zeyi Li, Chongkai Gao, and Hao Dong. Et-seed: Efficient trajectory-level se(3) equivariant diffusion policy. In *International Conference on Learning Representations*, volume 2025, pp. 60114–60132, 2025.

Hiroyasu Tsukamoto, Soon-Jo Chung, and Jean-Jacques E. Slotine. Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview. *Annual Reviews in Control*, 52: 135–169, 2021.

Dian Wang, Stephen Hart, David Surovik, Tarik Kelestemur, Haojie Huang, Haibo Zhao, Mark Yeatman, Jiuguang Wang, Robin Walters, and Robert Platt. Equivariant diffusion policy. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 48–69. PMLR, 06–09 Nov 2025.

Wei Wang and Jean-Jacques E Slotine. On partial contraction analysis for coupled nonlinear oscillators. *Biological cybernetics*, 92(1):38–53, 2005.

Junwei Wu, Yihang Liu, Ruixuan Yu, and Jian Sun. Flow for future: Geometric SE(3)-equivariant flow matching for 3d trajectory prediction. In *Forty-third International Conference on Machine Learning*, 2026.

Kacper Wyrwal, Ismail I Ceylan, and Alexander Tong. Topological flow matching. In *International Conference on Learning Representations*, volume 2026, pp. 21075–21100, 2026.

Ge Yan, Jiyue Zhu, Yuquan Deng, Shiqi Yang, Ri-Zhao Qiu, Xuxin Cheng, Marius Memmel, Ranjay Krishna, Ankit Goyal, Xiaolong Wang, and Dieter Fox. Maniflow: A general robot manipulation policy via consistency flow training. In *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 2268–2293. PMLR, 27–30 Sep 2025.

Jeongyong Yang, Seunghwan Jang, and SooJean Han. Safeflowmatcher: Safe and fast planning using flow matching with control barrier functions. In *International Conference on Learning Representations*, volume 2026, pp. 54058–54087, 2026.

Jingyun Yang, Ziang Cao, Congyue Deng, Rika Antonova, Shuran Song, and Jeannette Bohg. Equibot: Sim(3)-equivariant diffusion policy for generalizable and data efficient learning. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 1048–1068. PMLR, 06–09 Nov 2025.

Olga Zaghen, Floor Eijkelboom, Alison Pouplin, and Erik J Bekkers. Towards variational flow matching on general geometries. In *ICLR 2025 Workshop on Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*, 2025.

Olga Zaghen, Floor Eijkelboom, Alison Pouplin, Cong Liu, Max Welling, Jan-Willem van de Meent, and Erik Bekkers. Riemannian variational flow matching for material and protein design. In *International Conference on Learning Representations*, volume 2026, pp. 34243–34286, 2026.

Qinglun Zhang, Shen Cheng, Tian Dan, Haoqiang Fan, Guanghui Liu, and Shuaicheng Liu. Efficient hybrid se(3)-equivariant visuomotor flow policy via spherical harmonics for robot manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 27989–27998, June 2026.

Zichen Zhong, Haoliang Sun, Yukun Zhao, Yongshun Gong, and Yilong Yin. Riemannian meanflow for one-step generation on manifolds. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=DeOm4Axr9W.

Xupeng Zhu, Fan Wang, Robin Walters, and Jane Shi. SE(3)-equivariant diffusion policy in spherical Fourier space. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 80187–80206. PMLR, 13–19 Jul 2025.

Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Kevin Lin, Abhiram Maddukuri, Soroush Nasiriany, and Yifeng Zhu. robosuite: A modular simulation framework and benchmark for robot learning. *arXiv preprint arXiv:2009.12293v3*, 2020.

A Rotation-Matrix Specialization

For an ambient matrix $X \in \mathbb{R}^{3\times 3}$, let $R = \Pi_{\mathrm{SO}(3)}(X)$ be the nearest proper rotation. For a matrix direction Z, the analytical projectors are

$$P_R(Z) = R\,\mathrm{skew}(R^\top Z),$$
$$Q_R(Z) = R\,\mathrm{sym}(R^\top Z).$$

Here $\mathrm{skew}(B) = (B - B^\top)/2$ and $\mathrm{sym}(B) = (B + B^\top)/2$, with orthogonality measured by the Frobenius inner product. The ambient MSFM specialization is

$$\dot{X} = P_R(W_\theta(t, X)) - \alpha(t)(X - R), \qquad R = \Pi_{\mathrm{SO}(3)}(X).$$

Here W_θ denotes the network's 3×3 matrix output. Theorem 1 applies wherever the nearest-rotation projection is smooth, and the solution remains in its domain. Square uses this field for each rotation block; the translation and gripper components are unconstrained.

Explicit conditional path. For an ambient source matrix X_0 and target rotation R_1, set $R_0 = \Pi_{\mathrm{SO}(3)}(X_0)$, $S_0 = R_0^\top(X_0 - R_0)$, and $\Omega = \log(R_0^\top R_1)$, using a chosen rotation-logarithm branch. The matrix S_0 is symmetric and Ω is skew-symmetric. Along the geodesic $R_t = R_0 e^{t\Omega}$, the solution of (6) is

$$r_t = q_\perp(t) R_0 e^{t\Omega/2} S_0 e^{t\Omega/2}, \qquad X_t = R_t + r_t. \tag{15}$$

Indeed, $R_t^\top r_t = q_\perp(t) e^{-t\Omega/2} S_0 e^{t\Omega/2}$ is symmetric, and differentiation gives $Q_{R_t}\dot{r}_t = -\alpha(t) r_t$. Differentiating $Q_{R_t} r_t = r_t$ then recovers the full projector ODE. The training velocity is $U_t = \dot{X}_t$, evaluated analytically; no auxiliary numerical ODE solve is needed for this geometry. Compatibility with the sampling field requires $\Pi_{\mathrm{SO}(3)}(X_t) = R_t$, as in the tubular-neighborhood assumptions.

B Compatible Probability Paths in Tubular Coordinates

Lemma 1 (Normality and compatibility of tubular paths)**.** *Let $y : [0, 1) \to \mathcal{M}$ be a continuously differentiable base curve whose orthogonal normal projector $Q_t = Q_{y_t}$ is continuously differentiable, and set $P_t = I - Q_t$. Let $\alpha : [0, 1) \to (0, \infty)$ be locally integrable and $r_0 \in N_{y_0}\mathcal{M}$. Then*

$$\dot{r}_t = \dot{Q}_t r_t - \alpha(t) r_t, \qquad r(0) = r_0,$$

has a unique locally absolutely continuous solution. For every $t < 1$,

$$Q_t r_t = r_t, \qquad \|r_t\| = q_\perp(t)\|r_0\|, \qquad q_\perp(t) = \exp\left(-\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right).$$

For $x_t = y_t + r_t$, its velocity $u_t = \dot{x}_t$ satisfies, almost everywhere,

$$P_t u_t = \dot{y}_t + \dot{Q}_t r_t, \tag{16}$$
$$Q_t u_t = -\alpha(t) r_t.$$

If, in addition, $x_t \in \mathcal{U}$ and $\Pi(x_t) = y_t$ for all $t < 1$, then $\mathrm{dist}(x_t, \mathcal{M}) = q_\perp(t)\|r_0\|$, and the path has the same normal velocity as the sampling field (4)*. Finally, if $y_t \to x_1 \in \mathcal{M}$ and $\int_0^1 \alpha(t)\,\mathrm{d}t = \infty$, then $x_t \to x_1$ as $t \to 1^-$.*

Proof. On every compact interval $[0, T] \subset [0, 1)$, the coefficient $\dot{Q}_t - \alpha(t) I$ is integrable. Standard existence and uniqueness for linear ODEs therefore give a unique absolutely continuous solution; uniqueness makes these solutions agree on overlapping intervals.

Differentiating $Q_t^2 = Q_t$ and multiplying the resulting identity on the left and right by Q_t gives

$$\dot{Q}_t Q_t + Q_t \dot{Q}_t = \dot{Q}_t, \qquad Q_t \dot{Q}_t Q_t = 0.$$

Define the normality defect $z_t = (I - Q_t) r_t$. Almost everywhere,

$$\dot{z}_t = -Q_t \dot{Q}_t r_t - \alpha(t) z_t = -\big(Q_t \dot{Q}_t + \alpha(t) I\big) z_t, \qquad z_0 = 0,$$

where $r_t = Q_t r_t + z_t$ and $Q_t \dot{Q}_t Q_t = 0$ were used. Uniqueness implies $z_t = 0$ for every $t < 1$, proving $Q_t r_t = r_t$. It follows that $Q_t \dot{Q}_t r_t = 0$, so $\dot{Q}_t r_t$ is tangent at y_t and is orthogonal to r_t. Hence

$$\frac{\mathrm{d}}{\mathrm{d}t}\|r_t\|^2 = -2\alpha(t)\|r_t\|^2$$

almost everywhere. Integrating yields $\|r_t\|^2 = q_\perp(t)^2\|r_0\|^2$ and thus the stated norm identity.

Differentiating $x_t = y_t + r_t$ gives $u_t = \dot{y}_t + \dot{Q}_t r_t - \alpha(t) r_t$. Since $\dot{y}_t$ and $\dot{Q}_t r_t$ are tangent while r_t is normal, applying P_t and Q_t proves the velocity identities. Under the additional projection assumption, $r(x_t) = x_t - \Pi(x_t) = r_t$ and $Q_{\Pi(x_t)} = Q_t$. Consequently,

$$\operatorname{dist}(x_t, \mathcal{M}) = \|r_t\|, \qquad Q_{\Pi(x_t)} u_t = -\alpha(t) r(x_t) = Q_{\Pi(x_t)} v_\theta(t, x_t),$$

which establishes compatibility in the normal direction, without requiring the learned tangential velocity to equal the target.

Finally, the divergent integral implies $q_\perp(t) \to 0$. Therefore

$$\|x_t - x_1\| \le \|y_t - x_1\| + q_\perp(t)\|r_0\| \to 0 \qquad \text{as } t \to 1^-.$$

□

Normality alone does not imply $\Pi(x_t) = y_t$; the projection condition remains a separate geometric assumption. The term $\dot{Q}_t r_t$ accounts for the changing normal space and generally cannot be omitted. For the schedule (10), $q_\perp(t) = e^{-\alpha_0 t}(1-t)^\beta$.

For a fixed affine proxy, $\dot{Q}_t = 0$ and $r_t = q_\perp(t) r_0$. With a general tangential schedule $q_\parallel$, the path is $x_t = x_1 + q_\parallel(t) s_0 + q_\perp(t) r_0$ and its velocity is $u_t = \dot{q}_\parallel(t) s_0 - \alpha(t) q_\perp(t) r_0$. Choosing $q_\parallel(t) = 1 - t$ yields the path used in the main text. For (10), the normal target magnitude is

$$\|Q_i u_t\| = e^{-\alpha_0 t}\big[\alpha_0(1-t) + \beta\big](1-t)^{\beta-1}\|r_0\|.$$

For nonzero r_0, this tends to zero as $t \to 1^-$ when $\beta > 1$, tends to $e^{-\alpha_0}\|r_0\|$ when $\beta = 1$, and diverges when $0 < \beta < 1$. It is constant in time in the linear-path special case $\alpha_0 = 0$, $\beta = 1$. On $0 \le t \le 1 - \delta$, a uniform upper bound is $(\alpha_0 + \beta)\max\{1, \delta^{\beta-1}\}\|r_0\|$.

B.1 Minimum Kinetic Action of the Normal Motion

Fix a smooth base path y_t on $[0, 1]$, with continuously differentiable normal projector Q_t, and an initial displacement $Q_0 r_0 = r_0$. For any absolutely continuous normal-displacement curve $\eta(t)$ satisfying $Q_t \eta(t) = \eta(t)$, define its normal velocity by $D_t^\perp \eta = Q_t \dot{\eta}$. The identity

$$\dot{\eta} = \dot{Q}_t \eta + D_t^\perp \eta$$

follows by differentiating $Q_t \eta = \eta$. The first term is tangent, while the second is normal. Thus $\|D_t^\perp \eta\|^2$ measures only normal motion, excluding the tangential velocity required by the changing normal space. This is the coordinate-free form of kinetic energy in a normal frame with no internal rotation.

Variational characterization. For a fixed smooth base path and $\alpha_0 = 0$, define the weighted normal kinetic action

$$\mathcal{E}_\beta[\eta] = \frac{1}{2}\int_0^1 (1-t)^{1-\beta}\|Q_t \dot{\eta}(t)\|^2 \,\mathrm{d}t, \tag{17}$$

on absolutely continuous normal-displacement curves $\eta(t) \in N_{y_t}\mathcal{M}$ satisfying $\eta(0) = r_0$ and $\eta(1) = 0$. The projection excludes the tangential velocity required by the changing normal space. For $\beta = 1$, the weight is one; only for fixed affine geometry does this also equal the ordinary ambient kinetic action of the displacement.

Take $\alpha_0 = 0$ and $\beta > 0$. Let r_t solve (6), and write $q_\perp(t) = (1-t)^\beta$. The vector $\xi_t = r_t / q_\perp(t)$, defined for $t < 1$, satisfies $\dot{\xi}_t = \dot{Q}_t \xi_t$, $Q_t \xi_t = \xi_t$, and $\|\xi_t\| = \|r_0\|$; it extends continuously to $t = 1$. Consequently,

$$D_t^\perp r_t = -\beta(1-t)^{\beta-1}\xi_t, \qquad \mathcal{E}_\beta[r] = \frac{\beta}{2}\|r_0\|^2.$$

Consider any other absolutely continuous normal curve η with $\eta(0) = r_0$, $\eta(1) = 0$, and finite action (17). Set $h_t = \eta(t) - r_t$, so $Q_t h_t = h_t$ and $h_0 = h_1 = 0$. With $w_\beta(t) = (1-t)^{1-\beta}$, expanding the action gives

$$\mathcal{E}_\beta[\eta] - \mathcal{E}_\beta[r] = -\beta \int_0^1 \xi_t^\top D_t^\perp h_t \,\mathrm{d}t + \frac{1}{2}\int_0^1 w_\beta(t)\|D_t^\perp h_t\|^2 \,\mathrm{d}t.$$

Because $\dot\xi_t$ is tangent and h_t is normal, $\xi_t^\top D_t^\perp h_t = \mathrm{d}(\xi_t^\top h_t)/\mathrm{d}t$. The first integral therefore vanishes, leaving

$$\mathcal{E}_\beta[\eta] = \frac{\beta}{2}\|r_0\|^2 + \frac{1}{2}\int_0^1 w_\beta(t)\|D_t^\perp h_t\|^2 \,\mathrm{d}t.$$

Equality requires $D_t^\perp h_t = 0$ almost everywhere. Then $\dot h_t = \dot Q_t h_t$, and $h_0 = 0$ implies $h_t = 0$ by uniqueness. Hence r_t is the unique minimizer. For a fixed affine proxy, Q_t is constant and $D_t^\perp \eta = \dot\eta$; the minimizer reduces to $(1-t)^\beta r_0$, recovering ordinary kinetic-action minimization when $\beta = 1$.

Relaxation and ambient curvature. For the full schedule (10), let $V(t) = \frac{1}{2}\|r_t\|^2$. If $x_t = y_t + r_t$ remains in the tubular neighborhood with $\Pi(x_t) = y_t$, this is also $\frac{1}{2}\operatorname{dist}(x_t, \mathcal{M})^2$. The normal dynamics give

$$\dot V(t) = -2\alpha(t)V(t), \qquad V(t) = e^{-2\alpha_0 t}(1-t)^{2\beta}V(0).$$

Under $\tau = -\log(1-t)$, $\mathrm{d}V/\mathrm{d}\tau = -2(\alpha_0 e^{-\tau} + \beta)V$: the normal potential decays at a bounded rate that is constant when $\alpha_0 = 0$. Meanwhile, $\dot r_t = \dot Q_t r_t - \alpha(t) r_t$ includes the change in normal direction. Together with the moving base point, this can produce a curved ambient trajectory even when $\beta = 1$. Tangential transport and normal contraction occur simultaneously, not as two successive phases.

The variational statement fixes y_t and minimizes only the weighted normal action. It neither selects the tangential transport nor establishes optimality of the full generative path. Switching between affine proxies requires a separate analysis and does not inherit a global minimum-action guarantee from this result.

Including a constant contraction term. For the full schedule (10), use $q_\perp(t) = e^{-\alpha_0 t}(1-t)^\beta$ and the action

$$\mathcal{E}_q[\eta] = \frac{1}{2}\int_0^1 \frac{\|D_t^\perp \eta(t)\|^2}{\alpha(t) q_\perp(t)} \,\mathrm{d}t.$$

The unique minimizer is again the corresponding solution r_t of (6). Indeed, $\xi_t = r_t/q_\perp(t)$ still satisfies $\dot\xi_t = \dot Q_t \xi_t$, while $D_t^\perp r_t = -\alpha(t) q_\perp(t)\xi_t$. With $h = \eta - r$, the same expansion and endpoint argument yield

$$\mathcal{E}_q[\eta] = \frac{1}{2}\|r_0\|^2 + \frac{1}{2}\int_0^1 \frac{\|D_t^\perp h_t\|^2}{\alpha(t) q_\perp(t)} \,\mathrm{d}t,$$

where $\int_0^1 \alpha(t) q_\perp(t)\,\mathrm{d}t = 1$ was used. Uniqueness follows as above. When $\alpha_0 = 0$, this action equals $\mathcal{E}_\beta/\beta$ and has the same minimizer.

C Proofs

C.1 Proof of Proposition 1

At x_t, take P, Q, and r from either geometry in (9) in Proposition 1. In both cases, P and $Q = I - P$ are complementary orthogonal projectors, with $Pr = 0$ and $Qr = r$. Decompose $u_t = Pu_t + Qu_t$ to obtain

$$\begin{aligned} v_\theta - u_t &= Pw_\theta - \alpha r - Pu_t - Qu_t \\ &= P(w_\theta - u_t) - \big(Qu_t + \alpha r\big). \end{aligned}$$

The first term lies in the tangent space and the second lies in the normal space. Their inner product is zero, so the Pythagorean identity gives (11). The normal term vanishes if and only if $Qu_t = -\alpha(t) r$ holds.

C.2 Proof of Theorem 1

If $x \in \mathcal{M}$, then $\Pi(x) = x$ and $r(x) = 0$, so

$$v_\theta(t, x) = P_x w_\theta(t, x) \in T_x\mathcal{M}.$$

Thus $\mathcal{M}$ is invariant.

For the transverse bound, use $V_\mathcal{M}$ from (2). Along (4),

$$\dot{V}_\mathcal{M} = r^\top (P_y w_\theta - \alpha r) = -\alpha(t)\|r\|^2 = -2\alpha(t)V_\mathcal{M},$$

because $r \perp T_{\Pi(x)}\mathcal{M}$. Therefore,

$$V_\mathcal{M}(x(t)) = \exp\left(-2\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right) V_\mathcal{M}(x(0)).$$

Taking square roots gives (12). Divergence of the integral at $t = 1$ implies terminal convergence.

C.3 Proof of Theorem 2

Let $r_i(t) = Q_i(x(t) - x_i^*)$. Since $Q_i P_i = 0$ and $Q_i^2 = Q_i$, we have $\dot{r}_i = Q_i\dot{x} = -\alpha(t)r_i$. Hence

$$r_i(t) = \exp\left(-\int_0^t \alpha(\tau)\,\mathrm{d}\tau\right) r_i(0),$$

which proves (13).

C.4 Proof of Theorem 3

For each proxy, define $V_i(x) = \frac{1}{2}\|Q_i(x - x_i^*)\|^2$. On a compact time interval within $[0, 1)$, the compositions $V_i(x(t))$ and their finite minimum are absolutely continuous. At almost every time, $\dot{x}$ exists and satisfies the selected field. For an active minimizing index i, $\nabla V_i^\top \dot{x} = -2\alpha(t)V_i$. The derivative of the minimum, whenever it exists, is bounded above by this active derivative. Hence $\dot{V}_{\mathrm{proxy}} \le -2\alpha(t)V_{\mathrm{proxy}}$ almost everywhere, and the integral comparison inequality yields (14). Moreover,

$$\mathrm{dist}(x, \widehat{\mathcal{M}}_{\mathrm{proxy}}) = \min_i \|Q_i(x - x_i^*)\| = \sqrt{2V_{\mathrm{proxy}}(x)},$$

which gives the distance bound.

C.5 Proof of Corollary 1

For (10), $\int_0^t \alpha(\tau)\,\mathrm{d}\tau = \alpha_0 t - \beta\log(1-t)$. Substituting $e^{-\int_0^t \alpha(\tau)\,\mathrm{d}\tau} = e^{-\alpha_0 t}(1-t)^\beta$ into (12) and (13) gives the stated bound for $\mathcal{S} = \mathcal{M}$ and $\mathcal{S} = \widehat{\mathcal{M}}_i$, respectively; the latter uses $\mathrm{dist}(x, \widehat{\mathcal{M}}_i) = \|Q_i(x - x_i^*)\|$. Taking square roots in (14) gives the same bound for $\mathcal{S} = \widehat{\mathcal{M}}_{\mathrm{proxy}}$. Setting $t = 1 - \delta$ gives the remaining case.

D Algorithm and Implementation Details

Local proxy preprocessing. Algorithm 1 forms the neighbor-difference matrix $Y_i = [x_{i,1}^* - x_i^*, \ldots, x_{i,k}^* - x_i^*]$ for each anchor. Its singular value decomposition $Y_i = U_i\Sigma_i V_i^\top$ gives $T_i = U_i[:, 1 : d_i]$. The rank d_i is fixed or selected by explained variance. The resulting $P_i = T_i T_i^\top$ and $Q_i = I - P_i$ are held fixed during network training.

Training and compatibility. Algorithm 2 draws $x_0 \sim p_0$, $x_1 \sim p_{\mathrm{data}}$, and $t \sim \mathcal{U}[0, 1 - \delta]$. The constructed (x_t, u_t) and the geometry used to evaluate the field are independent of the network parameters. Proposition 1 therefore gives identical parameter gradients for full-field regression and its tangential term, but any nonzero normal residual must still be reported. For a target-associated affine path, exact compatibility must hold with the state-selected sampling proxy, not merely the target's proxy. Observation-conditioned policies use the same conditioning and proxy selector during training and sampling.

Sampling and time transformation. Algorithm 3 integrates the learned field. With $\tau = -\log(1-t)$, use $t = 1 - e^{-\tau}$ and

$$\frac{\mathrm{d}x}{\mathrm{d}\tau} = (1-t)v_\theta(t,x), \qquad (1-t)\alpha(t) = \alpha_0 e^{-\tau} + \beta.$$

The normal-feedback coefficient lies between β and $\alpha_0 + \beta$ and is constant when $\alpha_0 = 0$. Stopping at $\tau = -\log\delta$ corresponds to $t = 1 - \delta$, with experiment-specific cutoffs given below. The scale satisfies $q_\perp(0) = 1$ and $\dot{q}_\perp = -\alpha q_\perp$, so changing the rate requires updating both path states and velocities.

Linear-path compatibility. For a fixed affine proxy containing x_1, the compatible base path is $y_t = x_1 + (1-t)s_0$. Setting $\alpha_0 = 0$ and $\beta = 1$ recovers the ambient linear conditional path. Keeping that linear path for other schedule parameters instead produces the normal residual $Qu_t + \alpha(t)r_t = [\alpha_0(1-t) + \beta - 1]r_0$. This fixed-proxy calculation does not extend automatically to changing proxies or curved manifolds.

E Additional Ellipse Details

For a generated point $z = (z_1, z_2)$, the analytic ellipse residual is $|(z_1/2)^2 + z_2^2 - 1|$. The off-proxy metric is $\|Q_{i(z)}(z - x^*_{i(z)})\|$, while nearest-data distance measures closeness to the sampled arc rather than to its extended affine proxies.

The analytic ellipse is used only to evaluate generated samples. The training data contain 800 uniformly spaced angles on $[-3\pi/4, \pi/2)$, mapped to $(2\cos\varphi, \sin\varphi)$. Each affine proxy has dimension one and uses 31 nearest neighbors; the active proxy is selected by nearest anchor, not minimum normal residual. The prior distribution is uniform on $[-3, 3] \times [-2, 2]$.

The network has two hidden layers of width 64 with tanh activations (4,546 parameters). Adam runs for 1,600 updates with learning rate 0.003, and batch size 256. Training uses linear conditional paths with $\beta = 1$. Validation uses 256 fixed source/target/time tuples drawn from the same 800 target points; it is not a held-out target-data split. The plotted full-field losses include the normal compatibility residual. Their nonzero plateau therefore cannot be interpreted as tangential regression error alone.

Sampling uses RK4 in log time to $t = 1 - 10^{-6}$, with 700 independent source draws. Mean terminal off-proxy error is 1.10×10^{-6}, median 1.03×10^{-6}, and maximum 2.57×10^{-6}. Mean nearest-data distance is 2.75×10^{-3}; the mean absolute analytic ellipse residual is 9.74×10^{-5}. The latter two metrics distinguish attraction to the union of extended affine lines from closeness to the sampled arc.

F Push-T Experimental Details

Data and representation. We use the state-based `pusht_cchi_v7_replay` demonstrations. The 206 demonstrations are split into 164 training, 21 validation, and 21 held-out test episodes. The condition is $c = (p_{a,x}, p_{a,y}, p_{b,x}, p_{b,y}, \cos\theta_b, \sin\theta_b)$, where p_a and p_b are world-frame pusher and block positions and θ_b is block orientation. The targets are eight successive absolute pusher positions, not object-frame poses or angular actions. Horizons never cross episode boundaries; 1,442 incomplete horizons are discarded. The episode-disjoint training/validation/test split contains 19,232/2,469/2,507 complete horizons. Normalization and the proxy atlas use training data only; held-out test horizons are not used for checkpoint selection.

Proxy field and training objective. For each normalized training horizon a_i and condition c_i, we select 64 neighbors using the joint squared distance $\|c - c_i\|^2 + 0.1\|x - a_i\|^2$. PCA of horizon differences defines the tangent projector P_i by retaining at least 95% of their squared singular-value energy; $Q_i = I - P_i$ is its normal complement. The median retained dimension is three. The active index $i = i(x, c)$ minimizes the same joint distance, using the current generated horizon and observed condition, without access to the target action. With flow time t, MSFM uses $v_\theta = P_i w_\theta(t, x, c) - (1-t)^{-1} Q_i(x - a_i)$. This run regresses against the linear-path velocity $u = x_1 - x_0$ at $x_t = (1-t)x_0 + tx_1$, with demonstrated horizon x_1 and $x_0 \sim \mathcal{N}(0, I_{16})$. Because the selected proxy

need not contain x_1, this is not an exactly normal-compatible probability path: its total loss includes a model-independent transverse residual. The plotted parallel loss is $\|P_i(v_\theta - u)\|^2/16$, recovered by subtracting that residual using the original seeded samples. Fixed validation samples make this subtraction constant across epochs, preserving checkpoint selection. The MSFM selected-checkpoint validation loss is approximately 0.173 after subtraction, versus 31.520 before subtraction.

Optimization and sampling. Each temporal U-Net has 156,338 trainable parameters and channel widths $(16, 32, 64)$, with a learning rate of 5×10^{-3}. Both networks use ELU activations, kernel size five, eight normalization groups, one residual block per level, one middle block, and a 32-dimensional time embedding. Push-T uses episode-disjoint data splits and AdamW with batch size 256 and weight decay 10^{-5}. Training uses gradient-norm clipping at 10 and early-stopping patience of 500 epochs within a 3,000-epoch maximum. FM and MSFM stop at epochs 537 and 543, selecting epochs 37 and 43, respectively. Validation evaluates velocity regression, not integrated-action prediction. Each rollout plan starts from standard Gaussian noise and uses 20 RK4 steps to $t = 1 - 10^{-6}$; FM integrates in t, whereas MSFM uses $s = -\log(1 - t)$. Thus, field-evaluation counts match, but integration grids differ. Generated targets are denormalized, clipped to $[0, 512]^2$, and only the first target is executed before replanning.

Evaluation and interpretation. The reported task evaluation is restricted to disturbed rollouts to probe performance under repeated external perturbations. Both methods use initial-state seeds 0–49 with a 1,000-step limit. Success means maximum environment reward of at least 0.95; the reported score is the episode maximum reward averaged over rollouts. Episodes schedule up to ten block translations before step 900, drawing displacements from $\mathcal{N}(0, 100I_2)$ in raw coordinates (standard deviation 10 per axis), clipping block position to the workspace and resetting linear/angular velocity. Success counts are 37/50 (FM) and 41/50 (MSFM), with mean maximum rewards of 0.862 and 0.922, respectively. Success bars show Wilson 95% intervals; reward bars show means. These results demonstrate higher observed success and reward for MSFM under the tested disturbance protocol.

G Robomimic Square Experimental Details

Data, conditioning, and controller. The experiment uses `square_ph_low_dim.hdf5` and `NutAssemblySquare` in Robosuite 1.4.1, with a Panda robot, 20 Hz control, and an absolute `OSC_POSE` controller. The proficient-human dataset contains 200 demonstrations, split into 180 training and 20 validation episodes, yielding 27,213 and 2,941 horizons; there is no separate offline test partition. Two successive 59-dimensional observations give a 118-dimensional condition containing object state, end-effector position/quaternion and velocities, gripper positions/velocities, and joint positions, sine/cosine encodings, and velocities. The dataset and simulator use the same observation ordering and training-derived normalization; images are not used. Demonstration targets reconstruct the intended absolute controller goals from scaled incremental commands, not from subsequently achieved poses. For current end-effector pose (p_a, R_a) and future world-frame goal (p_h, R_h), each token is $(R_a^\top(p_h - p_a)/\sigma_p, \mathrm{vec}(R_a^\top R_h), g_h)$, where g_h is the gripper command and $\sigma_p = 0.106773$ m is the training-derived translation scale. The anchor remains fixed within each plan. Horizons are padded at episode ends, with padded tokens excluded from losses.

Geometry and probability paths. Both methods learn in the same ambient space $\mathbb{R}^{64\times 13}$, with target geometry $(\mathbb{R}^3 \times \mathrm{SO}(3) \times \mathbb{R})^{64}$. FM uses linear interpolation from standard Gaussian noise. MSFM uses the tubular path $M_t = R_t + r_t$, where R_t is the geodesic from the projected source matrix to the demonstrated rotation and r_t is the explicit solution in (15), with $q_\perp(t) = e^{-\alpha_0 t}(1 - t)^\beta$. The saved run uses $\alpha_0 = 0$ and $\beta = 1$; translation and gripper paths remain linear. For a current matrix M, let $R = \Pi_{\mathrm{SO}(3)}(M)$ be its nearest proper rotation. The analytical tangent projector acts on a matrix direction Z as $P_R(Z) = R\,\mathrm{skew}(R^\top Z)$, and the rotational field is $P_R(w_\theta) - \alpha(t)(M - R)$ with $\alpha(t) = 1/(1-t)$. Training regresses the learned tangent component and reports the analytical normal residual separately. Losses mask padded tokens and use equal position/rotation/gripper weights, with rotational matrix errors weighted by the canonical factor $1/2$. The tubular compatibility statement applies within the smooth projection neighborhood; ambient Gaussian initialization does not place intermediate states exactly on $\mathrm{SO}(3)$.

Optimization and checkpoint selection. Each temporal U-Net has 859,347 trainable parameters and channel widths $(32, 60, 124)$. Both networks use SiLU activations, kernel size three, four normalization groups, two residual blocks per level, two middle blocks, a 78-dimensional time embedding, and a 74-dimensional conditioning embedding. Square task uses episode-disjoint data splits and batch size 256. Training uses identical initial parameters, AdamW, a 1,000-step learning-rate warmup followed by cosine decay from 3×10^{-4} to 3×10^{-5}, gradient clipping at 10, and exponential-moving-average weights with decay 0.999. The budget is 25,000 updates (234 epochs, with the last epoch partial). Rollout checkpoints minimize the integrated validation score over the first four actions, using 256 fixed validation examples and four Gaussian draws per example. This score averages squared normalized position error, squared rotation geodesic error after projection, and squared clipped-gripper error, divided by seven. The velocity-loss minima marked in Figure 4(a,b) occur at steps 10,593 and 11,021.

Sampling and paired evaluation. Both policies use RK4. FM integrates uniformly to $t = 1$; MSFM integrates to $1 - 10^{-10}$. Raw generated matrices are composed with the fixed observation anchor; the decoder applies the same hand-to-legacy-controller rotation correction and one final $\mathrm{SO}(3)$ projection before converting to absolute axis-angle commands. The first four of 64 generated actions are executed before replanning. Evaluation uses the same saved bank of 50 initial simulator states, Gaussian source, and a 500-step limit, with no injected disturbances. Success is the simulator's nut-assembly success flag.

Results and scope. FM succeeds on 30/50 episodes and MSFM on 36/50; 25 pairs succeed with both, five only with FM, and eleven only with MSFM. Geometric deviation is $\|M - \Pi_{\mathrm{SO}(3)}(M)\|_F$, averaged over the generated horizon and replanning times within each episode, then over episodes, before execution projection. Its mean is 0.0451 for FM and 6.40×10^{-7} for MSFM, with lower MSFM deviation in all 50 pairs. The residual measures rotational validity, not accuracy relative to the expert action.