HyRef-SB: Hyperbolic Schrödinger Bridges Guided by Geometric References

Anonymous authors
Paper under double-blind review

Abstract

Endpoint distributions specify where stochastic transport starts and ends, while reference dynamics provide a geometric inductive bias for the resulting stochastic paths between them. We introduce **Hy**perbolic **Ref**erence-guided **S**chrödinger **B**ridges (HyRef-SB), combining reference diffusion, stochastic control, and endpoint distribution constraints on a hyperbolic manifold. An invertible calibration preserves the input representation, while a prescribed base metric defines diffusion and control costs. An endpoint-preserving path teacher calibrates the reference drift; differentiable rollouts then jointly minimize intrinsic residual control energy and a soft endpoint penalty under the fixed reference. A shared backbone and conditional adapters model multiple terminal distributions. Our analysis characterizes the radial and transverse costs induced by calibration and quantifies, for regular teacher marginals, how reference regression error and density regularity bound reference marginal error and a feasible residual-control budget. Continuous and finite-step entropy identities identify the regularized path laws, and a population penalty limit connects soft endpoint fitting to hard marginal constraints. Experiments yield low-collision navigation and recover tree-structure statistics, while improving cellular distribution matching over the evaluated baselines on most terminal and held-out intermediate metrics. Hyperbolic geometry reduces terminal errors relative to scale-matched Euclidean controls on three cellular tasks, and the Trametinib ablation shows further gains from the learned reference over a Brownian reference. The complete code will be made publicly available.

1 Introduction

Generating stochastic evolution paths from observed distributions is a common problem in navigation, structured generation, and scientific modeling. Navigation connects starting regions to goals through obstacles; graph generation connects distributions of structures; and cellular data describe population changes across states or perturbation conditions. In each setting, endpoint matching leaves the intermediate paths underdetermined, and many paths connecting the same endpoints violate physical constraints. Learning meaningful trajectories therefore requires specifying the geometry that organizes the data and the reference dynamics that guide paths through physically plausible regions.

Schrödinger bridges (SBs) select paths by minimizing their relative entropy with respect to a reference process (Léonard, 2014). The state space, diffusion, and baseline drift jointly define this reference. Geometry determines distances, stochastic perturbations, and the cost of control deviations, providing a common basis for modeling the dynamics. For example, negative curvature amplifies transverse separation at a fixed radial depth, introducing a directional bias suited to hierarchical expansion (Lee, 2018; Nickel & Kiela, 2017). The hierarchical, tree-like, and branching states of cellular differentiation are naturally represented in hyperbolic geometry (Klimovskaia et al., 2020). Learning *geometric guidance* from data can therefore help paths avoid obstacles and regions with little observational support while retaining the structure of the data manifold.

As illustrated in Figure 1, HyRef-SB follows the principle of defining the reference through hyperbolic geometry and constraining control through path relative entropy. A full-dimensional invertible calibration maps observations to a product manifold and preserves the input information across geometric comparisons. An endpoint-preserving path teacher incorporates data support into a path prior, which is converted into a frozen reference drift through velocity regression. In the final stage, we roll out

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.

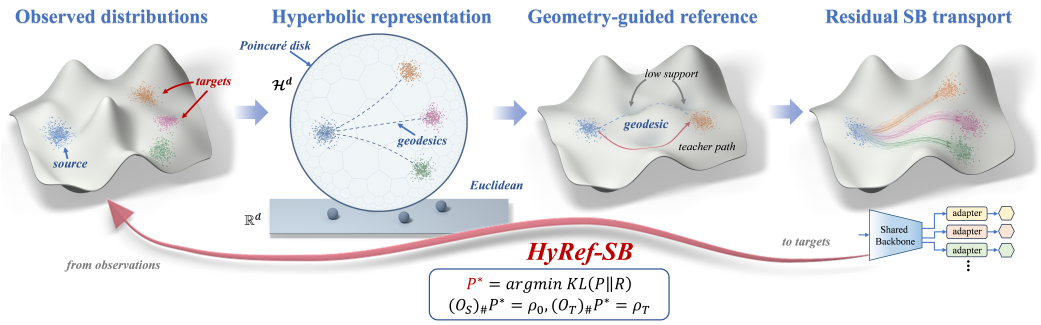

Figure 1. Geometric calibration and a support-aware teacher define the frozen reference; a shared backbone and conditional adapters provide multi-target residual SB control.

stochastic dynamics from the initial distribution, combining the fixed reference drift with a learnable residual control, and jointly optimize terminal distribution matching and control energy. For multiple terminal conditions, a shared control backbone captures common dynamics and conditional adapters express condition-specific corrections. Separate optimization weights allow difficult or underrepresented conditions to receive additional attention, while a fixed terminal prior governs generation.

Our work has three main contributions. First, we develop a hyperbolic SB architecture that connects geometric calibration, reference construction, and residual control, with a common base metric defining noise, control energy, and stochastic updates. Second, we characterize the pullback geometry and derive a teacher-to-control bound that links reference regression error and teacher density regularity to a feasible path-entropy budget. We apply the classical Girsanov identity to labeled geometric controls, establish the corresponding finite-step identity and safeguard bound, and connect population soft endpoint problems to hard marginal constraints through a penalty limit. Third, experiments on navigation, trees, and cellular data evaluate geometric feasibility, structural matching, and conditional prediction. Representation-dynamics comparisons and module ablations examine the contributions of geometry and the reference.

2 Related Work

Bridges and references. Diffusion Schrödinger Bridge (DSB) uses forward–backward diffusion, and Diffusion Schrödinger Bridge Matching (DSBM) uses iterative Markovian fitting (De Bortoli et al., 2021; Shi et al., 2023). Simulation-free score and flow matching (SF^2M) combines conditional flow and score matching (Tong et al., 2024). GSBM adds state costs (Liu et al., 2024), while reference refinement alternates bridge and reference updates across observed times (Shen et al., 2025). Twisted Schrödinger Bridge Matching derives matching objectives from Feynman–Kac tilted references (Noble et al., 2026); Smooth Schrödinger Bridges use smooth Gaussian reference processes to constrain temporal regularity across observed marginals (Hong et al., 2025). We freeze a teacher-fitted, support-aware reference before control, linking its quality to residual energy.

Manifold transport. Riemannian Flow Matching and Riemannian Diffusion Schrödinger Bridge (RDSB) learn intrinsic flows and stochastic bridges (Chen & Lipman, 2024; Thornton et al., 2022). Our navigation and Tree baseline RDSB uses forward–backward learning with geodesic random walks. Hyperbolic embeddings and mixed-curvature products encode hierarchies (Nickel & Kiela, 2017; Gu et al., 2019), including single-cell structure in Poincaré Maps (Klimovskaia et al., 2020). Metric Flow Matching and Riemannian Metric Matching relate data-dependent metrics to low-energy paths and local geometry (Kapuśniak et al., 2024; Bamberger et al., 2026). Our calibrated product metric consistently defines reference diffusion, transport updates, and control costs.

Cellular transport. FreeBridge uses empirical support costs (Wang et al., 2026); Departures models continuous–discrete perturbation responses (Chi et al., 2026), with Departures-PC denoting our continuous-PC adaptation. BranchSBM uses branch-specific velocities and growth fields (Tang et al., 2026), while scDiffEq learns drift and diffusion for fate prediction (Vinyard et al., 2025). Related unbalanced models include DeepRUOT for net growth and CytoBridge for population interactions (Zhang et al., 2025a;b). Our fixed generation prior and conditional adapters express heterogeneous responses under a shared geometric reference. Appendix A provides a more detailed review.

3 PRELIMINARIES

On a Riemannian manifold $(\mathcal{M}, g)$, let $T_z\mathcal{M}$ denote the tangent space at z, and let $\|\cdot\|_g$, ∇_g, and Δ_g denote the intrinsic norm, gradient, and Laplace–Beltrami operator. Given a reference drift b_0, an additional control u, and noise intensity $\sigma > 0$, the controlled diffusion is characterized by the generator

$$\mathcal{L}_t^u f = \langle b_0 + u, \nabla_g f \rangle_g + \frac{\sigma^2}{2} \Delta_g f. \tag{1}$$

Setting $u = 0$ gives the reference law R; also setting $b_0 = 0$ gives scaled manifold Brownian motion.

Let $P_t = (\omega \mapsto \omega_t)_{\#} P$ be the time-t marginal of a continuous path law P. Given a reference probability law R and endpoints ν_0, ν_1, the SB minimizes path relative entropy (Léonard, 2014):

$$P^{\mathrm{SB}} \in \operatorname*{arg\,min}_{P:\, P_0=\nu_0,\, P_1=\nu_1} \mathrm{KL}(P\|R), \qquad \mathrm{KL}(P\|R) = \mathbb{E}_P \log \frac{dP}{dR}.$$

Kullback–Leibler (KL) divergence is infinite unless $P \ll R$. Under the Girsanov assumptions of Theorem 5.2, SB admits a minimum-energy control formulation (Léonard, 2012). We use a finite-weight endpoint penalty; Section 5 establishes its population limit.

The Poincaré ball models hyperbolic space, a manifold of constant negative curvature (Lee, 2018). For $c > 0$, its domain, Riemannian metric, and position-dependent conformal factor are

$$\mathbb{B}_c^d = \{z \in \mathbb{R}^d : c\|z\|^2 < 1\}, \qquad g_z^H = \lambda_c(z)^2 I_d, \qquad \lambda_c(z) = \frac{2}{1 - c\|z\|^2}.$$

For $d \geq 2$, curvature is $-c$ and the boundary is infinitely distant. The exponential map Exp_z sends tangent vectors to geodesic endpoints, with inverse Log_z; parallel transport PT preserves inner products. Here $g_0^H = 4I$ and $d_H(0, \mathrm{Exp}_0^H(v)) = 2\|v\|$. Product geometry adds the squared factor norms and applies the component exponential maps factorwise (Appendix B).

4 HYREF-SB: GEOMETRIC REFERENCES AND CONDITIONAL TRANSPORT

4.1 GEOMETRIC REPRESENTATION CALIBRATION

Given source μ_0, targets $\{\mu_1^k\}_{k=1}^K$, and fixed prior $\pi_k > 0$, $\sum_k \pi_k = 1$, we seek conditional path laws and their mixture. An invertible calibration preserves all preprocessed input dimensions. Terminal directions enter the hyperbolic factor; other coordinates enter the Euclidean factor with weight $w_E > 0$:

$$\mathcal{M} = \mathbb{B}_c^{d_H} \times \mathbb{R}^{d_E}, \qquad d_H + d_E = D, \qquad g_z = \mathrm{diag}\big(\lambda_c(z_H)^2 I_{d_H}, w_E I_{d_E}\big).$$

Training endpoints determine center m, scale $s > 0$, complete orthogonal basis Q, and radius coefficient $\alpha_{\mathrm{cal}} > 0$. With $v = (\alpha_{\mathrm{cal}}/s)Q^\top(x - m)$, the encoder and inverse decoder are

$$E(x) = \big(\mathrm{Exp}_0^H(v_H), v_E\big), \qquad D(z) = m + \frac{s}{\alpha_{\mathrm{cal}}} Q\big(\mathrm{Log}_0^H(z_H), z_E\big). \tag{2}$$

The pullback $g_{\mathrm{cal}} = E^* g$ determines input-space distance, noise, and control scales. Lemma B.2 establishes unclipped invertibility; Proposition 5.1 characterizes transverse amplification. Appendix B gives calibration and safeguards, and Section 6.2 describes the Tree representation.

4.2 CALIBRATING REFERENCE DYNAMICS ON THE MANIFOLD

Let Γ_k be the endpoint-pair distribution from minibatch optimal permutations under squared intrinsic distance, and let $\chi_k > 0$, $\sum_k \chi_k = 1$, be fixed teacher-training sampling probabilities. For $(z_0, z_1) \sim \Gamma_k$, Equation (20) defines an endpoint-preserving correction $\psi_\eta^k(t) = \psi_\eta^k(t; z_0, z_1)$ of the product geodesic. The teacher objective balances intrinsic kinetic energy with support cost:

$$\mathcal{L}_{\mathrm{teach}}(\eta) = \sum_k \chi_k \mathbb{E}_{\Gamma_k} \int_0^1 \left[\frac{\lambda_{\mathrm{T,kin}}}{2} \|\dot{\psi}_\eta^k(t)\|_g^2 + \lambda_{\mathrm{T,sup}} V_{\mathcal{U}}(t, \psi_\eta^k(t), k) \right] dt, \tag{3}$$

with $\lambda_{\mathrm{T,kin}} > 0$ regulating intrinsic motion and $\lambda_{\mathrm{T,sup}} \geq 0$ weighting training-endpoint support (Appendix C.12). Velocity regression fits a state-dependent drift to the teacher:

$$\mathcal{L}_{\mathrm{warm}}(\omega) = \sum_k \chi_k \mathbb{E}_{\Gamma_k,\, T\sim\mathrm{Unif}[0,1]} \big\| b_\omega^k(T, \psi_\eta^k(T)) - \dot{\psi}_\eta^k(T) \big\|_g^2. \tag{4}$$

The teacher's conditional mean derivative minimizes this loss (Theorem C.16). We freeze b_{ref} and start from zero residual control (Lemma C.2). Proposition 5.3 bounds reference marginal error and feasible residual-control cost; Algorithm 2 specifies the complete calibration and reference-fitting procedure.

4.3 CONDITIONAL CONTROL AND ENDPOINT CONSTRAINTS

An origin-tangent backbone, multiple terminal prototypes, a shared conditional residual, and adapters capture conditional heterogeneity; parallel transport gives $u_\theta^k(t,z)$ (Appendix C.13). The initial joint law preserves the source marginal and the prescribed condition prior (Appendix C.6).

Fixing the representation and reference, we set $b_0 = b_{\mathrm{ref}}^k$ in Equation (1) and learn u_θ^k. Decoded terminal error for $\rho_{1,\theta}^k = D_\# P_1^k$ combines squared maximum mean discrepancy (MMD^2) and debiased Sinkhorn divergence (Gretton et al., 2012; Feydy et al., 2019):

$$e_k = \sum_{j=1}^{2} \frac{\lambda_j}{s_j} d_j(\rho_{1,\theta}^k, \mu_1^k), \qquad d_1 = \mathrm{MMD}^2, \quad d_2 = S_{\varepsilon,p=1}.$$

Positive scales s_j are calibrated once on training data and fixed. With $\beta > 0$, $\lambda_V \geq 0$, and $0 \leq \alpha \leq 1$, the mean–worst risk emphasizes difficult conditions without changing the generation prior:

$$\mathcal{J}(\theta) = \mathcal{R}_\alpha(e) + \beta\mathcal{K}(\theta) + \lambda_V \mathcal{V}(\theta), \qquad \mathcal{R}_\alpha(e) = (1-\alpha)\sum_k \pi_k e_k + \alpha \max_k e_k, \tag{5}$$

$$\mathcal{K} = \sum_k \pi_k \mathbb{E}_{P^k} \int_0^1 \frac{\|u_\theta^k\|_g^2}{2\sigma^2}\, dt, \qquad \mathcal{V} = \sum_k \pi_k \mathbb{E}_{P^k} \int_0^1 V_{\mathcal{U}}(t, Z_t, k)\, dt.$$

Path costs and generation use π, and Theorem 5.2 identifies $\mathcal{K}$ as reference-relative entropy. For LARRY only, $\mathcal{J}_{\mathrm{LARRY}} = \mathcal{J} + 0.5\mathcal{L}_{\mathrm{fate}}/s_{\mathrm{fate}}$ adds supervision from training-clone terminal descendants; Appendix F.8 gives the loss and fixed training-only scale.

For stochastic exploration under the same metric, choose $F(z)F(z)^\top = g_z^{-1}$ and independent standard Gaussian vectors ξ_n. With a uniform step size $h = 1/N$, the controlled exponential update is

$$Z_{n+1} = \mathrm{Exp}_{Z_n}\left[h\{b_{\mathrm{ref}}^k + u_\theta^k\}(t_n, Z_n) + \sigma\sqrt{h}\, F(Z_n)\xi_n\right]. \tag{6}$$

The frames are $\lambda_c^{-1} I$ and $w_E^{-1/2} I$. Backpropagation learns controls along the path (Algorithm 1); Appendix C.14 establishes local generator consistency and specifies the shared numerical safeguards.

5 THEORETICAL ANALYSIS: GEOMETRY, REFERENCE, AND CONTROL

The continuous identity applies Girsanov theory (Léonard, 2012; 2014), with a finite-step counterpart in Proposition C.19. Proofs appear in Appendices B.3, C.3, C.11, and C.8.

Proposition 5.1 (Radial and transverse geometry). *For the unclipped map (2), let $a = \|v_H\|$ and, for $a > 0$, let $P_\parallel, P_\perp$ be radial and transverse projections. Set $\kappa_c(a) = \sinh(2\sqrt{c}\,a)/(2\sqrt{c}\,a)$, continuously extended by $\kappa_c(0) = 1$. The resulting metric on the full input representation is*

$$E^* g = (\alpha_{\mathrm{cal}}/s)^2 Q\, \mathrm{diag}\big(4(P_\parallel + \kappa_c(a)^2 P_\perp), w_E I\big)\, Q^\top. \tag{7}$$

Two points at geodesic radius r with angular separation $\vartheta \in [0,\pi]$ have intrinsic distance

$$d_H = \frac{2}{\sqrt{c}} \operatorname{arsinh}\big[\sinh(\sqrt{c}r)\sin(\vartheta/2)\big] \geq 2r\sin(\vartheta/2). \tag{8}$$

Transverse amplification grows with radius while the radial coefficient stays constant, favoring hierarchical separation and motivating origin-matched Euclidean ablations. Path KL is invariant under invertible coordinate changes (Corollary B.3); geometry acts through reference dynamics and costs. Section 6.4 tests these effects through representation–dynamics and scale-matched comparisons.

For the fixed reference $\bar{R}(k, d\omega) = \pi_k R^k(d\omega)$ and correction $u = b - b_{\mathrm{ref}}$, labeled control satisfies:

Theorem 5.2 (Girsanov identity for labeled geometric controls). *Suppose P^k, R^k satisfy Assumption C.3, share initial law ν_0^k and nondegenerate diffusion $\sigma^2\Delta_g/2$, and have drift difference u^k under the canonical path filtration. For the labeled controlled law $\bar{P}(k, d\omega) = \pi_k P^k(d\omega)$,*

$$\mathrm{KL}(\bar{P}\|\bar{R}) = \sum_k \pi_k\, \mathrm{KL}(P^k\|R^k) = \sum_k \pi_k \mathbb{E}_{P^k} \int_0^1 \frac{\|u_t^k\|_g^2}{2\sigma^2} dt = \mathcal{K}.$$

Marginalizing labels bounds mixture entropy (Corollary C.6); state costs induce initial-law-preserving reference tilts (Corollary C.7). Proposition C.10 characterizes endpoint-factorized reference changes that preserve the hard-marginal optimizer and their contribution to the soft objective.

For ideal exponential Gaussian chains (6), Proposition C.19 gives $\mathrm{KL}(P_h \| R_h) = \mathcal{K}_h$, where $\mathcal{K}_h$ is the discrete control energy. Applying common numerical safeguards yields the corresponding entropy upper bound, with residual energy evaluated along the safeguarded controlled chain under its path law.

Let $\varrho_t^k = p_t^k \, d\mathrm{vol}_g$ be the teacher marginals and b_*^k their conditional mean velocity (Theorem C.16). The reference regression error and the teacher's time-integrated Fisher information are

$$\varepsilon_{\mathrm{reg}} = \sum_k \pi_k \int_0^1 \mathbb{E}_{\varrho_t^k} \| b_{\mathrm{ref}}^k - b_*^k \|_g^2 dt, \qquad \mathcal{I}_{\mathrm{T}} = \sum_k \pi_k \int_0^1 \mathbb{E}_{\varrho_t^k} \| \nabla_g \log p_t^k \|_g^2 dt.$$

Write $\mathcal{V}_{\mathrm{T}} = \sum_k \pi_k \int_0^1 \mathbb{E}_{\varrho_t^k} V_{\mathcal{U}}(t, z, k) dt$ for the teacher support cost.

Proposition 5.3 (Reference calibration and a feasible control budget). *Under Assumption C.17, a law $P^\dagger \in \mathcal{P}_0$ has conditional marginals ϱ_t^k and satisfies both entropy inequalities:*

$$\mathrm{KL}(P^\dagger \| \bar{R}) \le B_\sigma, \qquad \sup_{t \in [0,1]} \sum_k \pi_k \, \mathrm{KL}(\varrho_t^k \| R_t^k) \le B_\sigma,$$
$$B_\sigma := \frac{1}{2\sigma^2} \left(\sqrt{\varepsilon_{\mathrm{reg}}} + \frac{\sigma^2}{2} \sqrt{\mathcal{I}_{\mathrm{T}}} \right)^2. \tag{9}$$

Let $\mathcal{H} = \{P \in \mathcal{P}_0 : D_\# P_1^k = \mu_1^k \ \forall k\}$. If $D_\# \varrho_1^k = \mu_1^k$ for every k and $\mathcal{V}_{\mathrm{T}} < \infty$, then $\inf_{P \in \mathcal{H}} F(P) \le F(P^\dagger) \le C_{\mathrm{T}} := \beta B_\sigma + \lambda_V \mathcal{V}_{\mathrm{T}}$.

The bound separates regression error from the diffusion-induced density correction. Its comparison law is analytical; training learns the reference drift and residual control.

With geometry, reference, prior, and F fixed, consider $J_\eta(P) = F(P) + \eta \mathcal{D}(P)$ for $P \in \mathcal{P}_0$. Equation (5) has population counterpart $\eta = 1$, $\mathcal{D} = \mathcal{R}_\alpha(e)$; varying η connects soft and hard constraints.

Theorem 5.4 (Penalty limit of the ideal soft endpoint problem). *On the path-measure space of Definition C.1, assume $I_V \ge 0$ is lower semicontinuous, and $\mathcal{D} \ge 0$ is weakly lower semicontinuous and zero exactly when all target marginals match. If a hard-feasible $P^\dagger$ has $F(P^\dagger) < \infty$, then $F + \eta\mathcal{D}$ attains its minimum for every $\eta > 0$. For $\eta_n \to \infty$, $\epsilon_n \ge 0$, $\epsilon_n \to 0$, and approximate minimizers P_n with objective gap at most ϵ_n, the endpoint discrepancy satisfies*

$$\mathcal{D}(P_n) \le \frac{F(P^\dagger) + \epsilon_n}{\eta_n}.$$

The sequence is weakly relatively compact and converges to the unique hard-constrained minimizer.

Proposition 5.3 replaces $F(P^\dagger)$ by C_{T}. Squared MMD gives rate $O(\eta_n^{-1/2})$ in MMD; Corollary C.14 separates model-class, optimization, and objective-estimation errors.

6 Experiments

We evaluate our method on navigation, tree generation, and cellular distribution transport. All main benchmarks use random seed 804 and task-specific metrics, with settings in Appendix F. Supplementary material includes evolution videos for navigation and all cellular experiments.

6.1 Obstacle navigation: geometric feasibility

On the two-dimensional Poincaré ball $\mathbb{B}_1^2$, we transport a source distribution to three terminal regions arranged around two hyperbolic ball obstacles. Evaluation uses 128 integration steps and 16 Monte Carlo paths. RDSB serves as the baseline and fits endpoint marginals relative to manifold Brownian motion through alternating time reversal (Thornton et al., 2022).

Table 1. Navigation safety and path length.

Method	Safe (%) ↑	Collision (‰) ↓	Length
RDSB	63.476	365.240	3.413
HyRef-SB	**99.980**	**0.198**	**2.951**

We evaluate safety, collision rate, and path length in this hyperbolic obstacle setting. Safety is determined by collisions on the discrete time grid, while path length is computed using intrinsic distance. Appendix F.5.3 details the evaluation metrics, and Appendix F.3 specifies the obstacle configuration.

Both methods separate terminal populations, while HyRef-SB substantially reduces collisions, reflecting obstacle-aware motion (Figure 2 and Table 1). Joint gains in safety and intrinsic path length indicate efficient obstacle avoidance and support goal-directed transport through feasible regions under the learned reference.

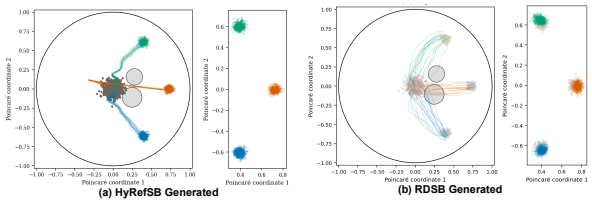

Figure 2. Navigation paths and terminal populations: HyRef-SB (left pair), RDSB (right pair).

Both methods incorporate obstacle costs during training (Appendix F.4). These results show that HyRef-SB's learned reference motion, path-risk optimization, and conditional corrections jointly guide transport more effectively in this navigation task.

6.2 Tree Generation: Geometric Transport

The Tree task starts from 16 source nodes and adds 28 nodes for three target tree families, with all nodes represented within 48 masked slots. We again use RDSB as the baseline. Hyperbolic vector quantization encodes nodes, conditional transport acts on their masked product space, and a branch-conditioned decoder reconstructs the target-specific adjacency.

Table 2. Tree generation; MMD^2 in 10^{-2}.

Method	Degree ↓	Spectral ↓	Edge F1 ↑
RDSB	1.227	0.187	0.419
HyRef-SB	**0.050**	**0.003**	**0.903**

HyRef-SB reduces degree and spectral discrepancies and improves paired Edge F1 (Table 2). Figure 6 shows closer recovery of a long backbone, multiple branching centers, and a dominant hub for the respective growth families. These observations support target-specific structure beyond the decoder's guaranteed tree validity (Proposition C.22). Agreement across degree, spectrum, and paired edges suggests coordinated local branching and global organization with target-specific node relationships.

6.3 Cellular Distribution Transport

We evaluate drug response and developmental transport with FreeBridge, Departures-PC, and BranchSBM (Wang et al., 2026; Chi et al., 2026; Tang et al., 2026), using common endpoints, features, conditions, and splits. Selection uses endpoint validation; intermediate observations are held out. LARRY methods, including scDiffEq (Vinyard et al., 2025), share training-clone associations, descendant labels, and auxiliary fate supervision (Appendix F.4).

Table 3. Drug-response errors (↓). MMD^2 uses 10^{-2} units; bold marks column bests.

	Clonidine			Trametinib		
Method	MMD^2	W_1	W_2	MMD^2	W_1	W_2
FreeBridge	3.957	1.308	1.469	7.156	1.390	1.659
Departures-PC	8.956	1.724	1.948	7.670	1.329	1.654
BranchSBM	2.854	0.955	1.024	**3.134**	1.003	1.747
HyRef-SB	**2.410**	**0.728**	**0.864**	4.290	**0.711**	**0.854**

On both drugs, HyRef-SB gives the lowest Wasserstein errors; kernel discrepancy favors HyRef-SB on Clonidine and BranchSBM on Trametinib (Table 3). Condition-directed paths retain terminal spread (Figure 3), supporting improved placement with metric-dependent shape recovery (Figure 7).

Mouse Hematopoiesis favors HyRef-SB in all terminal and most intermediate metrics, except intermediate W_2 (Table 9). The curved lower branch in Figure 3 reflects nonlinear geometry in the cellular population. HyRef-SB better preserves this structure than the compared methods, consistent with the role of hyperbolic geometry in organizing branching states and of the learned support-aware reference in guiding stochastic transport along the intrinsic geometry. On Veres, HyRef-SB retains the lowest Wasserstein errors at every reported time; competing methods lead

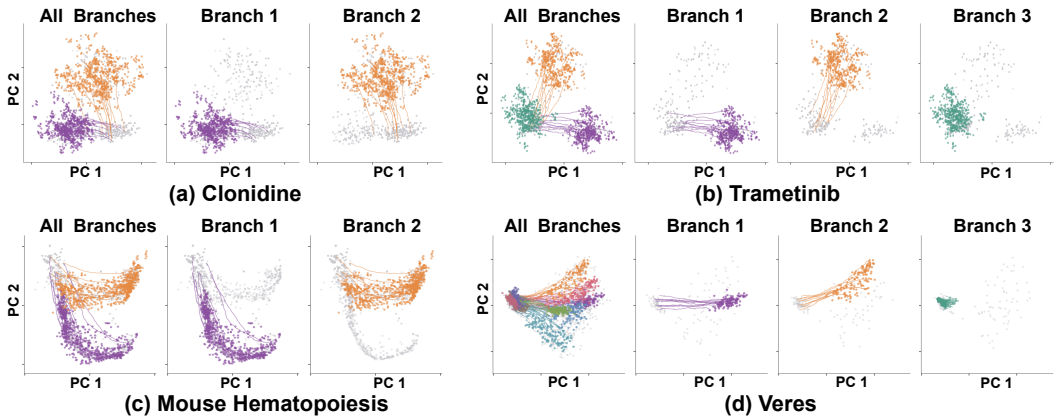

Figure 3. Basic cellular paths in PC 1–PC 2: drug responses (top), Mouse Hematopoiesis and Veres (bottom). Gray denotes observations; colored paths show all or individual terminal branches.

some intermediate kernel discrepancies (Table 10). Progressive branch separation supports temporal placement, with complementary metrics resolving differences in distributional shape.

Table 4. LARRY and Chicken distribution errors (↓); MMD^2 in 10^{-2}. Intermediate metrics average held-out times; bold marks the lowest errors within each task and metric.

	LARRY						Chicken Cardiogenesis					
	Terminal			Intermediate mean			Terminal			Intermediate mean		
Method	MMD^2	W_1	W_2	MMD^2	W_1	W_2	MMD^2	W_1	W_2	MMD^2	W_1	W_2
FreeBridge	1.947	0.863	12.230	2.349	0.541	10.144	1.925	0.571	16.747	4.782	1.019	20.783
Departures-PC	2.345	0.895	12.177	2.461	0.564	10.273	2.743	0.668	14.612	3.952	0.871	18.248
BranchSBM	**1.109**	0.999	7.402	1.280	0.546	8.141	1.008	1.578	13.964	5.245	3.080	20.162
HyRef-SB	1.165	**0.762**	**5.898**	**0.091**	**0.139**	**7.624**	**0.441**	**0.256**	**9.258**	**0.381**	**0.334**	**12.914**

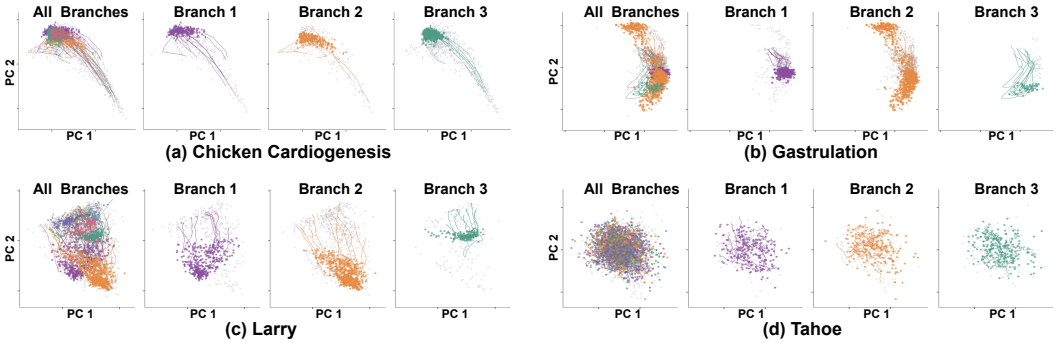

Figure 4. Advanced cellular paths in PC 1–PC 2: Chicken and Gastrulation (top), LARRY and dose-ordered Tahoe (bottom). Gray denotes observations; colored paths show all or selected branches.

A frozen training-terminal KDE reads fates from generated endpoints (Equation (25)), averaging source cells within each test clone into broad fate classes. The matched clone-fate objective uses training-terminal descendants (Appendix F.8). Higher accuracy than both the source-independent prior and scDiffEq (Table 5) supports source-specific descendant information beyond population composition, complementing the distribution-level comparisons.

Table 5. LARRY accuracy.

Method	Accuracy (%) ↑
Prior	54.098
scDiffEq	63.934
HyRef-SB	**80.328**

Chicken Cardiogenesis and Gastrulation favor HyRef-SB in every terminal and mean intermediate metric (Tables 4 and 11). Curved contraction in Chicken and expansion into separated Gastrulation populations (Figures 4 and 8) support recovery of contrasting developmental geometries. LARRY favors HyRef-SB across intermediate metrics and terminal Wasserstein errors; BranchSBM retains lower terminal MMD^2. Its elongated paths and conditional spread connect these gains to intervening population evolution. Since intermediate observations are excluded

from model selection, their recovery indicates that the learned path prior transfers endpoint supervision into useful temporal structure across distinct developmental geometries.

Tahoe favors HyRef-SB across intermediate metrics and terminal MMD^2, while competitors retain lower terminal Wasserstein errors (Table 11). Overlapping aggregates with distinct conditional populations (Figure 4) illustrate heterogeneous responses, supporting effective interpolation across held-out doses and condition-dependent population responses.

6.4 Ablation studies

LARRY and Gastrulation compare representation–dynamics pairs $H \to H$, $H \to E$, and $E \to E$ with matched splits and budgets. H denotes hyperbolic–Euclidean product geometry and E Euclidean geometry; $H \to E$ retains hyperbolic calibration and decoding (Appendix F.3).

$H \to E$ improves every metric over $E \to E$, and $H \to H$ improves most further (Table 6). LARRY terminal W_1 and Gastrulation intermediate MMD^2 favor $H \to E$. Hyperbolic representation thus supplies useful organization, while matched intrinsic dynamics add metric-dependent gains.

Table 6. Representation–dynamics ablation (↓); MMD^2 in 10^{-2}, $H \to H$: full model.

	LARRY				Gastrulation			
	Terminal		Intermediate		Terminal		Intermediate	
Setting	MMD^2	W_1	MMD^2	W_1	MMD^2	W_1	MMD^2	W_1
$H \to H$	**1.165**	0.762	**0.091**	**0.139**	**0.735**	**0.639**	0.709	**0.409**
$H \to E$	1.265	**0.706**	0.133	0.171	0.805	0.694	**0.650**	0.490
$E \to E$	1.592	0.833	0.162	0.191	0.946	0.782	0.801	0.541

Table 7. Scale-matched geometry ablation (↓). MMD^2 values are in units of 10^{-2}.

		Terminal			Intermediate	
Task	Geometry	MMD^2	W_1	W_2	W_1	W_2
Clonidine	Hyperbolic	**2.410**	**0.728**	**0.864**	–	–
	Euclidean	3.613	0.848	1.032	–	–
Trametinib	Hyperbolic	**4.290**	**0.711**	**0.854**	–	–
	Euclidean	5.260	0.848	0.979	–	–
Mouse Hematopoiesis	Hyperbolic	**1.084**	**0.150**	**0.196**	**0.387**	0.476
	Euclidean	2.847	0.229	0.279	0.394	**0.475**

With Euclidean controls matched in dimension, scale, and origin metric, hyperbolic geometry improves all terminal metrics on drug and Mouse tasks (Table 7). Mouse intermediate W_1 improves and W_2 slightly worsens. These results support hyperbolic models of branching populations.

The terminal gains are consistent with the transverse amplification in Proposition 5.1, which increases transverse separation away from the chart center while keeping the radial coefficient fixed. This mechanism can help organize branching populations under comparable local scales. The intermediate rankings further motivate assessing endpoint recovery and temporal placement together, since they probe the same dynamics at observed endpoints and held-out intermediate times.

With Trametinib geometry fixed, the learned reference outperforms both Brownian motion and a decaying teacher anchor (Table 8), supporting a motion prior retained during residual optimization. Single prototypes worsen every metric, linear adapters underperform nonlinear corrections, and removing adapters gives the largest errors. These trends suggest complementary roles for terminal prototypes and nonlinear adapters: prototypes organize distinct target directions, while adapters provide expressive corrections for condition-dependent population structure. Together, the ablations support contributions from curvature, reference calibration, and conditional flexibility.

Table 8. Trametinib modules; MMD^2 in 10^{-2}.

Setting	MMD^2	W_1	W_2
HyRef-SB	**4.290**	**0.711**	**0.854**
Brownian reference	4.600	0.792	1.015
Decaying teacher anchor	4.890	0.891	0.898
Single terminal prototype	4.790	0.860	1.084
Linear adapter	4.680	0.891	1.147
No adapter	7.072	0.998	1.172

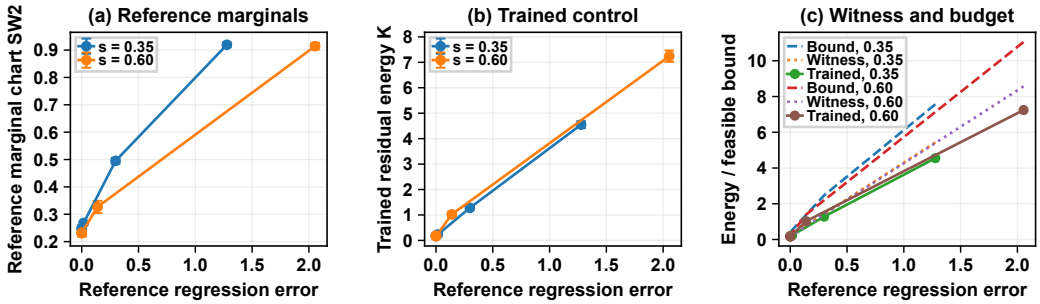

Figure 5. Reference calibration. Left to right: marginal chart SW_2, trained control energy, and budget/witness comparison versus regression error. Error bars show five-seed standard deviations.

6.5 Reference calibration and control energy

A five-seed $\mathbb{H}^2$ diagnostic fixes teacher and diffusion within each of two widths. Controls selected for low energy under a common endpoint tolerance undergo independent testing (Appendix D.4). Figure 5 links lower regression error to smaller reference-marginal discrepancy and lower residual energy. Most energy reduction occurs during early calibration; later checkpoints reach a similar low-energy range at both widths, with modest variation relative to seed variability. The late plateau suggests diminishing returns from additional regression updates once reference marginals stabilize. Jointly tracking marginal discrepancy and control energy therefore helps identify the stage at which the learned reference supplies effective guidance. Independent finer-grid retesting preserves feasibility and this trend for frozen controls, supporting reduced correction effort.

The feasible budget and analytical witness approach a positive Fisher contribution, reflecting the density correction that maintains the teacher marginal flow under diffusion. At later checkpoints, mean trained energy lies below the witness: meeting an endpoint tolerance allows additional freedom in the intermediate marginals. Proposition 5.3 bounds the full-flow witness, and the diagnostic illustrates how a better calibrated reference reduces the correction required for endpoint fitting. Chart SW_2 measures distributional discrepancy in the standardized evaluation coordinates; the theoretical budget bounds marginal KL and the witness's path entropy.

7 Discussion

The results connect three aspects of transport quality: feasible paths, distributional placement, and conditional specificity. Navigation and tree generation show that geometric reference guidance can recover task structure. Cellular interpolation and clone-fate prediction provide complementary evidence about population evolution and source-dependent outcomes. Differences between kernel and Wasserstein rankings highlight complementary aspects of distributional recovery. The ablations associate improvements with hyperbolic representation, intrinsic dynamics, and a learned motion prior; the calibration diagnostic links reference quality to reduced residual effort at both teacher widths. Denser lineage measurements and repeated benchmark training offer useful next steps for assessing biological path recovery and training variability.

8 Conclusion

We introduced HyRef-SB, a geometric SB framework that couples invertible calibration, support-aware reference learning, and conditional residual control, with a common metric governing diffusion, transport updates, and control energy. The analysis characterizes curvature-dependent transport costs and links reference accuracy and teacher density regularity to a feasible control budget, with continuous and finite-step entropy identities and a population endpoint-penalty limit connecting soft fitting to hard marginal constraints. Experiments support improved feasibility, structural recovery, and distributional transport across the evaluated tasks; ablations support complementary contributions from geometry, reference calibration, and conditional modules. Together, these results show how geometric reference design can guide stochastic transport when endpoint observations alone leave the intervening dynamics and trajectories underdetermined.

AI USE STATEMENT

AI tools assisted with language editing, LaTeX formatting, and plotting code used to redraw the reference diagnostic from recorded results. All methodological decisions, mathematical arguments, experimental designs, results, and cited material were determined, verified, and validated by the authors. The authors take full responsibility for the content of this paper.

REPRODUCIBILITY STATEMENT

The appendices provide model assumptions and proofs, stagewise algorithms, task and data settings, training configurations, and metric definitions. Ablations use the same metric conventions as the main experiments and distinguish base geometry, the reference process, conditional representation, and endpoint objectives. Benchmark results use seed 804, and the reference-calibration diagnostic uses five seeds. Repeated sampling is distinguished from independent training. Appendices F.4 and F.6 describe retrained baselines, the common comparison protocol, and reporting conventions.

ETHICS STATEMENT

This work studies distribution transport. Cellular endpoint predictions and clonal fate readouts are evaluated as computational models of observed populations. Clinical use and causal interpretation would require dedicated validation. Reuse of the data is subject to the licenses and privacy requirements of the original datasets, and the conclusions concern the data and evaluation settings specified here.

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. URL https://jmlr.org/papers/v26/23-1605.html.

Jacob Bamberger, Adam Gosztolai, Pierre Vandergheynst, Michael Bronstein, and Iolo Jones. Riemannian metric matching for scalable geometric modeling of distributions. In *International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2606.14334.

Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The Schrödinger bridge between Gaussian measures has a closed form. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 5802–5833. PMLR, 2023. URL https://proceedings.mlr.press/v206/bunne23a.html.

Ricky T. Q. Chen and Yaron Lipman. Flow matching on general geometries. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/d1f9936d3be6997ffffab692977eebe6-Abstract-Conference.html.

Changxi Chi, Yufei Huang, Jun Xia, Jiangbin Zheng, Yunfan Liu, Zelin Zang, and Stan Z. Li. Departures: Distributional transport for single-cell perturbation prediction with neural Schrödinger bridges. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(25):20544–20552, 2026. doi: 10.1609/aaai.v40i25.39190. URL https://ojs.aaai.org/index.php/AAAI/article/view/39190.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, volume 34, pp. 17695–17709, 2021. URL https://papers.nips.cc/paper_files/paper/2021/hash/940392f5f32a7ade1cc201767cf83e31-Abstract.html.

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, pp. 2406–2422, 2022. doi: 10.52202/068431-0175. URL https://arxiv.org/abs/2202.02763.

Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouvé, and Gabriel Peyré. Interpolating between optimal transport and MMD using Sinkhorn divergences. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2681–2690. PMLR, 2019. URL https://proceedings.mlr.press/v89/feydy19a.html.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25): 723–773, 2012. URL https://jmlr.org/papers/v13/gretton12a.html.

Albert Gu, Frederic Sala, Beliz Gunel, and Christopher Ré. Learning mixed-curvature representations in products of model spaces. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/pdf?id=HJxeWnCcF7.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830. URL https://doi.org/10.1080/01621459.1963.10500830.

Wanli Hong, Yuliang Shi, and Jonathan Niles-Weed. Trajectory inference with smooth Schrödinger bridges. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 23722–23750. PMLR, 2025. URL https://proceedings.mlr.press/v267/hong25f.html.

Elton P. Hsu. *Stochastic Analysis on Manifolds*, volume 38 of *Graduate Studies in Mathematics*. American Mathematical Society, 2002. URL https://www.ams.org/books/gsm/038/.

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, 2024. doi: 10.52202/079017-4291. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/f381114cf5aba4e45552869863deaaa7-Abstract-Conference.html.

Anna Klimovskaia, David Lopez-Paz, Léon Bottou, and Maximilian Nickel. Poincaré maps for analyzing complex hierarchies in single-cell data. *Nature Communications*, 11:2966, 2020. doi: 10.1038/s41467-020-16822-4. URL https://www.nature.com/articles/s41467-020-16822-4.

Hugo Lavenant, Stephen Zhang, Young-Heon Kim, and Geoffrey Schiebinger. Toward a mathematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A): 428–500, 2024. doi: 10.1214/23-AAP1969. URL https://doi.org/10.1214/23-AAP1969.

John M. Lee. *Introduction to Riemannian Manifolds*, volume 176 of *Graduate Texts in Mathematics*. Springer, second edition, 2018. doi: 10.1007/978-3-319-91755-9. URL https://sites.math.washington.edu/~lee/Books/RM/.

Christian Léonard. Girsanov theory under a finite entropy condition. In *Séminaire de Probabilités XLIV*, volume 2046 of *Lecture Notes in Mathematics*, pp. 429–465. Springer, 2012. doi: 10.1007/978-3-642-27461-9_20. URL https://arxiv.org/abs/1101.3958.

Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574, 2014. doi: 10.3934/dcds.2014.34.1533. URL https://www.aimsciences.org/article/doi/10.3934/dcds.2014.34.1533.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A. Theodorou, and Ricky T. Q. Chen. Generalized Schrödinger bridge matching. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2310.02233.

Madhav Mantri, Gaetano J. Scuderi, Roozbeh Abedini-Nassab, Michael F. Z. Wang, David McKellar, Hao Shi, Benjamin Grodner, Jonathan T. Butcher, and Iwijn De Vlaminck. Spatiotemporal single-cell RNA sequencing of developing chicken hearts identifies interplay between cellular differentiation and morphogenesis. *Nature Communications*, 12(1):1771, 2021. doi: 10.1038/s41467-021-21892-z. URL https://www.nature.com/articles/s41467-021-21892-z.

Mattia Mosso, Jaemoo Choi, and Heng Yang. Hard-constrained sampling on embedded Riemannian manifolds via adjoint Schrödinger bridges. arXiv preprint arXiv:2608.25838, 2026. URL https://arxiv.org/abs/2608.25838.

Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://arxiv.org/abs/1705.08039.

Maxence Noble, Marie Scheid, Yazid Janati, Eric Moulines, and Alain Durmus. Twisted Schrödinger bridge matching. arXiv preprint arXiv:2607.16987, 2026. URL https://arxiv.org/abs/2607.16987.

Blanca Pijuan-Sala, Jonathan A. Griffiths, Carolina Guibentif, Tom W. Hiscock, Wajid Jawaid, Fernando J. Calero-Nieto, Carla Mulas, Ximena Ibarra-Soria, Richard C. V. Tyser, Debbie Lee Lian Ho, Wolf Reik, Shankar Srinivas, Benjamin D. Simons, Jennifer Nichols, John C. Marioni, and Berthold Göttgens. A single-cell molecular map of mouse gastrulation and early organogenesis. *Nature*, 566:490–495, 2019. doi: 10.1038/s41586-019-0933-9. URL https://www.nature.com/articles/s41586-019-0933-9.

Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal Schrödinger bridges with iterative reference refinement. In *Proceedings of The 28th International Conference on*

Artificial Intelligence and Statistics, volume 258 of *Proceedings of Machine Learning Research*, pp. 3817–3825. PMLR, 2025. URL https://proceedings.mlr.press/v258/shen25b.html.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger bridge matching. In *Advances in Neural Information Processing Systems*, volume 36, pp. 62183–62223, 2023. doi: 10.52202/075280-2717. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/c428adf74782c2092d254329b6b02482-Abstract-Conference.html.

Sophia Tang, Yinuo Zhang, Alexander Tong, and Pranam Chatterjee. Branched Schrödinger bridge matching. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2506.09007.

James Thornton, Michael Hutchinson, Emile Mathieu, Valentin De Bortoli, Yee Whye Teh, and Arnaud Doucet. Riemannian diffusion Schrödinger bridge. In *ICML Workshop on Continuous Time Methods for Machine Learning*, 2022. URL https://arxiv.org/abs/2207.03024.

Alexander Y. Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 1279–1287. PMLR, 2024. URL https://proceedings.mlr.press/v238/tong24a.html.

Adrian Veres, Aubrey L. Faust, Henry L. Bushnell, Elise N. Engquist, Jennifer Hyoje-Ryu Kenty, George Harb, Yeh-Chuin Poh, Elad Sintov, Mads Gürtler, Felicia W. Pagliuca, Quinn P. Peterson, and Douglas A. Melton. Charting cellular identity during human in vitro β-cell differentiation. *Nature*, 569(7756):368–373, 2019. doi: 10.1038/s41586-019-1168-5. URL https://www.nature.com/articles/s41586-019-1168-5.

Michael E. Vinyard, Anders W. Rasmussen, Ruitong Li, Allon M. Klein, Gad Getz, and Luca Pinello. Learning cell dynamics with neural differential equations. *Nature Machine Intelligence*, 7(12):1969–1984, 2025. doi: 10.1038/s42256-025-01150-3. URL https://www.nature.com/articles/s42256-025-01150-3.

Xurui Wang, Qin Ren, Jun Ma, Haibin Ling, and Chenyu You. FreeBridge: Variational Schrödinger bridges for cellular transition dynamics. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2026*, volume 16891 of *Lecture Notes in Computer Science*. Springer Nature Switzerland, 2026. URL https://papers.miccai.org/miccai-2026/0394-Paper2873.html.

Caleb Weinreb, Samuel Wolock, Betsabeh K. Tusi, Merav Socolovsky, and Allon M. Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proceedings of the National Academy of Sciences*, 115(10):E2467–E2476, 2018. doi: 10.1073/pnas.1714723115. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC5878004/.

Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D. Camargo, and Allon M. Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479):eaaw3381, 2020. doi: 10.1126/science.aaw3381. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC7608074/.

Jesse Zhang, Airol A. Ubas, Valentine Svensson, Nicole Thomas, Vishvak Subramanyam, Christopher Carpenter, Aidan Winters, Richard de Borja, Noam Teyssier, Neha Thakar, Umair Khan, Pooja V. Akella, Harper J. Green, John D. Thompson, Vuong Tran, Joseph Pangallo, Efthymia Papalexi, Ajay Sapre, Hoai Nguyen, Oliver Sanderson, Maria Nigos, Olivia Kaplan, Sarah Schroeder, Bryan Hariadi, Simone Marrujo, Crina Curca, Alec Salvino, Guillermo Gallareta Olivares, Ryan Koehler, Alexander Rosenberg, Charles Roco, Chiara Ricci-Tam, Matthew G. Jones, Alexander Dobin, Brian S. Plosky, Danielle E. Lyons, Daniele Merico, Nima Alidoust, Hani Goodarzi, and Johnny Yu. Tahoe-100M: Mapping drug-induced molecular phenotypes at single-cell resolution. *Cell*, 189(19):5945–5961.e8, 2026. doi: 10.1016/j.cell.2026.08.035. URL https://www.sciencedirect.com/science/article/pii/S0092867426010081.

Zhenyi Zhang, Tiejun Li, and Peijie Zhou. Learning stochastic dynamics from snapshots through regularized unbalanced optimal transport. In *International Conference on Learning Representations*, 2025a. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/32b8a612105de5c22db337b774ce7b61-Abstract-Conference.html.

Zhenyi Zhang, Zihan Wang, Yuhao Sun, Tiejun Li, and Peijie Zhou. Modeling cell dynamics and interactions with unbalanced mean field Schrödinger bridge. In *Advances in Neural Information Processing Systems*, volume 38, pp. 139518–139558, 2025b. doi: 10.52202/085713-4659. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/cbf552bd72c0dd301605d3f620fe0c3a-Abstract-Conference.html.

APPENDIX

The appendix reviews related work and scope (Appendix A), develops the calibration geometry and theoretical guarantees (Appendices B–C), and reports supplementary experiments (Appendix D). Algorithms and reproducible evaluation protocols follow in Appendices E–F.

A RELATED WORK AND SCOPE

This section expands the literature review across path-space formulations, geometric transport, reference design, and cellular applications. It also outlines the theoretical levels used below: continuous path laws, ideal finite-step transitions, and safeguarded numerical updates.

A.1 SCHRÖDINGER BRIDGES AND REFERENCE PROCESSES

Léonard's survey connects endpoint-constrained path entropy to entropy-regularized transport, while his Girsanov analysis identifies the control-energy interpretation of drift changes (Léonard, 2014; 2012). These formulations underpin several computational approaches. Diffusion Schrödinger Bridge (DSB) implements iterative proportional fitting through learned forward and backward diffusions (De Bortoli et al., 2021); Diffusion Schrödinger Bridge Matching (DSBM) uses iterative Markovian fitting and bridge matching (Shi et al., 2023). Simulation-free score and flow matching (SF^2M) combines conditional flow and score objectives to learn stochastic transport without training-time simulation (Tong et al., 2024). For Gaussian marginals, Bunne et al. (2023) derive closed-form bridges, providing an analytical account of the distributional evolution.

A.2 STOCHASTIC BRIDGES ON MANIFOLDS AND CONDITIONAL TRANSPORT

Manifold methods incorporate the state-space geometry into generative dynamics. Riemannian Score-Based Generative Modelling develops manifold diffusion and time reversal for score-based generation (De Bortoli et al., 2022); Riemannian Flow Matching learns intrinsic vector fields from conditional paths (Chen & Lipman, 2024). Riemannian Diffusion Schrödinger Bridge (RDSB) fits endpoint distributions through forward–backward learning with geodesic random walks (Thornton et al., 2022). Conditional structure raises a complementary modeling issue: populations can diverge toward several outcomes. Branched Schrödinger Bridge Matching (BranchSBM) addresses this through branch-specific velocity and growth fields, together with data-dependent state costs (Tang et al., 2026).

A.3 STATE POTENTIALS, SMOOTH REFERENCES, AND CONSTRAINTS

A complementary line of work shapes intermediate paths through costs or reference design. Generalized Schrödinger Bridge Matching (GSBM) incorporates state costs (Liu et al., 2024), while Twisted Schrödinger Bridge Matching (TSBM) derives matching objectives for references obtained through Feynman–Kac tilting (Noble et al., 2026). Smooth Schrödinger Bridges impose temporal regularity through smooth Gaussian reference processes and solve tractable multi-marginal cases by lifting them to phase space (Hong et al., 2025). For hard state constraints, Mosso et al. (2026) formulate adjoint bridge sampling intrinsically on embedded manifolds. Together, these works show how reference laws and feasible state spaces encode additional information about the desired trajectories.

A.4 CELLULAR DYNAMICS, EMPIRICAL SUPPORT, AND UNBALANCED TRANSPORT

Cellular applications require models of heterogeneous responses and developmental trajectories. Departures uses minibatch optimal-transport pairing to jointly train bridges for discrete gene activation and continuous expression, predicting perturbation responses from unpaired populations (Chi et al., 2026). For temporal dynamics and fate prediction, scDiffEq learns state-dependent drift and diffusion (Vinyard et al., 2025). FreeBridge addresses endpoint-only cellular transport with empirical-support regularization in a fixed single-cell representation (Wang et al., 2026).

Unbalanced transport further accounts for changing population mass. DeepRUOT learns state transitions and net growth from sparse temporal snapshots through regularized unbalanced optimal transport (Zhang et al., 2025a). CytoBridge introduces an unbalanced mean-field bridge that additionally models cellular interactions (Zhang et al., 2025b). These approaches capture growth and interaction

mechanisms alongside state transitions. Our focus is a geometric reference that incorporates data support, with one hyperbolic–Euclidean metric defining diffusion and residual control energy.

B Geometric Calibration and Analysis of Negative Curvature

We establish the algebraic structure, invertibility, and pullback metric of the full-dimensional calibration. The first two properties ensure that all input coordinates are retained, while the pullback metric describes how calibration changes the cost of motion. The final path-entropy invariance result connects the coordinate transformation to the control interpretation in Appendix C.3.

B.1 Notation and basic geometry

Let $\mathbb{B}_c^d = \{z \in \mathbb{R}^d : c\|z\|^2 < 1\}$, $c > 0$, and $g_z^H = \lambda_c(z)^2 I$, where $\lambda_c(z) = 2/(1 - c\|z\|^2)$. For $d \geq 2$, the sectional curvature is $-c$; $d = 1$ gives a line model with the same metric expression. The exponential and logarithmic maps at the origin are

$$\mathrm{Exp}_0^H(v) = \frac{\tanh(\sqrt{c}\|v\|)}{\sqrt{c}\|v\|} v, \qquad \mathrm{Log}_0^H(z) = \frac{\operatorname{artanh}(\sqrt{c}\|z\|)}{\sqrt{c}\|z\|} z. \tag{10}$$

Both coefficients are defined to be 1 at the origin by continuity. Since $g_0^H = 4I$,

$$\|v\|_{g_0} = 2\|v\|, \qquad d_H(0, \mathrm{Exp}_0^H(v)) = 2\|v\|.$$

On the product manifold $\mathcal{M} = \mathbb{B}_c^{d_H} \times \mathbb{R}^{d_E}$ used in the main text, exponential maps, logarithmic maps, and parallel transport are computed factorwise, and

$$d_g(z, y)^2 = d_H(z_H, y_H)^2 + w_E\|z_E - y_E\|^2.$$

This product is complete and simply connected, with nonpositive sectional curvature. By the Cartan–Hadamard theorem, each $\mathrm{Exp}_z : T_z\mathcal{M} \to \mathcal{M}$ is a global diffeomorphism (Lee, 2018). We also use this property below to bound the teacher radius and calculate discrete relative entropy.

B.2 Calibration from training endpoints and a complete coordinate basis

Let $\{x_i\}_{i=1}^n$ be the training endpoints. Using their overall center m and a global root mean square (RMS) scale $s > 0$, define $\tilde{x} = (x - m)/s$. In standardized coordinates, let a_0 be the source mean and a_k the conditional terminal means. Set $d_k = a_k - a_0$ and

$$S = \frac{1}{K}\sum_{k=1}^{K} d_k d_k^\top + \rho C, \qquad \rho = 10^{-3},$$

where C is the training-endpoint covariance. When branch alignment is enabled, a full eigendecomposition gives $Q \in \mathbb{R}^{D\times D}$. Eigenvalues are sorted in decreasing order, and each eigenvector is oriented so that its component with the largest absolute value is positive. The first d_H directions enter the hyperbolic factor, and the remaining directions enter the Euclidean factor. With branch alignment disabled, we use $Q = I$. Within each eigenspace with a repeated eigenvalue, we use and cache the orthonormal basis returned by the numerical decomposition.

Lemma B.1 (Calibration-matrix decomposition and subspace objective)**.** *Let $\bar{a} = K^{-1}\sum_k a_k$. The displacement matrix then has the decomposition*

$$\frac{1}{K}\sum_k (a_k - a_0)(a_k - a_0)^\top = \frac{1}{K}\sum_k (a_k - \bar{a})(a_k - \bar{a})^\top + (\bar{a} - a_0)(\bar{a} - a_0)^\top.$$

If Q_H has the first d_H orthonormal eigenvectors of S, it maximizes $\mathrm{tr}(U^\top S U)$ over $U^\top U = I_{d_H}$.

Proof. Write $\delta_k = a_k - \bar{a}$ and $v = \bar{a} - a_0$. Expanding the displacement products gives the stated decomposition because $\sum_k \delta_k = 0$ cancels both cross terms. For $S = V\,\mathrm{diag}(\lambda_1, \ldots, \lambda_D)V^\top$ with decreasing eigenvalues, let $b_i = \|(V^\top U)_{i,:}\|^2$. Orthogonality gives $0 \leq b_i \leq 1$ and $\sum_i b_i = d_H$, hence

$$\mathrm{tr}(U^\top S U) = \sum_i \lambda_i b_i \leq \sum_{i=1}^{d_H} \lambda_i.$$

The bound follows by placing the available weight on the largest eigenvalues and is attained by $U = Q_H$. If $d_H = D$, every orthonormal basis attains $\operatorname{tr} S$. □

Thus, S combines the common displacement $vv^\top$, the dispersion of terminal means, and the global covariance. Subspace selection maximizes their total projected strength, while the complete matrix Q retains the remaining coordinates.

Let $q_{.95}$ and $q_{\max}$ denote the training 95th percentile and maximum of the hyperbolic-coordinate norms after rotation and before scaling. In the nondegenerate case, set

$$\alpha_{\text{cal}} = \min\left\{\frac{\operatorname{artanh}(r_{.95})}{\sqrt{c}\, q_{.95}}, \frac{\operatorname{artanh}(r_{\max})}{\sqrt{c}\, q_{\max}}\right\}, \qquad r_{.95} = 0.45, \quad r_{\max} = 0.8.$$

For the prescaling norm q, Equation (10) gives the normalized radius $\sqrt{c}\|z_H\| = \tanh(\sqrt{c}\,\alpha_{\text{cal}} q)$. Monotonicity bounds the training 95th-percentile and maximum radii by $r_{.95}$ and $r_{\max}$, respectively. The implementation bounds the denominators away from zero and freezes the fitted α_{cal}.

Lemma B.2 (Full-dimensional invertibility)**.** *If $s, \alpha_{\text{cal}}, c, w_E > 0$ and Q is orthogonal, the untruncated map in Equation (2) is a diffeomorphism between $\mathbb{R}^D$ and $\mathcal{M}$, with inverse D.*

Proof. The affine map $x \mapsto (\alpha_{\text{cal}}/s)Q^\top(x-m)$ has inverse $v \mapsto m+(s/\alpha_{\text{cal}})Qv$. By the Hadamard property, Exp_0^H has smooth global inverse Log_0^H, including at the origin. The Euclidean factor is the identity, so the decoder D in Equation (2) satisfies $D \circ E = \operatorname{Id}_{\mathbb{R}^D}$ and $E \circ D = \operatorname{Id}_{\mathcal{M}}$. □

The domain of this lemma contains all principal component (PC) coordinates supplied to the transport model, establishing the exact correspondence between analytical calibration and decoding. Appendix C.14 and Proposition C.19 describe the maps with numerical safeguards.

Let $J_E(x)$ denote the Jacobian of the calibration in the product coordinates. Under Lemma B.2, define

$$g_{\text{cal}}(x) = J_E(x)^\top g_{E(x)} J_E(x).$$

The matrix $J_E(x)$ is invertible and $g_{E(x)}$ is positive definite, so g_{cal} is a smooth positive-definite metric. For any tangent vectors $\xi, \zeta \in T_x\mathbb{R}^D$,

$$\xi^\top g_{\text{cal}}(x)\zeta = \langle dE_x[\xi], dE_x[\zeta]\rangle_{g_{E(x)}}.$$

Consequently, every piecewise C^1 curve $\gamma : [0,1] \to \mathbb{R}^D$ satisfies

$$\operatorname{Length}_{g_{\text{cal}}}(\gamma) = \int_0^1 \sqrt{\dot\gamma_t^\top g_{\text{cal}}(\gamma_t)\dot\gamma_t}\, dt = \int_0^1 \|dE_{\gamma_t}[\dot\gamma_t]\|_{g_{E(\gamma_t)}}\, dt = \operatorname{Length}_g(E\circ\gamma).$$

Since E is a diffeomorphism, it gives a bijection between the curves joining x, y and those joining $E(x), E(y)$. Taking infima of their lengths yields

$$d_{g_{\text{cal}}}(x,y) = d_g(E(x), E(y)).$$

Thus $E : (\mathbb{R}^D, g_{\text{cal}}) \to (\mathcal{M}, g)$ is an isometry. Full-dimensional invertibility supplies an information-preserving coordinate identification, while the prescribed metric and its calibrated pullback supply the geometric inductive bias. Proposition 5.1 computes this pullback explicitly.

B.3 Proof of Proposition 5.1

Invertibility provides common coordinates for comparison. We now compute the metric in these coordinates to describe how the costs of radial motion and transverse separation vary with position.

Proof. First, let $d_H \geq 2$ and $v_H = a\omega$, where $a > 0$ and $\|\omega\| = 1$. Write $P_\parallel = \omega\omega^\top$ and $P_\perp = I - P_\parallel$. In hyperbolic polar coordinates, the geodesic radius is $r = 2a$, and the metric is

$$ds_H^2 = dr^2 + \frac{\sinh^2(\sqrt{c}\, r)}{c}\, ds^2_{\mathbb{S}^{d_H-1}}.$$

For an arbitrary tangent perturbation δv_H,

$$\delta a = \omega^\top \delta v_H, \qquad \delta\omega = \frac{P_\perp \delta v_H}{a}, \qquad \delta r = 2\omega^\top \delta v_H.$$

Substituting into the polar metric yields

$$\begin{aligned}\|d\operatorname{Exp}_0^H(v_H)[\delta v_H]\|_{g_H}^2 &= 4\|P_\parallel \delta v_H\|^2 + \frac{\sinh^2(2\sqrt{c}\,a)}{ca^2}\|P_\perp \delta v_H\|^2 \\ &= 4\,\delta v_H^\top (P_\parallel + \kappa_c(a)^2 P_\perp)\delta v_H.\end{aligned}$$

The Euclidean factor contributes $w_E\|\delta v_E\|^2$. Since $\delta v = (\alpha_{\mathrm{cal}}/s)Q^\top \delta x$, we obtain

$$E^* g = \left(\frac{\alpha_{\mathrm{cal}}}{s}\right)^2 Q\,\mathrm{diag}\big(4(P_\parallel + \kappa_c(a)^2 P_\perp), w_E I\big)\,Q^\top,$$

which is Equation (7). The expansion $\sinh x/x = 1 + x^2/6 + O(x^4)$ gives $\kappa_c(a)^2 = 1 + 4ca^2/3 + O(a^4)$. Hence the hyperbolic block tends to $4I$ as $a \to 0$, independently of the direction of approach. For $d_H = 1$, $P_\perp = 0$, and the same expression reduces to the radial metric.

For two points at radius r with included angle ϑ, the hyperbolic law of cosines gives

$$\cosh(\sqrt{c}\,d_H) = \cosh^2(\sqrt{c}\,r) - \sinh^2(\sqrt{c}\,r)\cos\vartheta = 1 + 2\sinh^2(\sqrt{c}\,r)\sin^2(\vartheta/2).$$

Using $\cosh(2\operatorname{arsinh} u) = 1 + 2u^2$ and $d_H \geq 0$ yields

$$d_H = \frac{2}{\sqrt{c}}\operatorname{arsinh}\big(\sinh(\sqrt{c}\,r)\sin(\vartheta/2)\big).$$

Set $x = \sqrt{c}\,r$ and $s_\vartheta = \sin(\vartheta/2)$. The function $f(u) = \operatorname{arsinh} u$ satisfies $f(0) = 0$ and $f''(u) = -u(1+u^2)^{-3/2} \leq 0$. Thus,

$$f(s_\vartheta \sinh x) \geq s_\vartheta f(\sinh x) + (1 - s_\vartheta)f(0) = s_\vartheta x.$$

Multiplication by $2/\sqrt{c}$ gives the lower bound in Equation (8). Strict concavity gives a strict inequality for $r > 0$ and $0 < \vartheta < \pi$; equality holds when $r = 0$ or $\vartheta \in \{0, \pi\}$. □

These calculations identify the transverse coefficient $\kappa_c(a)^2$ as the source of the geometric effect. The Euclidean control in Table 7 retains the same $m, s, Q, \alpha_{\mathrm{cal}}$, scales the coordinates corresponding to the hyperbolic factor by 2, and scales the Euclidean coordinates by $\sqrt{w_E}$. Its pullback metric is

$$g_{\mathrm{E}}^{\mathrm{cal}} = \left(\frac{\alpha_{\mathrm{cal}}}{s}\right)^2 Q\,\mathrm{diag}(4I_{d_H}, w_E I_{d_E})Q^\top.$$

This equals $E^* g$ at the origin, fixing the local units of noise and control while retaining the curvature-dependent differences away from the origin.

B.4 Corollary B.3: Path-entropy invariance under invertible decoding

The metric calculation describes how geometry changes the reference and motion cost. The following identity gives the relative entropy of fixed path laws under invertible coordinates.

Corollary B.3 (Path-entropy invariance under full-dimensional decoding)**.** *Under Lemma B.2, let* $\mathcal{T}_D(k, \omega) = (k, D \circ \omega)$. *For path laws* P, R, $\mathrm{KL}((\mathcal{T}_D)_\# P \| (\mathcal{T}_D)_\# R) = \mathrm{KL}(P\|R)$.

Proof. The path maps $\mathcal{T}_D$ and $\mathcal{T}_E$ are continuous inverses in the uniform topology. Set $P' = (\mathcal{T}_D)_\# P$ and $R' = (\mathcal{T}_D)_\# R$. If $P \ll R$, change of variables sends $f = dP/dR$ to $dP'/dR' = f \circ \mathcal{T}_E$, giving

$$\mathrm{KL}(P'\|R') = \int (f \circ \mathcal{T}_E)\log(f \circ \mathcal{T}_E)\,dR' = \int f \log f\,dR = \mathrm{KL}(P\|R).$$

As $x \log x \geq -e^{-1}$, the identity allows $+\infty$. If $P \not\ll R$, a measurable A with $R(A) = 0 < P(A)$ has measurable image $\mathcal{T}_D(A)$ satisfying $R'(\mathcal{T}_D(A)) = 0 < P'(\mathcal{T}_D(A))$, so both sides are infinite. □

Path relative entropy is therefore invariant when the controlled and reference laws are transformed together. The geometric design enters the learning objective through the reference process and the control norm; Appendix C.3 develops the corresponding energy interpretation.

C Path-Space Theory and Algorithmic Analysis

This section formulates the geometric control objective as a variational problem on path measures. We first use the generator and the Girsanov transformation to establish the entropy interpretation of control energy. We then analyze the state potential and endpoint risk, and use entropy compactness to prove existence for the soft endpoint problem and its penalty limit. Throughout, the calibration map, reference process, and condition prior are fixed. The analysis characterizes the common distribution-transport objective in Equation (5); Appendix F.8 specifies its experimental augmentation for LARRY.

C.1 Variational formulation of conditional transport

Definition C.1 (Variational formulation of the conditional hyperbolic Schrödinger bridge)**.** Let $\Omega = \{1, \dots, K\} \times C([0,1], \mathcal{M})$, and let $\bar{R}(k, d\omega) = \pi_k R^k(d\omega)$ be a fixed reference, with $\pi_k > 0$ and $\sum_k \pi_k = 1$. Let $\mathcal{P}_0$ be the set of path probability measures whose initial joint law is $\pi_k \nu_0^k$. For $P \in \mathcal{P}_0$, denote its conditional path laws by P^k, their marginals at time t by P_t^k, and set $\rho_1^k = D_{\#} P_1^k$. The ideal problem with hard marginal constraints is

$$\inf_{\substack{P \in \mathcal{P}_0 \\ \rho_1^k = \mu_1^k,\ \forall k}} F(P), \qquad F(P) = \beta\, \mathrm{KL}(P \| \bar{R}) + \lambda_V \mathbb{E}_P I_V, \qquad I_V = \int_0^1 V_{\mathcal{U}}(t, Z_t, k)\, dt,$$

where $\beta > 0$ and $\lambda_V \geq 0$. The potential term is defined to be zero when $\lambda_V = 0$.

The set $\mathcal{P}_0$ fixes the source distributions and condition masses, while the terminal constraints fix each condition's target distribution. For $\lambda_V = 0$, the objective is path relative entropy. Corollary C.7 absorbs a nonnegative state potential into a tilted reference that preserves the initial law. Theorem 5.4 connects this hard-constrained problem to the soft endpoint objective in the main text.

To explain the fixed reference, consider a labeled path law P and write $\mathcal{C}(P) = \mathcal{R}_\alpha(e(P)) + \lambda_V \mathbb{E}_P I_V$. This collects terms that depend on the generated law rather than its drift decomposition.

Lemma C.2 (Residual cost under drift reallocation)**.** *Fix the geometry, diffusion coefficient, and initial joint law. Suppose each admissible total drift b determines a unique labeled path law P^b, with conditional laws $P^{b,k}$. For drift fields a and u, define*

$$\mathcal{J}_{\mathrm{joint}}(a, u) = \mathcal{C}(P^{a+u}) + \beta \mathcal{K}_{a+u}(u), \qquad \mathcal{K}_b(u) = \sum_k \pi_k \mathbb{E}_{P^{b,k}} \int_0^1 \frac{\|u^k(t, Z_t)\|_g^2}{2\sigma^2}\, dt,$$

where $\beta > 0$. If (a, u) and $(a + u, 0)$ are both admissible and the objective is finite, then

$$\mathcal{J}_{\mathrm{joint}}(a + u, 0) = \mathcal{J}_{\mathrm{joint}}(a, u) - \beta \mathcal{K}_{a+u}(u).$$

Consequently, when a jointly optimized reference class can absorb the control and has no separate penalty, every attainable path law admits a decomposition with zero residual cost.

Proof. Equal total drift $a + u$, diffusion, and initial law imply the same P^{a+u} and $\mathcal{C}$ by uniqueness in law. Substituting $\mathcal{K}_{a+u}(0) = 0$ proves both claims for each admissible pair. □

In HyRef-SB, reference fitting determines $a = b_{\mathrm{ref}}$ before bridge optimization. Keeping this field fixed makes the residual energy a cost relative to that learned path prior.

C.2 Probability space and basic assumptions

We equip path space with the topology of uniform convergence and its Borel σ-algebra, and use the usual augmentation of the canonical filtration. With the reference fixed, the controlled generator is

$$\mathcal{L}_t^{k,u} f = \langle b_{\mathrm{ref}}^k + u^k, \nabla_g f \rangle_g + \frac{\sigma^2}{2} \Delta_g f, \qquad \sigma > 0. \tag{11}$$

The reference law R^k corresponds to $u^k = 0$, and P^k and R^k share the initial law ν_0^k. The analysis allows controls $u^k(t, \omega)$ that are predictable with respect to the canonical filtration. The main text

uses their Markov parameterization $u_\theta^k(t, Z_t)$. The generator and change of measure follow the standard framework of stochastic analysis on manifolds (Hsu, 2002; Léonard, 2012).

Assumption C.3 (Conditions for the relative-entropy analysis)**.** The canonical martingale problems for the reference and controlled diffusions admit nonexplosive weak solutions and satisfy uniqueness in law. In the global coordinates of our product manifold, choose the smooth orthonormal frame $F = \mathrm{diag}(\lambda_c^{-1} I, w_E^{-1/2} I)$ and set $a_t^k = \sigma^{-1} F(Z_t)^{-1} u_t^k$. Assume that $\int_0^1 \|a_t^k\|^2 dt < \infty$ almost surely under the given reference law, and further assume that

$$L_t^k = \exp\left\{ \int_0^t \langle a_s^k, dW_s^{R,k} \rangle - \frac{1}{2} \int_0^t \|a_s^k\|^2 ds \right\}$$

is a true martingale with mean 1. The drift and this density are measurable with respect to the canonical path. In addition, $\mathbb{E}_{P^k} \int_0^1 \|a_t^k\|^2 dt < \infty$.

This assumption provides existence and integrability for the path-space change of measure. We derive coordinate equations, then give verifiable growth conditions ensuring these properties.

Lemma C.4 (Coordinate generator on the product space)**.** *Let $d = d_H$, $\lambda = \lambda_c(z_H)$, and $b = b_{\mathrm{ref}} + u$. Equation (11) is the generator of the following coordinate Itô equations:*

$$dZ_H = \left[b_H + \frac{\sigma^2}{2}(d-2) c \lambda^{-1} Z_H \right] dt + \sigma \lambda^{-1} dW_H, \qquad dZ_E = b_E dt + \sigma w_E^{-1/2} dW_E. \quad (12)$$

Proof. Since $|g|^{1/2} = \lambda^d w_E^{d_E/2}$,

$$\begin{aligned} \Delta_g f &= |g|^{-1/2} \partial_i (|g|^{1/2} g^{ij} \partial_j f) = \lambda^{-d} \sum_{i=1}^{d} \partial_i (\lambda^{d-2} \partial_i f) + w_E^{-1} \Delta_E f \\ &= \lambda^{-2} \Delta_H^{\mathrm{coord}} f + (d-2) \lambda^{-3} \nabla \lambda \cdot \nabla_H f + w_E^{-1} \Delta_E f. \end{aligned}$$

Direct differentiation gives $\partial_i \lambda = 4 c z_i / (1 - c\|z_H\|^2)^2 = c \lambda^2 z_i$, so

$$\frac{\sigma^2}{2} \Delta_g f = \frac{\sigma^2}{2} \lambda^{-2} \Delta_H^{\mathrm{coord}} f + \frac{\sigma^2}{2} (d-2) c \lambda^{-1} z_H \cdot \nabla_H f + \frac{\sigma^2}{2 w_E} \Delta_E f.$$

The coordinate Itô generator consists of the first-order drift term and the contraction of half the covariance with the Hessian. Adding $b \cdot \nabla f$ gives Equation (12). □

Denote the geometric drift correction by $G(z)$. The canonical path determines the reference noise as

$$W_t^{R,k} = \frac{1}{\sigma} \int_0^t F(Z_s)^{-1} [dZ_s - \{ b_{\mathrm{ref}}^k(s, Z_s) + G(Z_s) \} ds].$$

Thus, path-predictable controls yield path-measurable Girsanov densities. Appendix C.14 shows that exponential-map random steps recover the same generator through a geodesic expansion.

Lemma C.5 (Sufficient conditions for nonexplosion and a martingale density)**.** *On the complete product manifold above, suppose that b_{ref}, u are continuous in time and locally Lipschitz in the state, with local Lipschitz constants uniform on compact time intervals. Fix $o \in \mathcal{M}$. Suppose that a constant $C < \infty$ bounds the reference drift growth and the control norm as*

$$\|b_{\mathrm{ref}}(t, z)\|_g \le C(1 + d_g(o, z)), \qquad \|u(t, z)\|_g \le C,$$

and the initial law has a finite second geodesic moment, then both diffusions are nonexplosive, unique in law, and satisfy all the conditions of Assumption C.3.

Proof. Smooth nondegenerate diffusion coefficients and locally Lipschitz drifts give a pathwise unique solution up to the explosion time. Let $r_H = d_H(o_H, z_H)$ and $r^2 = d_g(o, z)^2 = r_H^2 + w_E \|z_E - o_E\|^2$. The Hadamard property makes r^2 smooth. The hyperbolic radial Laplacian gives

$$\Delta_g r^2 = 2 + 2(d_H - 1) \sqrt{c}\, r_H \coth(\sqrt{c}\, r_H) + 2 d_E.$$

At $r_H = 0$, use the limit $x \coth x \to 1$; when $d_H = 1$, the middle term vanishes. Since $x \coth x \le 1 + x$ for $x \ge 0$, there is a constant C_1 such that $\Delta_g r^2 \le C_1(1 + r^2)$. Moreover,

$$|d(r^2)[b]| \le 2r\|b\|_g \le C_2(1 + r^2), \qquad \mathcal{L}_t(1 + r^2) \le C_3(1 + r^2)$$

for both the reference and controlled drifts. Set $\tau_R = \inf\{t : r(Z_t) \ge R\}$. Applying Itô's formula and Gronwall's inequality to the stopped process gives

$$\mathbb{E}[1 + r^2(Z_{t\wedge\tau_R})] \le \mathbb{E}[1 + r^2(Z_0)] + C_3 \int_0^t \mathbb{E}[1 + r^2(Z_{s\wedge\tau_R})]\, ds \le \mathbb{E}[1 + r^2(Z_0)]e^{C_3 t}.$$

Consequently, $\Pr(\tau_R \le 1) \le \mathbb{E}[1 + r^2(Z_0)]e^{C_3}/(1 + R^2) \to 0$. The solution extends to $[0, 1]$, and local pathwise uniqueness gives global uniqueness in law.

The orthonormal frame satisfies $F^\top g F = I$, so

$$\int_0^1 \|a_t\|^2 dt = \frac{1}{\sigma^2} \int_0^1 \|u_t\|_g^2 dt \le C^2/\sigma^2, \qquad \mathbb{E}_R \exp\left(\frac{1}{2} \int_0^1 \|a_t\|^2 dt\right) \le e^{C^2/(2\sigma^2)}.$$

Novikov's condition ensures that L_t is a true martingale, and the same deterministic energy bound holds after the change of measure. The noise-recovery formula above establishes measurability with respect to the canonical path, completing the verification of Assumption C.3. □

This lemma gives sufficient conditions for the continuous model. Its second-moment estimate also helps verify integrability of the state potential in Appendix C.12.

C.3 Theorem 5.2: The entropy interpretation of control energy

We apply the canonical-path Girsanov identity (Léonard, 2012; Hsu, 2002) to the fixed geometric reference and then decompose the contribution of the condition labels. This identifies the residual control as the quantity measured by relative entropy.

Proof. Fix a condition k and temporarily omit the superscript. By Assumption C.3, $dP/dR = L_1$, and Girsanov's theorem gives $W_t^P = W_t^R - \int_0^t a_s ds$. Since $\mathbb{E}_R[L_1 \mid \mathcal{F}_0] = L_0 = 1$, the initial law is preserved. Uniqueness in law identifies the law obtained by this change of measure with the controlled law. Substituting $dW^R = dW^P + a\, dt$ into the density yields

$$\log \frac{dP}{dR} = \int_0^1 \langle a_t, dW_t^R \rangle - \frac{1}{2} \int_0^1 \|a_t\|^2 dt = \int_0^1 \langle a_t, dW_t^P \rangle + \frac{1}{2} \int_0^1 \|a_t\|^2 dt.$$

Finite energy and the Itô isometry give

$$\mathbb{E}_P \left| \int_0^1 \langle a_t, dW_t^P \rangle \right|^2 = \mathbb{E}_P \int_0^1 \|a_t\|^2 dt < \infty.$$

This stochastic integral has mean zero. Since $(F^{-1})^\top F^{-1} = g$,

$$\mathrm{KL}(P^k \| R^k) = \frac{1}{2} \mathbb{E}_{P^k} \int_0^1 \|a_t^k\|^2 dt = \mathbb{E}_{P^k} \int_0^1 \frac{\|u_t^k\|_g^2}{2\sigma^2} dt.$$

The joint laws with labels satisfy $d\bar{P}/d\bar{R}(k, \omega) = dP^k/dR^k(\omega)$, hence

$$\mathrm{KL}(\bar{P} \| \bar{R}) = \sum_k \pi_k \int \log \frac{dP^k}{dR^k}\, dP^k = \sum_k \pi_k \mathbb{E}_{P^k} \int_0^1 \frac{\|u_t^k\|_g^2}{2\sigma^2} dt = \mathcal{K}.$$ □

Corollary C.6 (Entropy gap after marginalizing branch labels). *Write $P_{\mathrm{mix}} = \sum_k \pi_k P^k$ and $R_{\mathrm{mix}} = \sum_k \pi_k R^k$, and let $p(k \mid \omega), r(k \mid \omega)$ be their respective label posteriors. Under the conditions of Theorem 5.2,*

$$\mathcal{K} = \mathrm{KL}(P_{\mathrm{mix}} \| R_{\mathrm{mix}}) + \mathbb{E}_{P_{\mathrm{mix}}} \mathrm{KL}\big(p(\cdot \mid \omega) \| r(\cdot \mid \omega)\big).$$

Thus, the Kullback–Leibler (KL) divergence between the mixture path laws is bounded by the control energy, with equality precisely when the two label posteriors agree P_{mix}-almost everywhere.

Proof. The Polish structure ensures the existence of regular conditional probabilities for the labels. Let $f = dP_{\mathrm{mix}}/dR_{\mathrm{mix}}$. Disintegrating the joint measures gives

$$\frac{d\bar{P}}{d\bar{R}}(k,\omega) = f(\omega)\frac{p(k \mid \omega)}{r(k \mid \omega)} \quad \bar{P}\text{-almost everywhere.}$$

Taking logarithms and integrating,

$$\begin{aligned}\mathrm{KL}(\bar{P}\|\bar{R}) &= \int \sum_k p(k \mid \omega)\left[\log f(\omega) + \log\frac{p(k \mid \omega)}{r(k \mid \omega)}\right] P_{\mathrm{mix}}(d\omega) \\ &= \mathrm{KL}(P_{\mathrm{mix}}\|R_{\mathrm{mix}}) + \int \mathrm{KL}(p(\cdot \mid \omega)\|r(\cdot \mid \omega))\, P_{\mathrm{mix}}(d\omega).\end{aligned}$$

The second term is nonnegative and vanishes exactly when the two label-posterior probability vectors agree almost everywhere under the mixture path law. □

The common initial law removes the initial entropy term. The general initial-state decomposition is

$$\mathrm{KL}(P\|R) = \mathrm{KL}(P_0\|R_0) + \int \mathrm{KL}(P(\cdot \mid Z_0 = z)\|R(\cdot \mid Z_0 = z))\, P_0(dz),$$

which explains the role of fixing the initial joint law in the main text.

C.4 A REFERENCE TILTED BY THE STATE COST AT EACH INITIAL POINT

Control energy measures deviation from the reference. The state potential also defines an exponential weight on reference paths. Let $a = \lambda_V/\beta$, assume that $I_V \geq 0$ is measurable, and define

$$Z_V(z,k) = \mathbb{E}_{R^k(\cdot|Z_0=z)} e^{-aI_V}, \qquad \frac{dR^{V,k}}{dR^k}(\omega) = \frac{e^{-aI_V(\omega,k)}}{Z_V(Z_0,k)}.$$

For $a = 0$, both the exponential weight and Z_V are defined to be 1. At any initial point for which $Z_V(z,k) > 0$, the conditional density has the form

$$\frac{dR^{V,k}(\cdot \mid Z_0 = z)}{dR^k(\cdot \mid Z_0 = z)}(\omega) = \frac{\exp[-\lambda_V I_V(\omega,k)/\beta]}{Z_V(z,k)}. \tag{13}$$

Corollary C.7 (Potential-tilted reference preserving the initial law)**.** *Suppose that $\beta > 0$, $\lambda_V \geq 0$, and $Z_V(z,k) > 0$ almost everywhere under the initial joint law, with $\sum_k \pi_k \int |\log Z_V(z,k)|\, \nu_0^k(dz) < \infty$. If $\bar{P} \in \mathcal{P}_0$ has finite relative entropy and finite weighted potential energy, then $\bar{R}^V(k,d\omega) = \pi_k R^{V,k}(d\omega)$ preserves the same initial joint law and*

$$\beta\,\mathrm{KL}(\bar{P}\|\bar{R}) + \lambda_V \mathbb{E}_{\bar{P}} I_V = \beta\,\mathrm{KL}(\bar{P}\|\bar{R}^V) - \beta \sum_k \pi_k \int \log Z_V(z,k)\, \nu_0^k(dz). \tag{14}$$

Proof. For any measurable set A of initial states,

$$R^{V,k}(Z_0 \in A) = \int_A \mathbb{E}_{R^k(\cdot|Z_0=z)} \frac{e^{-aI_V}}{Z_V(z,k)}\, \nu_0^k(dz) = \nu_0^k(A).$$

If $a > 0$, finite potential energy gives $I_V < \infty$ almost surely under $\bar{P}$. Together with $\bar{P} \ll \bar{R}$ and the positive partition function, this implies $\bar{P} \ll \bar{R}^V$. The relation also holds when $a = 0$, giving

$$\log\frac{d\bar{P}}{d\bar{R}^V} = \log\frac{d\bar{P}}{d\bar{R}} + aI_V + \log Z_V(Z_0,k).$$

By integrability, taking expectations under $\bar{P}$ with its fixed initial law gives Equation (14). □

The final term depends only on the initial law and the reference. Minimizing the sum of relative entropy and potential energy is therefore equivalent to minimizing relative entropy with respect to the tilted reference. Normalization at each initial point preserves both the condition masses and the source distributions. If $0 \leq I_V \leq C$, then $e^{-aC} \leq Z_V \leq 1$. More generally, conditional Jensen's inequality gives

$$0 \leq -\log Z_V(z,k) \leq a\,\mathbb{E}_{R^k}[I_V \mid Z_0 = z].$$

Finite expected reference potential energy ensures initial-law integrability of the log partition function.

C.5 Markov interpretation of the tilted reference and static feasibility

The preceding corollary is an identity on path measures. For a Markov reference, the Feynman–Kac formula translates the exponential tilt into a drift correction. We then use endpoint disintegration to construct a feasible law with finite entropy, as required by the penalty limit.

Proposition C.8 (Doob drift induced by a state potential)**.** *Suppose that the reference is a nonexplosive Markov diffusion, $V(t,z,k)$ is bounded and nonnegative, and*

$$h_k(t,z) = \mathbb{E}_{R^k_{t,z}} \exp\left[-\frac{\lambda_V}{\beta}\int_t^1 V(s,Z_s,k)ds\right]$$

is a positive $C^{1,2}$ function satisfying the Feynman–Kac equation. Then $h_k(1,z) = 1$, $(\partial_t + \mathcal{L}_t^{k,0})h_k = (\lambda_V/\beta)Vh_k$, and the reference in Equation (13) has generator

$$\mathcal{L}_t^{V,k} f = \mathcal{L}_t^{k,0} f + \sigma^2\langle \nabla_g \log h_k, \nabla_g f\rangle_g.$$

Proof. Omit k, set $a = \lambda_V/\beta$, and define the Feynman–Kac operator

$$Q_{s,t}f(z) = \mathbb{E}_{R_{s,z}}\left[e^{-a\int_s^t V(r,Z_r)dr} f(Z_t)\right].$$

The Markov property gives $Q_{s,t}h(t,\cdot) = h(s,\cdot)$. Under the tilted law, for $s < t$,

$$\mathbb{E}_{R^V}[f(Z_t) \mid Z_s = z] = \frac{Q_{s,t}(h(t,\cdot)f)(z)}{h(s,z)}.$$

The right-hand side equals 1 when $f = 1$, defining a conservative transition kernel. For smooth compactly supported f, a local generator expansion gives

$$\begin{aligned}\mathcal{L}_t^V f &= h^{-1}\big[(\partial_t + \mathcal{L}_t)(hf) - aVhf\big] \\ &= \mathcal{L}_t f + fh^{-1}(\partial_t h + \mathcal{L}_t h - aVh) + \sigma^2 h^{-1}\langle \nabla_g h, \nabla_g f\rangle_g = \mathcal{L}_t f + \sigma^2\langle \nabla_g \log h, \nabla_g f\rangle_g.\end{aligned}$$

We used the generator product rule and the backward equation for h. The identity $h(0,z) = Z_V(z,k)$ connects this kernel to the preceding initial-point normalization. □

The state potential therefore induces the reference drift $b_{\mathrm{ref}} + \sigma^2\nabla_g \log h_k$, the Feynman–Kac/Doob transformation for the fixed diffusion used here (Léonard, 2014; Hsu, 2002).

Lemma C.9 (Static sufficient conditions for feasibility with finite entropy)**.** *For each k, let the endpoint coupling γ^k have initial marginal ν_0^k, terminal marginal $E_\# \mu_1^k$, and $\mathrm{KL}(\gamma^k \| R_{01}^k) < \infty$. Denote the regular conditional bridges of the reference by R^{k,z_0,z_1}. Then*

$$P^{\dagger,k}(d\omega) = \int R^{k,z_0,z_1}(d\omega)\,\gamma^k(dz_0,dz_1), \qquad \mathrm{KL}(P^{\dagger,k}\|R^k) = \mathrm{KL}(\gamma^k\|R_{01}^k).$$

If $\mathbb{E}_{P^{\dagger,k}} I_V < \infty$, their π-mixture has $F(P^\dagger) < \infty$ and satisfies all hard marginal constraints.

Proof. Let $f_k = d\gamma^k/dR_{01}^k$. For any bounded measurable path function φ, disintegration gives

$$\int \varphi\, dP^{\dagger,k} = \int\left[\int \varphi\, dR^{k,z_0,z_1}\right] f_k(z_0,z_1) R_{01}^k(dz_0,dz_1) = \int \varphi(\omega) f_k(Z_0,Z_1)\, R^k(d\omega).$$

Hence $dP^{\dagger,k}/dR^k = f_k(Z_0,Z_1)$, its joint endpoint law is γ^k, and

$$\mathrm{KL}(P^{\dagger,k}\|R^k) = \int f_k(Z_0,Z_1)\log f_k(Z_0,Z_1)\, dR^k = \int f_k \log f_k\, dR_{01}^k.$$

The identity $D_\# E_\# \mu_1^k = \mu_1^k$ ensures the endpoint constraint in observation space. Finitely many labels, finite relative entropy, and finite potential energy make the joint objective finite. □

The same disintegration gives the classical endpoint entropy decomposition (Léonard, 2014):

$$\mathrm{KL}(P^k\|R^k) = \mathrm{KL}(P^k_{01}\|R^k_{01}) + \int \mathrm{KL}(P^{k,z_0,z_1}\|R^{k,z_0,z_1})\, P^k_{01}(dz_0, dz_1).$$

For a fixed endpoint coupling, retaining the reference conditional bridges therefore minimizes path relative entropy. The construction above reduces hard feasibility to integrability conditions on the endpoint coupling and the potential.

The endpoint entropy decomposition separates endpoint couplings from conditional paths. Endpoint reweighting (Léonard, 2014) makes this distinction explicit.

Proposition C.10 (Endpoint-factorized reference changes)**.** *Let $\bar{R}(k, d\omega) = \pi_k R^k(d\omega)$ and $\widetilde{\bar{R}}(k, d\omega) = \pi_k \widetilde{R}^k(d\omega)$ be reference probability laws with the same initial joint law. Suppose*

$$\frac{d\widetilde{R}^k}{dR^k}(\omega) = a_k(Z_0)b_k(Z_1),$$

where $a_k, b_k : \mathcal{M} \to (0,\infty)$ are finite measurable functions. Set $\nu^k_1 = E_\# \mu^k_1$ and assume $\log a_k \in L^1(\nu^k_0)$ and $\log b_k \in L^1(\nu^k_1)$. Define the corresponding change in the objective by

$$\widetilde{F}(P) = \beta\,\mathrm{KL}(P\|\widetilde{\bar{R}}) + \lambda_V \mathbb{E}_P I_V, \qquad C_{a,b} = \sum_k \pi_k \left[\int \log a_k\, d\nu^k_0 + \int \log b_k\, d\nu^k_1\right].$$

For every $P \in \mathcal{H}$ with finite $F(P)$,

$$\widetilde{F}(P) = F(P) - \beta C_{a,b}. \tag{15}$$

The two hard-marginal problems have the same finite-valued minimizers. Moreover, R^k and $\widetilde{R}^k$ have identical conditional path laws given (Z_0, Z_1), for R^k_{01}-almost every endpoint pair.

Proof. The density $a_k(Z_0)b_k(Z_1)$ is strictly positive and finite, so the two references are equivalent. For a finite-entropy $P \in \mathcal{H}$, the chain rule for Radon–Nikodym derivatives gives

$$\log \frac{dP^k}{d\widetilde{R}^k} = \log \frac{dP^k}{dR^k} - \log a_k(Z_0) - \log b_k(Z_1).$$

The last two terms are integrable under P^k because its endpoint marginals are ν^k_0 and ν^k_1. Taking expectations under each conditional law and weighting by π_k yields

$$\mathrm{KL}(P\|\widetilde{\bar{R}}) = \mathrm{KL}(P\|\bar{R}) - C_{a,b}.$$

The state-cost term is unchanged, proving Equation (15). Reversing the identity shows that the finite-objective domains agree; subtracting a fixed constant preserves their minimizers.

For any bounded measurable path function φ, conditional change of measure gives

$$\mathbb{E}_{\widetilde{R}^k}[\varphi \mid Z_0, Z_1] = \frac{\mathbb{E}_{R^k}[a_k(Z_0)b_k(Z_1)\varphi \mid Z_0, Z_1]}{\mathbb{E}_{R^k}[a_k(Z_0)b_k(Z_1) \mid Z_0, Z_1]} = \mathbb{E}_{R^k}[\varphi \mid Z_0, Z_1].$$

This identifies the conditional path laws. □

For the soft problems $J_\eta = F + \eta\mathcal{D}$ and $\widetilde{J}_\eta = \widetilde{F} + \eta\mathcal{D}$, the same calculation gives, whenever the displayed logarithmic expectations and objectives are finite,

$$\begin{aligned} \widetilde{J}_\eta(P) &= J_\eta(P) - \beta C_a - \beta \sum_k \pi_k \int \log b_k(z)\, P^k_1(dz), \\ C_a &= \sum_k \pi_k \int \log a_k(z)\, \nu^k_0(dz), \qquad P \in \mathcal{P}_0. \end{aligned} \tag{16}$$

Here the terminal marginals vary with the control. Endpoint-factorized reweighting contributes a terminal potential to the soft objective, even when the hard-marginal optimizer is unchanged. General reference changes can modify the endpoint coupling, the endpoint-conditioned paths, or both.

C.6 Compatibility of the initial joint law with a fixed prior

Fixed source-conditioned probabilities $q(k \mid x)$ define the initial joint law $\gamma_0(dx, k) = \mu_0(dx)q(k \mid x)$. Let $\tilde{\pi}_k = \int q(k \mid x)\mu_0(dx) > 0$. Then

$$\nu_0^k = E_\#\left(\frac{q(k \mid x)}{\tilde{\pi}_k}\mu_0(dx)\right), \qquad \sum_k \tilde{\pi}_k \nu_0^k = E_\#\mu_0.$$

The second identity follows by testing against an arbitrary bounded measurable f:

$$\sum_k \tilde{\pi}_k \int f\, d\nu_0^k = \int f(E(x)) \sum_k q(k \mid x)\, \mu_0(dx) = \int f\, d(E_\#\mu_0).$$

For source-independent conditions, $q(k \mid x) = \pi_k$, and all conditions share $\nu_0^k = E_\#\mu_0$. For an optimal transport (OT) coupling of training endpoints, the column masses are set to the training prior π, so the empirical joint law satisfies $\tilde{\pi} = \pi$ up to solver error. At test time, the fixed conditional probabilities determine the actual $\tilde{\pi}$. When applying the entropy decomposition to this test population, the corresponding weights and conditional initial laws are given by the formulas above.

C.7 Proposition C.11: endpoint risk and subgradients

Path energy is aggregated using fixed condition masses, and endpoint risk emphasizes difficult conditions within the same collection of condition scores. We give a support-function representation of the risk and bounds for each condition, which are then used in the penalty limit.

Proposition C.11 (Distributionally robust form of endpoint risk)**.** *Let $\alpha \in [0, 1]$ and $\mathcal{W}_\alpha = (1 - \alpha)\pi + \alpha\Delta_K$. Then $\mathcal{R}_\alpha(e) = \max_{w \in \mathcal{W}_\alpha} w^\top e$. Distributing the mass α equally among conditions tied for the largest error gives a valid subgradient with respect to e.*

Proof. Writing $w = (1 - \alpha)\pi + \alpha v$ with $v \in \Delta_K$ gives

$$\max_{w \in \mathcal{W}_\alpha} w^\top e = (1 - \alpha)\pi^\top e + \alpha \max_k e_k.$$

For $A(e) = \{k : e_k = \max_j e_j\}$, the maximizing weights form $\partial\mathcal{R}_\alpha(e) = \{(1 - \alpha)\pi + \alpha v : v \in \Delta_K,\ \operatorname{supp} v \subseteq A(e)\}$: each such w satisfies $\mathcal{R}_\alpha(e') \geq w^\top e' = \mathcal{R}_\alpha(e) + w^\top(e' - e)$. Taking $v_k = \mathbf{1}_{\{k \in A(e)\}}/|A(e)|$ proves the stated selection. For differentiable $e_k(\theta)$,

$$D\{\mathcal{R}_\alpha \circ e\}(\theta; v) = (1 - \alpha) \sum_k \pi_k \nabla e_k(\theta)^\top v + \alpha \max_{k \in A(e(\theta))} \nabla e_k(\theta)^\top v.$$

The corresponding weighted branch gradient lies in the Clarke subdifferential of this locally Lipschitz function and equals the ordinary gradient when the maximizer is unique. □

The main text first forms the joint score $e_k = \sum_j \lambda_j d_{jk}/s_j$ using fixed scales, then selects a common difficult condition. Subadditivity of finite maxima gives

$$\max_k \sum_j \frac{\lambda_j d_{jk}}{s_j} \leq \sum_j \max_k \frac{\lambda_j d_{jk}}{s_j},$$

clarifying the order of aggregation in the joint risk.

Corollary C.12 (Deterministic bound on each branch error)**.** *If every $e_k \geq 0$, the aggregate risk controls each branch error: for every condition k, $e_k \leq \mathcal{R}_\alpha(e)/[\alpha + (1 - \alpha)\pi_k]$.*

Proof. Since $\sum_j \pi_j e_j \geq \pi_k e_k$ and $\max_j e_j \geq e_k$,

$$\mathcal{R}_\alpha(e) \geq [\alpha + (1 - \alpha)\pi_k]e_k, \qquad e_k \leq \frac{\mathcal{R}_\alpha(e)}{\alpha + (1 - \alpha)\pi_k}. \tag{17}$$

The denominator is strictly positive because $\pi_k > 0$ and $\alpha \in [0, 1]$. □

C.8 THEOREM 5.4: EXISTENCE AND THE PENALTY LIMIT

The preceding results establish the nonnegative path objective and endpoint risk. Fix

$$F(P) = \beta\,\mathrm{KL}(P\|\bar{R}) + \lambda_V \mathbb{E}_P I_V, \qquad J_\eta(P) = F(P) + \eta\mathcal{D}(P), \qquad P \in \mathcal{P}_0.$$

Assume that I_V is nonnegative and lower semicontinuous, that $\mathcal{D}$ is nonnegative and weakly lower semicontinuous, and that its zero set consists precisely of the laws satisfying all target endpoint constraints. Also assume that a law $P^\dagger$ satisfying the hard constraints has $F(P^\dagger) < \infty$; Lemma C.9 gives one way to verify this condition.

Lemma C.13 (Compactness of relative-entropy sublevel sets)**.** *For a fixed probability measure R on a Polish space and finite $C \geq 0$, the set $\{P : \mathrm{KL}(P\|R) \leq C\}$ is weakly compact.*

Proof. The variational representation

$$\mathrm{KL}(P\|R) = \sup_{f \in C_b} \left\{ \int f\,dP - \log \int e^f\,dR \right\}$$

expresses relative entropy as a supremum of weakly continuous functions, proving weak lower semicontinuity. For measurable A, write $p = P(A)$ and $r = R(A) \in (0,1)$. The partition gives

$$\begin{aligned}\mathrm{KL}(P\|R) &\geq p\log\frac{p}{r} + (1-p)\log\frac{1-p}{1-r} \\ &= p\log\frac{1}{r} + (1-p)\log\frac{1}{1-r} + p\log p + (1-p)\log(1-p) \geq p\log(1/r) - \log 2.\end{aligned}$$

Hence $P(A) \leq (C + \log 2)/\log(1/R(A))$. When $R(A) = 0$, finite relative entropy implies $P(A) = 0$. Given $\epsilon > 0$, tightness of R provides a compact set K_ϵ such that

$$R(K_\epsilon^c) < \exp[-(C + \log 2)/\epsilon].$$

Every measure in the sublevel set satisfies $P(K_\epsilon^c) \leq \epsilon$, so the set is uniformly tight. Prokhorov's theorem gives relative compactness, and weak lower semicontinuity gives closedness. Together these two properties establish weak compactness of the relative-entropy sublevel set. □

Proof of Theorem 5.4. First, we verify closedness of the constraints. The map $(k, \omega) \mapsto (k, \omega_0)$ is continuous, so the set $\mathcal{P}_0$ defined by fixing its pushforward is weakly closed. By the Portmanteau theorem, $P \mapsto \mathbb{E}_P I_V$ is lower semicontinuous; hence both F and J_η are lower semicontinuous. The set $\mathcal{H} = \{P \in \mathcal{P}_0 : \mathcal{D}(P) = 0\}$ of laws satisfying the hard constraints is also weakly closed.

Fix $\eta > 0$ and choose a minimizing sequence $P^{(j)}$ such that

$$J_\eta(P^{(j)}) \leq \inf_{\mathcal{P}_0} J_\eta + 1/j \leq F(P^\dagger) + 1, \qquad \mathrm{KL}(P^{(j)}\|\bar{R}) \leq \frac{F(P^\dagger) + 1}{\beta}.$$

Lemma C.13 provides a weakly convergent subsequence $P^{(j_\ell)} \Rightarrow P_\eta$. Closedness implies $P_\eta \in \mathcal{P}_0$. Lower semicontinuity of the objective therefore gives

$$J_\eta(P_\eta) \leq \liminf_{\ell\to\infty} J_\eta(P^{(j_\ell)}) = \inf_{\mathcal{P}_0} J_\eta.$$

Thus, the soft minimum is attained. The argument for F on $\mathcal{H}$ also gives a hard minimizer.

Next, let $\eta_n \to \infty$, $\epsilon_n \geq 0$, and $\epsilon_n \to 0$, and let P_n be approximate minimizers with error at most ϵ_n. Comparing their objective values with that of the feasible law $P^\dagger$ gives

$$F(P_n) + \eta_n\mathcal{D}(P_n) \leq F(P^\dagger) + \epsilon_n, \quad \mathcal{D}(P_n) \leq \frac{F(P^\dagger) + \epsilon_n}{\eta_n}. \tag{18}$$

Nonnegativity also gives a uniform relative-entropy bound, so $\{P_n\}$ is relatively compact in the weak topology. For any convergent subsequence $P_{n_j} \Rightarrow P_*$,

$$P_* \in \mathcal{P}_0, \qquad 0 \leq \mathcal{D}(P_*) \leq \liminf_j \mathcal{D}(P_{n_j}) = 0.$$

For any $Q \in \mathcal{H}$ with $F(Q) < \infty$,

$$F(P_{n_j}) \le J_{\eta_{n_j}}(P_{n_j}) \le F(Q) + \epsilon_{n_j}, \qquad F(P_*) \le \liminf_j F(P_{n_j}) \le F(Q).$$

Every accumulation point is therefore a minimizer of the hard-constrained problem.

Finally, we prove uniqueness. The initial joint law and conditional terminal marginals impose linear constraints on measures, so $\mathcal{H}$ is convex. Suppose that distinct $P, Q \in \mathcal{H}$ have finite objective values and attain the same minimum. Then $P, Q \ll \bar{R}$. Write $p = dP/d\bar{R}$ and $q = dQ/d\bar{R}$. For $0 < t < 1$, strict convexity gives

$$\int (tp + (1-t)q) \log(tp + (1-t)q)\, d\bar{R} < t \int p \log p\, d\bar{R} + (1-t) \int q \log q\, d\bar{R}.$$

Expected potential energy is linear, so $F(tP + (1-t)Q) < tF(P) + (1-t)F(Q)$, contradicting minimality. The hard-constrained minimizer is therefore unique. The weak topology on probability measures is metrizable. A relatively compact sequence with a single accumulation point converges, so the full sequence satisfies $P_n \Rightarrow P_*$. Comparison with P_* further gives

$$F(P_*) \le \liminf_n F(P_n) \le \limsup_n J_{\eta_n}(P_n) \le F(P_*),$$

so both $F(P_n)$ and $J_{\eta_n}(P_n)$ converge to the hard-constrained optimal value. □

For $\mathcal{D}(P) = \mathcal{R}_\alpha(e(P))$, combining Equations (17) and (18) gives

$$e_k(P_n) \le \frac{F(P^\dagger) + \epsilon_n}{\eta_n[\alpha + (1-\alpha)\pi_k]}.$$

If $e_k(P) \ge \zeta_k \,\mathrm{MMD}^2(\rho_1^k, \mu_1^k)$ for a fixed coefficient $\zeta_k > 0$, then

$$\mathrm{MMD}(\rho_{1,n}^k, \mu_1^k) \le \left[\frac{F(P^\dagger) + \epsilon_n}{\zeta_k \eta_n[\alpha + (1-\alpha)\pi_k]}\right]^{1/2}.$$

Under Proposition 5.3, $F(P^\dagger)$ can be replaced by $C_{\mathrm{T}} = \beta B_\sigma + \lambda_V \mathcal{V}_{\mathrm{T}}$. Appendix C.18 verifies the continuity and identification of the full-dimensional Gaussian-kernel MMD in the population problem.

Corollary C.14 (Error decomposition for a parameterized population problem)**.** *Assume the hypotheses of Theorem 5.4. For each $\eta > 0$, let $\mathcal{A}_\eta \subset \mathcal{P}_0$ be a nonempty model class on which J_η is finite, and define the gap between its optimum and the unrestricted population optimum by*

$$a_\eta := \inf_{P \in \mathcal{A}_\eta} J_\eta(P) - \inf_{P \in \mathcal{P}_0} J_\eta(P) \ge 0.$$

Let an estimated objective $\widehat{J}_\eta$ satisfy

$$\sup_{P \in \mathcal{A}_\eta} |\widehat{J}_\eta(P) - J_\eta(P)| \le \delta_\eta, \qquad \delta_\eta \ge 0,$$

and let $\widehat{P}_\eta \in \mathcal{A}_\eta$ satisfy

$$\widehat{J}_\eta(\widehat{P}_\eta) \le \inf_{P \in \mathcal{A}_\eta} \widehat{J}_\eta(P) + \epsilon_\eta, \qquad \epsilon_\eta \ge 0.$$

For every hard-feasible $P^\dagger$ with finite $F(P^\dagger)$,

$$\mathcal{D}(\widehat{P}_\eta) \le \frac{F(P^\dagger) + a_\eta + \epsilon_\eta + 2\delta_\eta}{\eta}. \tag{19}$$

If $\eta_n \to \infty$ and $a_{\eta_n} + \epsilon_{\eta_n} + 2\delta_{\eta_n} \to 0$, then $\widehat{P}_{\eta_n}$ converges weakly to the unique hard minimizer.

Proof. The uniform objective bound and approximate optimality give

$$J_\eta(\widehat{P}_\eta) \le \widehat{J}_\eta(\widehat{P}_\eta) + \delta_\eta \le \inf_{\mathcal{A}_\eta} \widehat{J}_\eta + \epsilon_\eta + \delta_\eta \le \inf_{\mathcal{A}_\eta} J_\eta + \epsilon_\eta + 2\delta_\eta$$
$$= \inf_{\mathcal{P}_0} J_\eta + a_\eta + \epsilon_\eta + 2\delta_\eta \le F(P^\dagger) + a_\eta + \epsilon_\eta + 2\delta_\eta.$$

Since $F \ge 0$, dropping $F(\widehat{P}_\eta)$ proves Equation (19). The same chain makes $\widehat{P}_{\eta_n}$ an approximate minimizer over $\mathcal{P}_0$, with gap at most $a_{\eta_n} + \epsilon_{\eta_n} + 2\delta_{\eta_n}$. Convergence follows from Theorem 5.4. □

The three error terms distinguish restriction to a model class, incomplete optimization, and estimation of the population objective. The corollary concerns admissible continuous path laws and one fixed population objective. The ideal finite-step model has its own path-entropy identity in Proposition C.19.

Finite-entropy feasibility identifies the target marginals for the penalty limit. Data processing gives

$$\mathrm{KL}(P^k \| R^k) \geq \mathrm{KL}(P_1^k \| R_1^k).$$

If R_1^k has a continuous density and the target is supported on a finite set A, then $R_1^k(A) = 0$ and $P_1^k(A) = 1$, so the right-hand side is infinite. Accordingly, the hard-limit analysis uses population or smoothed marginals with finite-entropy feasibility, whereas finite-sample training uses soft discrepancies against empirical samples.

This section develops the mathematical foundations of training and generation. An endpoint-preserving teacher constructs paired paths, and velocity regression turns these paths into a reference drift. Exponential stochastic steps and discrete entropy identities connect the continuous geometry to numerical paths. We then analyze the endpoint losses, fate probabilities, and tree readouts.

C.9 An Endpoint-Preserving Teacher

For a training endpoint coupling (z_0, z_1, k), let $\gamma_t = \mathrm{Exp}_{z_0}(t \,\mathrm{Log}_{z_0} z_1)$ and define the teacher by

$$\psi_\eta^k(t; z_0, z_1) = \mathrm{Exp}_{\gamma_t}[t(1-t)\,\mathrm{Sat}_L(\mathrm{PT}_{0\to\gamma_t} h_\eta(z_0, z_1, t, k))]\,, \tag{20}$$

where $\mathrm{Sat}_L(a) = a/\sqrt{1 + \|a\|_g^2/L^2}$ and $L > 0$. The correction field h_η is learned from training endpoints under the teacher objective in Equation (3). A route through a training-endpoint neighborhood can enter the correction direction before saturation.

Lemma C.15 (Teacher endpoints and correction radius)**.** *Let γ be the product geodesic above, with the correction field in Eq. (20) defined at every time. For the unclipped teacher,*

$$\psi_\eta^k(0; z_0, z_1) = z_0, \qquad \psi_\eta^k(1; z_0, z_1) = z_1, \qquad d_g(\psi_\eta^k(t), \gamma_t) \leq t(1-t)L \leq L/4.$$

Proof. Set $a_t = \mathrm{PT}_{0\to\gamma_t} h_\eta$ and $w_t = t(1-t)\,\mathrm{Sat}_L(a_t)$. Transport isometry and $\|\,\mathrm{Sat}_L(a_t)\|_g^2 = L^2\|a_t\|_g^2/(L^2 + \|a_t\|_g^2) \leq L^2$ give $\|w_t\|_g \leq t(1-t)L \leq L/4$. The Hadamard property gives $d_g(\gamma_t, \mathrm{Exp}_{\gamma_t} w_t) = \|w_t\|_g$, and $w_0 = w_1 = 0$ preserves both endpoints. □

This construction adjusts the intermediate path while preserving its endpoints. To specify the velocity regression target, assume that the correction field is differentiable in time and write $q_t = \sqrt{1 + \|a_t\|_g^2/L^2}$. Metric compatibility of the covariant derivative yields

$$\nabla_t\,\mathrm{Sat}_L(a_t) = \frac{\nabla_t a_t}{q_t} - \frac{a_t\langle a_t, \nabla_t a_t\rangle_g}{L^2 q_t^3}, \qquad \nabla_t w_t = (1-2t)\,\mathrm{Sat}_L(a_t) + t(1-t)\nabla_t\,\mathrm{Sat}_L(a_t).$$

The Levi–Civita horizontal–vertical decomposition of $T(T\mathcal{M})$ gives the full path derivative

$$\dot\psi_\eta^k(t) = d\,\mathrm{Exp}_{(\gamma_t, w_t)}[\dot\gamma_t, \nabla_t w_t].$$

It includes changes in the baseline, parallel transport, neural correction, and temporal envelope. The cell implementation computes this derivative by automatic differentiation and then applies the configured numerical safeguards. The warm-up endpoint coupling uses the optimal permutation under squared intrinsic distance between equal-sized batches, defining the teacher training distribution.

The coefficients and condition-sampling probabilities in Equation (3) are fixed throughout teacher fitting. The support pool, base metric, and minibatch pairing rule are independent of η; the gradient acts through the teacher position and its full time derivative. The sampling probabilities χ_k organize teacher fitting, whereas π_k weights the subsequent path-space objective.

For a specified teacher, let ϱ_t^k be its marginal and let b_*^k be the conditional mean velocity in Theorem C.16. Conditional Jensen's inequality gives

$$\sum_k \chi_k \int_0^1 \int_{\mathcal{M}} \|b_*^k(t,z)\|_g^2\, \varrho_t^k(dz)dt \leq \sum_k \chi_k \mathbb{E}_{\Gamma_k} \int_0^1 \|\dot\psi_\eta^k(t)\|_g^2 dt.$$

Moreover, by the definition of a marginal and Tonelli's theorem,

$$\mathbb{E}_{\Gamma_k} \int_0^1 V_{\mathcal{U}}(t, \psi_\eta^k(t), k) dt = \int_0^1 \int_{\mathcal{M}} V_{\mathcal{U}}(t, z, k)\, \varrho_t^k(dz) dt.$$

Thus the teacher objective controls its mean-field kinetic energy and expected marginal support cost.

C.10 Conditional-Expectation Projection of Teacher Velocities

The preceding lemma provides a family of endpoint-consistent paths. We now aggregate their velocities into a vector field that can be evaluated at any current state. This construction follows the projection principle of conditional flow matching (Chen & Lipman, 2024; Kapuśniak et al., 2024).

Theorem C.16 (Conditional-expectation decomposition of teacher velocity regression)**.** *Fix the teacher, endpoint coupling, and positive condition-sampling probabilities. Let $T \sim \mathrm{Unif}[0,1]$ be independent of the condition and endpoint samples, define $Y = \psi_\eta^k(T)$ and $U = \dot{\psi}_\eta^k(T)$, and assume $\mathbb{E}\|U\|_g^2 < \infty$. Using a measurable orthonormal frame, define $b^*(t,z,k) = \mathbb{E}[U \mid T = t, Y = z, k]$. Every square-integrable velocity field b then satisfies*

$$\mathcal{L}_{\mathrm{warm}}(b) = \mathbb{E}\|b(T,Y,k) - b^*(T,Y,k)\|_g^2 + C_\eta, \qquad C_\eta = \mathbb{E}\|U - b^*(T,Y,k)\|_g^2.$$

If the teacher paths are almost surely absolutely continuous and $\mathbb{E}[\int_0^1 \|\dot{\psi}_\eta^k(t)\|_g^2 dt \mid k] < \infty$, their marginals ϱ_t^k satisfy the weak continuity equation $\partial_t \varrho_t^k + \mathrm{div}_g(\varrho_t^k b^*) = 0$.

Proof. Let $\mathcal{G} = \sigma(T, Y, k)$ and write $\widehat{U} = F(Y)^{-1}U$ in an orthonormal frame. Set $\widehat{b}_* = \mathbb{E}[\widehat{U} \mid \mathcal{G}]$ and $b_* = F(Y)\widehat{b}_*$. The frame is fixed conditionally on $\mathcal{G}$, so orthonormal changes leave b_* unchanged. Conditional Jensen gives $\mathbb{E}\|\widehat{b}_*\|^2 \le \mathbb{E}\|\widehat{U}\|^2 < \infty$. For $\widehat{b} = F(Y)^{-1}b(T,Y,k)$, conditional orthogonality yields

$$\mathbb{E}\|\widehat{b} - \widehat{U}\|^2 = \mathbb{E}\|\widehat{b} - \widehat{b}_*\|^2 + \mathbb{E}\|\widehat{U} - \widehat{b}_*\|^2,$$

because $\mathbb{E}[\widehat{U} - \widehat{b}_* \mid \mathcal{G}] = 0$. The minimizer is unique up to $dt\, \varrho_t^k(dz)$-null sets.

Fix k and $f \in C_c^\infty(\mathcal{M})$. The chain rule along absolutely continuous teacher paths, followed by Fubini and conditioning on position, gives

$$\int f\, d\varrho_t^k - \int f\, d\varrho_s^k = \int_s^t \int \langle \nabla_g f(z), b_*^k(\tau, z)\rangle_g\, \varrho_\tau^k(dz)\, d\tau.$$

These operations are justified by Cauchy–Schwarz and the assumed finite kinetic energy, since $\nabla_g f$ is bounded. Equivalently, every $\varphi \in C_c^\infty((0,1) \times \mathcal{M})$ satisfies $\int_0^1 \int [\partial_t \varphi + \langle \nabla_g \varphi, b_*^k\rangle_g]\, d\varrho_t^k\, dt = 0$, the weak continuity equation. □

This result identifies the population regression target of the warm-up drift. With positive noise, the reference marginals $\varrho_t^{\mathrm{ref,k}}$ satisfy the weak Fokker–Planck equation

$$\frac{d}{dt} \int f\, d\varrho_t^{\mathrm{ref,k}} = \int \left[\langle b_{\mathrm{ref}}^k, \nabla_g f\rangle_g + \frac{\sigma^2}{2} \Delta_g f \right] d\varrho_t^{\mathrm{ref,k}}.$$

The diffusion term drives stochastic exploration, and the subsequent control adjusts the resulting distribution through endpoint losses. In practice, warm-up times are sampled uniformly from $(0.001, 0.999)$. Replacing the time law above by this sampling law gives the same conditional-expectation decomposition on the sampled time interval.

C.11 From Teacher Regression to a Feasible Control Budget

The conditional mean velocity transports the teacher marginals by a continuity equation. To compare these marginals with a noisy reference, we construct an auxiliary diffusion with the same marginals. The velocity–score correction used in this construction is standard in stochastic interpolants (Albergo et al., 2025); here it is expressed using the base Riemannian metric.

Assumption C.17 (Regularity for the teacher-to-control comparison)**.** For each k, the teacher marginal has a strictly positive density $\varrho_t^k(dz) = p_t^k(z)\,d\mathrm{vol}_g(z)$, with $p^k \in C^{1,2}([0,1]\times\mathcal{M})$ and $\varrho_0^k = \nu_0^k$. The marginal flow is narrowly continuous with uniformly bounded second geodesic moments and satisfies the continuity equation of Theorem C.16. Define

$$b_\dagger^k(t,z) = b_*^k(t,z) + \frac{\sigma^2}{2}\nabla_g \log p_t^k(z). \tag{21}$$

The diffusions with drifts b_{ref}^k and $b_\dagger^k$, common diffusion operator $\sigma^2\Delta_g/2$, and common initial law ν_0^k satisfy Assumption C.3. Their marginal flows have uniformly bounded second geodesic moments. The Fokker–Planck equation for $b_\dagger^k$ has a unique narrowly continuous probability solution in this moment class. The quantities $\varepsilon_{\mathrm{reg}}$ and $\mathcal{I}_{\mathrm{T}}$ in Proposition 5.3 are finite.

Proof of Proposition 5.3. Fix k. In the weak sense, since $p_t^k\nabla_g \log p_t^k = \nabla_g p_t^k$, Equation (21) gives

$$\begin{aligned} -\operatorname{div}_g(p_t^k b_\dagger^k) + \frac{\sigma^2}{2}\Delta_g p_t^k &= -\operatorname{div}_g(p_t^k b_*^k) - \frac{\sigma^2}{2}\operatorname{div}_g(\nabla_g p_t^k) + \frac{\sigma^2}{2}\Delta_g p_t^k \\ &= -\operatorname{div}_g(p_t^k b_*^k) = \partial_t p_t^k. \end{aligned}$$

Thus the teacher marginal flow solves the Fokker–Planck equation for $b_\dagger^k$. The diffusion with this drift starts from $\nu_0^k = \varrho_0^k$, so uniqueness of that equation identifies all its marginals with ϱ_t^k. Denote its path law by $P^{\dagger,k}$ and set $P^\dagger(k, d\omega) = \pi_k P^{\dagger,k}(d\omega)$. Its residual drift relative to the reference is

$$u_\dagger^k = b_*^k - b_{\mathrm{ref}}^k + \frac{\sigma^2}{2}\nabla_g \log p_t^k.$$

Applying Theorem 5.2 and using the identified marginals yields

$$\mathrm{KL}(P^\dagger\|\bar{R}) = \frac{1}{2\sigma^2}\sum_k \pi_k \int_0^1 \int_{\mathcal{M}} \left\| b_*^k - b_{\mathrm{ref}}^k + \frac{\sigma^2}{2}\nabla_g \log p_t^k \right\|_g^2 \varrho_t^k(dz)dt. \tag{22}$$

Minkowski's inequality in the weighted space $L^2(\pi_k\,dt\,\varrho_t^k(dz); T\mathcal{M})$ gives

$$\left(\sum_k \pi_k \int_0^1 \int_{\mathcal{M}} \|u_\dagger^k\|_g^2\,\varrho_t^k(dz)dt\right)^{1/2} \le \sqrt{\varepsilon_{\mathrm{reg}}} + \frac{\sigma^2}{2}\sqrt{\mathcal{I}_{\mathrm{T}}}.$$

Squaring and dividing by $2\sigma^2$ proves $\mathrm{KL}(P^\dagger\|\bar{R}) \le B_\sigma$. For every t, apply data processing to $(k,\omega)\mapsto(k,\omega_t)$. Because the label masses agree,

$$\sum_k \pi_k\,\mathrm{KL}(\varrho_t^k\|R_t^k) \le \mathrm{KL}(P^\dagger\|\bar{R}) \le B_\sigma.$$

This proves the uniform marginal bound.

If $D_\#\varrho_1^k = \mu_1^k$ for every k, then $P^\dagger$ satisfies the hard marginal constraints. For the nonnegative state potential, Tonelli's theorem expresses the expected path cost through its marginals:

$$\mathbb{E}_{P^\dagger} I_V = \sum_k \pi_k \int_0^1 \int_{\mathcal{M}} V_{\mathcal{U}}(t,z,k)\,\varrho_t^k(dz)dt = \mathcal{V}_{\mathrm{T}}.$$

Evaluating F at this feasible law establishes $\inf_{P\in\mathcal{H}} F(P) \le F(P^\dagger) \le \beta B_\sigma + \lambda_V \mathcal{V}_{\mathrm{T}}$. □

Define the conditional population regression loss by

$$\mathcal{L}_{\mathrm{warm}}^k(b) := \mathbb{E}_{\Gamma_k,\,T\sim\mathrm{Unif}[0,1]}\|b(T,\psi_\eta^k(T)) - \dot{\psi}_\eta^k(T)\|_g^2.$$

Theorem C.16 then gives

$$\varepsilon_{\mathrm{reg}} = \sum_k \pi_k\left[\mathcal{L}_{\mathrm{warm}}^k(b_{\mathrm{ref}}^k) - \mathcal{L}_{\mathrm{warm}}^k(b_*^k)\right].$$

These losses use the full time interval $[0,1]$. When warm-up samples only $(0.001, 0.999)$, the definition of $\varepsilon_{\mathrm{reg}}$ additionally includes the two omitted endpoint intervals. The score in Equation (21) is used only to construct the comparison law; reference regression and residual-control training do not require a score network. For a fixed teacher and noise level, reducing the excess regression loss tightens the bound B_σ, while $\mathcal{I}_{\mathrm{T}}$ quantifies the density correction associated with adding diffusion.

C.12 Reference Support Pools and State Costs

The support cost incorporates training-endpoint geometry into teacher and path objectives. The frozen pool uses only training endpoints. Its radial basis function (RBF) form is

$$V_{\mathrm{RBF}}(z) = 1 - \frac{1}{|\mathcal{S}|J} \sum_{s\in\mathcal{S}} \sum_{j=1}^{J} \exp[-d_g(z, y_j)^2/(2h_s^2)], \qquad h_s > 0.$$

Each kernel value lies in $(0, 1]$, so $0 \le V_{\mathrm{RBF}} \le 1$ and the associated path potential satisfies the boundedness condition of Corollary C.7. The local-scale form uses $e_j(z) = d_g(z, y_j)^2/(2\ell_j^2)$, with $\ell_j > 0$ and $\tau > 0$, and defines a smooth minimum over these scaled distances:

$$V_{\mathrm{local}}(z) = -\tau \log \left[\frac{1}{J} \sum_{j=1}^{J} e^{-e_j(z)/\tau} \right].$$

For $m(z) = \min_j e_j(z)$, the bounds $J^{-1}e^{-m/\tau} \le J^{-1} \sum_j e^{-e_j/\tau} \le e^{-m/\tau}$ together with Jensen's inequality yield the following lower and upper bounds:

$$0 \le m(z) \le V_{\mathrm{local}}(z) \le \min \left\{ m(z) + \tau \log J,\ \frac{1}{J} \sum_j e_j(z) \right\}.$$

A finite reference pool and positive scales also give $V_{\mathrm{local}}(z) \le C(1 + d_g(o, z)^2)$. Thus, the second-moment estimate in Lemma C.5 ensures that the potential is integrable under the reference law. A conditional Jensen bound then ensures that $\log Z_V$ is integrable under the initial law. Squared distance is smooth on a Hadamard manifold, so both costs are continuous and nonnegative. If $\omega_n \to \omega$ uniformly, pointwise continuity and Fatou's lemma give

$$I_V(\omega) = \int_0^1 \lim_n V(\omega_n(t))dt \le \liminf_n \int_0^1 V(\omega_n(t))dt = \liminf_n I_V(\omega_n),$$

establishing lower semicontinuity for the penalty limit. Tahoe uses a frozen pool per condition.

C.13 Shared Conditional Fields, Adapters, and Reference Freezing

After fitting the reference, we compute the control in the tangent space at the origin and parallel-transport it to the current state. Let $r(z)$ denote $\mathrm{Log}_0 z$ normalized by a fixed scale, and let $c_k(z)$ collect the training terminal prototypes, scales, and priors. The parameterization is

$$u_\theta^k(t, z) = \mathrm{PT}_{0\to z}[f_\theta(r(z), t) + a_k(t, z)\{h_\theta(r(z), t, c_k(z)) + A_k r(z) + n_k(r(z))\}].$$

Here f_θ is the shared backbone, h_θ is the shared conditional residual, $A_k = B_k C_k$ is a low-rank adapter, and n_k is a nonlinear adapter. The adapter type is selected for each task. Attention over multiple prototypes captures local heterogeneity within a condition, and a_k controls the strength of the conditional residual over time and state. Appendix F.3 specifies the networks and fixed scales.

Writing the tangent vector in brackets as $v_\theta^k(t, z)$, the isometry of parallel transport gives

$$\|u_\theta^k(t, z)\|_{g_z}^2 = \|v_\theta^k(t, z)\|_{g_0}^2 = 4\|(v_\theta^k)_H\|^2 + w_E\|(v_\theta^k)_E\|^2.$$

This equality relates the conditional network output to the intrinsic control energy. The warm-up drift is copied into the frozen b_{ref}, the control is initialized to zero, and the total drift is $b_{\mathrm{ref}} + u_\theta$. State derivatives of the reference remain active during stochastic unrolling.

C.14 Generator Consistency of Exponential Stochastic Steps

The following result connects the geometric fields above to the continuous generator in Appendix C.

Proposition C.18 (Local weak consistency). *Consider the unclipped update in Eq. (6) on a complete smooth manifold, with $b = b_{\mathrm{ref}} + u$, $FF^\top = g^{-1}$, and $\sigma > 0$. Assume that b is a deterministic Markov field bounded on the compact time and state sets under consideration, and*

that each Gaussian increment is independent of the past. For $f \in C_c^\infty(\mathcal{M})$, locally uniformly,

$$\frac{\mathbb{E}[f(Z_{n+1}) \mid Z_n = z] - f(z)}{h} \longrightarrow \langle b(t,z), \nabla_g f(z)\rangle_g + \frac{\sigma^2}{2}\Delta_g f(z).$$

Proof. In an orthonormal frame, write $\delta = hb(t,z) + \sigma\sqrt{h}\sum_i \xi_i e_i(z)$. On a compact state set K, smoothness on a fixed tangent disk gives the uniform Taylor expansion

$$f(\mathrm{Exp}_z\,\delta) = f(z) + df_z(\delta) + \tfrac{1}{2}\,\mathrm{Hess}_g\, f_z(\delta,\delta) + O_K(\|\delta\|_g^3), \qquad \|\delta\|_g \le \delta_0.$$

The Gaussian moments satisfy $\mathbb{E}[df_z(\delta)] = h\, df_z(b)$,

$$\mathbb{E}[\mathrm{Hess}_g\, f_z(\delta,\delta)] = \sigma^2 h\Delta_g f(z) + h^2\,\mathrm{Hess}_g\, f_z(b,b), \qquad \mathbb{E}\|\delta\|_g^3 = O_K(h^{3/2}).$$

Boundedness of b implies $\{\|\delta\|_g > \delta_0\} \subseteq \{\|\xi\| > \delta_0/(2\sigma\sqrt{h})\}$ for small h. The tail probability and its first two moments are $O_K(e^{-c_K/h})$; boundedness of f and the local Taylor coefficients therefore permits full expectations with error $o_K(h)$. Consequently, $\mathbb{E}[f(Z_{n+1}) \mid Z_n = z] - f(z) = h\mathcal{L}_t^{k,u} f(z) + O_K(h^{3/2}) + o_K(h)$, and division by h proves local uniform convergence. □

This calculation verifies the noise and metric conventions of the geodesic stochastic update (Hsu, 2002; Thornton et al., 2022). The orthonormal trace of the Hessian produces the Laplace–Beltrami operator, so both diffusion and drift follow the same underlying geometry. The common-map formulation in the next proposition covers discrete transitions with numerical safeguards.

C.15 Discrete Control Energy and Path Relative Entropy

The entropy identity for continuous paths follows from the Girsanov change of measure. In our Hadamard geometry, the relative entropy of finite-step chains can also be computed directly from the corresponding tangent-space Gaussian densities.

Proposition C.19 (Entropy identity for ideal discrete chains and bound under safeguards)**.** *On a Hadamard manifold, construct two unclipped exponential Gaussian Markov chains with the same initial joint law, step size $h > 0$, noise level $\sigma > 0$, and orthonormal frame F. Their drifts are the fixed measurable fields $b_{\mathrm{ref}} + u$ and b_{ref}, respectively, and each step uses an independent standard Gaussian variable. If the expected energy is finite, the relative entropy of the labeled path chains is*

$$\mathrm{KL}(P_h \| R_h) = \sum_k \pi_k \mathbb{E}_{P_h^k} \sum_{n=0}^{N-1} \frac{h\|u^k(t_n, Z_n)\|_g^2}{2\sigma^2}. \tag{23}$$

If, conditionally on the current (t_n, z, k), both tangent Gaussian laws pass through the same measurable deterministic map $T_{n,z,k}$, which may include the exponential map and numerical safeguards, the relative entropy of the resulting safeguarded chains is bounded above by this energy evaluated along the safeguarded controlled chain.

Proof. Fix (n, z, k) and define the orthonormal-coordinate means by

$$m_0 = hF(z)^{-1} b_{\mathrm{ref}}^k(t_n, z), \qquad \delta = hF(z)^{-1} u^k(t_n, z).$$

The tangent laws $G_0 = \mathcal{N}(m_0, \sigma^2 hI)$ and $G_1 = \mathcal{N}(m_0 + \delta, \sigma^2 hI)$ have

$$\log \frac{dG_1}{dG_0}(v) = \frac{\langle \delta, v - m_0\rangle}{\sigma^2 h} - \frac{\|\delta\|^2}{2\sigma^2 h}, \qquad \mathrm{KL}(G_1 \| G_0) = \frac{h\|u^k(t_n, z)\|_g^2}{2\sigma^2}.$$

Averaging under G_1 gives the second identity. The invertible map $v \mapsto \mathrm{Exp}_z(F(z)v)$ preserves this divergence for $K_n^{P,k}$ and $K_n^{R,k}$. With the common initial law, the Markov entropy chain rule yields

$$\mathrm{KL}(P_h^k \| R_h^k) = \sum_{n=0}^{N-1} \mathbb{E}_{P_h^k}\, \mathrm{KL}(K_n^{P,k}(Z_n, \cdot) \| K_n^{R,k}(Z_n, \cdot)).$$

Weighting by π_k proves Equation (23). For common safeguards, data processing bounds divergence by $\mathrm{KL}(G_1 \| G_0)$; apply the chain rule along the safeguarded controlled chain. □

The proposition specifies both the sampling law for the energy and the transition rule: expectations are taken along the controlled chain, invertible exponential transitions give the exact entropy, and common safeguards give the data processing bound. Antithetic sampling preserves each path's marginal law; Monte Carlo groups account for dependence across paths.

C.16 Gradients Through Stochastic Unrolling and Training Scales

For a local step $Z_{n+1} = T_\theta(Z_n, \xi_n)$ and $J_n = \partial_\theta Z_n$, the chain rule gives

$$J_{n+1} = \partial_z T_\theta(Z_n, \xi_n) J_n + \partial_\theta T_\theta(Z_n, \xi_n).$$

Before safeguards, differentiating the tangent increment $v_n = h[b_{\mathrm{ref}}(t_n, Z_n) + u_\theta(t_n, Z_n)] + \sigma\sqrt{h}F(Z_n)\xi_n$ gives

$$\begin{aligned}\partial_\theta v_n &= h[\partial_z b_{\mathrm{ref}} + \partial_z u_\theta]J_n + h\,\partial_\theta u_\theta + \sigma\sqrt{h}\,\{dF(Z_n)[J_n]\}\xi_n,\\ J_{n+1} &= d_1 \operatorname{Exp}_{(Z_n, v_n)}[J_n] + d_2 \operatorname{Exp}_{(Z_n, v_n)}[\partial_\theta v_n].\end{aligned}$$

Here $\partial_\theta u_\theta$ denotes the direct parameter derivative at a fixed state. Freezing the reference sets its direct parameter derivative to zero while retaining $\partial_z b_{\mathrm{ref}} J_n$. When these derivatives exist and are locally dominated by an integrable random variable, differentiation and expectation can be interchanged, yielding the pathwise gradient of the stochastic unrolling objective.

The cell setting uses $\sigma = \sigma_{\mathrm{total}}/\sqrt{D}$ and $\sigma_{\mathrm{total}} = 0.05$. Since $F^\top g F = I$,

$$\mathbb{E}\|\sigma\sqrt{h}F(Z_n)\xi_n\|_g^2 = \sigma^2 h \mathbb{E}\|\xi_n\|^2 = \sigma_{\mathrm{total}}^2 h.$$

The total noise magnitude is therefore normalized across dimensions. The control weight $\beta = 0.01\sigma^2$ gives the integrated cost $\beta/(2\sigma^2)\|u\|_g^2 = 0.005\|u\|_g^2$.

C.17 Teacher Regularity and Quadratic Variation of Stochastic Paths

Teacher velocities determine the regression target for the drift, while quadratic variation describes local fluctuations of the stochastic diffusion. For $f \in C_c^\infty(\mathcal{M})$, Itô's formula gives

$$f(Z_t) - f(Z_0) = \int_0^t \mathcal{L}_s f(Z_s) ds + M_t^f, \qquad M_t^f = \sigma \int_0^t df_{Z_s}(F(Z_s) dW_s).$$

The finite-variation term has zero quadratic variation; an orthonormal frame obeys $\sum_i [df(e_i)]^2 = \|\nabla_g f\|_g^2$. Hence

$$[f(Z)]_t = [M^f]_t = \sigma^2 \int_0^t \|\nabla_g f(Z_s)\|_g^2 ds.$$

Stop the process within a relatively compact coordinate neighborhood and choose a compactly supported test function that agrees with a coordinate function there. Its gradient norm has a positive lower bound on the neighborhood, so the quadratic variation has a strictly positive increment over every interval of positive length before the stopping time. For a continuous function of finite variation, the sum of squared increments along a partition satisfies

$$\sum_i |\Delta f_i|^2 \le \max_i |\Delta f_i| \sum_i |\Delta f_i| \longrightarrow 0.$$

These identities characterize the local roughness of positive-noise sample paths and clarify the role of smooth teacher paths as targets for drift learning (Hsu, 2002).

C.18 Population Endpoint Discrepancies and Finite-Batch Losses

Appendix C.8 uses a nonnegative, weakly lower semicontinuous discrepancy that identifies the target marginal. We verify these properties for a concrete choice on the full observed principal component (PC) space $\mathbb{R}^D$, then describe its empirical estimator and training gradients. For fixed $a_\ell, s_\ell > 0$, take

$$k(x, y) = \sum_{\ell=1}^{L} a_\ell \exp[-\|x - y\|^2/(2s_\ell^2)].$$

Proposition C.20 (Continuity and identification by a full-dimensional kernel discrepancy)**.** *The population squared maximum mean discrepancy (MMD^2),* $\mathrm{MMD}_k^2(\rho,\mu)$*, is nonnegative, jointly continuous under weak convergence of* (ρ,μ)*, and zero exactly when* $\rho=\mu$*. For fixed* $\pi_k>0$ *and* $\alpha\in[0,1]$*, the risk* $\mathcal{R}_\alpha$ *formed from conditional MMD scores satisfies the nonnegativity, weak lower semicontinuity, and zero-set conditions of Theorem 5.4.*

Proof. The bounded positive-definite kernel has a mean embedding for every probability measure, and the squared reproducing kernel Hilbert space (RKHS) norm gives (Gretton et al., 2012)

$$\mathrm{MMD}_k^2(\rho,\mu)=\int k\,d(\rho\otimes\rho)+\int k\,d(\mu\otimes\mu)-2\int k\,d(\rho\otimes\mu)\geq 0.$$

Weak convergence passes to product measures, so bounded continuity of k makes all three terms jointly continuous. For identification, use the Gaussian Fourier identity

$$k(x,y)=\int e^{i\omega^\top(x-y)}q(\omega)d\omega,\qquad q(\omega)=\sum_{\ell=1}^{L}\frac{a_\ell s_\ell^D}{(2\pi)^{D/2}}e^{-s_\ell^2\|\omega\|^2/2}>0.$$

With $\widehat{\rho}(\omega)=\int e^{i\omega^\top x}\rho(dx)$, integrability of q permits Fubini and yields $\mathrm{MMD}_k^2(\rho,\mu)=\int|\widehat{\rho}-\widehat{\mu}|^2q\,d\omega$. A zero value forces equality of the characteristic functions almost everywhere, then everywhere by continuity; uniqueness gives $\rho=\mu$. The converse is immediate.

For fixed positive label masses,

$$\int f\,d\rho_1^k=\frac{1}{\pi_k}\int \mathbf{1}_{\{j=k\}}f(D(\omega_1))\,P(dj,d\omega),\qquad f\in C_b(\mathbb{R}^D).$$

The discrete label space, continuous terminal evaluation, and continuous decoder make the integrand bounded and continuous. Thus $P\mapsto\rho_1^k$, every $e_k(P)$, and their finite mean–maximum risk are continuous. Nonnegativity and $\pi_k>0$ force every $e_k=0$ when the risk vanishes for $\alpha<1$; at $\alpha=1$, the zero maximum gives the same conclusion. □

Boundedness of the kernel makes these conclusions valid for all probability marginals. Adding nonnegative, weakly lower semicontinuous population terms that vanish at the target marginal preserves the zero set and lower semicontinuity when the MMD coefficient is positive.

Training uses the sum kernel with fixed bandwidths $\{0.01,0.1,1,10,100\}$ and computes the empirical V-statistic including diagonal terms. To describe its finite-sample bias, let $X_1,\dots,X_n$ be independent and identically distributed samples from ρ, and let $Y_1,\dots,Y_m$ be independent and identically distributed samples from μ, independently of the first group. Write $A_\rho=\mathbb{E}k(X,X')$ and $D_\rho=\mathbb{E}k(X,X)$, where X' is an independent copy, and define A_μ,D_μ analogously. Then

$$\mathbb{E}\left[\frac{1}{n^2}\sum_{i,i'}k(X_i,X_{i'})\right]=\left(1-\frac{1}{n}\right)A_\rho+\frac{1}{n}D_\rho,\qquad \mathbb{E}\left[\frac{1}{nm}\sum_{i,j}k(X_i,Y_j)\right]=\mathbb{E}k(X,Y).$$

Combining these expectations gives the finite-sample bias of the empirical V-statistic:

$$\mathbb{E}\widehat{\mathrm{MMD}}_k^{\,2}-\mathrm{MMD}_k^2(\rho,\mu)=\frac{D_\rho-A_\rho}{n}+\frac{D_\mu-A_\mu}{m}.\tag{24}$$

Equation (24) gives the bias under independent sampling. With shared sources or antithetic noise, the expectations of the corresponding kernel terms depend on their joint sampling law. We therefore organize error estimation by independent sampling groups.

For uniform positive weights $a_i=1/n$ and $b_j=1/m$, let $\Pi(a,b)=\{\Gamma\geq 0:\Gamma\mathbf{1}=a,\Gamma^\top\mathbf{1}=b\}$. Training uses regularized optimal transport (OT) and its debiased discrepancy (Feydy et al., 2019)

$$\mathrm{OT}_\varepsilon(a,b;C)=\min_{\Gamma\in\Pi(a,b)}\{\langle C,\Gamma\rangle+\varepsilon\,\mathrm{KL}(\Gamma\|a\otimes b)\},$$
$$S_\varepsilon(\widehat{\rho},\widehat{\mu})=\mathrm{OT}_\varepsilon(\widehat{\rho},\widehat{\mu})-\tfrac{1}{2}\,\mathrm{OT}_\varepsilon(\widehat{\rho},\widehat{\rho})-\tfrac{1}{2}\,\mathrm{OT}_\varepsilon(\widehat{\mu},\widehat{\mu}).$$

Here $C_{ij}=\|x_i-y_j\|^p$, $p\in\{1,2\}$, $\mathrm{KL}(\Gamma\|a\otimes b)=\sum_{ij}\Gamma_{ij}\log[\Gamma_{ij}/(a_ib_j)]$, and $0\log 0=0$. For $\varepsilon>0$, the feasible set is compact and the objective is continuous and strictly convex. The optimal plan $\Gamma^*(C)$ therefore exists and is unique.

Let $C_t = C + tH$ and $\Gamma_t = \Gamma^*(C_t)$. For $t > 0$, optimality gives

$$\langle \Gamma_t, H \rangle \le \frac{\mathrm{OT}_\varepsilon(C_t) - \mathrm{OT}_\varepsilon(C)}{t} \le \langle \Gamma_0, H \rangle.$$

By compactness, every accumulation point of Γ_t minimizes the objective at C, and uniqueness implies $\Gamma_t \to \Gamma_0$. The analogous argument for $t < 0$ gives the envelope derivative

$$D_C \,\mathrm{OT}_\varepsilon[H] = \langle \Gamma^*, H \rangle.$$

The gradient of the debiased value combines the derivatives of the cross term and both self terms; each self term includes derivatives with respect to both sample arguments. Replacing the optimal plan by an approximate plan $\widehat{\Gamma}$ gives the direct directional-gradient bound

$$|\langle \widehat{\Gamma} - \Gamma^*, H \rangle| \le \|\widehat{\Gamma} - \Gamma^*\|_1 \|H\|_\infty.$$

For $p = 1$, coincident points can be handled with a valid subgradient of the distance norm.

In the implementation, ε is determined by the median positive distance in the same target batch. The cross term and both self terms share this value, which is fixed with respect to the current model parameters. A finite number of log-domain iterations produces an approximate plan, and clipping at zero gives the batch loss used in practice. These computations provide statistical and numerical approximations to the population discrepancy. Evaluation W_1 and W_2 are computed separately using the unregularized empirical OT definitions in Appendix F.5.

C.19 Terminal Density Estimation and Clone Readouts

A probabilistic readout fitted to training terminal cells converts generated endpoints into source-conditioned fate distributions during training and evaluation. Let F be the number of retained fate labels and let $\tilde{y}$ be standardized using training terminal statistics. Each class has a nonempty reference set $\mathcal{A}_f$ and a positive prior ϖ_f, with $\sum_f \varpi_f = 1$. For a shared kernel density estimation (KDE) bandwidth $h_{\mathrm{KDE}} > 0$, define

$$\ell_f(y) = \frac{1}{|\mathcal{A}_f|} \sum_{a \in \mathcal{A}_f} e^{-\|\tilde{y} - \tilde{a}\|^2/(2h_{\mathrm{KDE}}^2)}, \qquad r_f(y) = \frac{\varpi_f \ell_f(y)}{\sum_{j=1}^F \varpi_j \ell_j(y)}.$$

All kernel values are positive, so the denominator is strictly positive, $r_f(y) > 0$, and $\sum_f r_f(y) = 1$. The Gaussian normalization constant for the shared bandwidth cancels in the ratio. The bandwidth is determined from within-class neighborhoods of training samples, with at most 256 training references per class. Priors are computed from training terminal counts. The readout is frozen before bridge optimization; derivatives with respect to generated endpoints remain active during training for the fate objective in Appendix F.8.

Training endpoints define a frozen kernel density estimation (KDE) label readout $r_f(y)$. With fixed source-conditioned weights $q(k \mid x)$, generated endpoints yield

$$p_\theta(f \mid x) = \sum_{k=1}^K q(k \mid x)\, \mathbb{E}\big[r_f(D(Z_1^k)) \mid Z_0 = E(x)\big]. \tag{25}$$

Proposition C.21 (Probability conservation and Monte Carlo error)**.** *If $q(\cdot \mid x)$ and $r(\cdot)$ are probability vectors, then Eq. (25) is also a probability vector. For a fixed source x, sampling endpoints from the fixed discrete controlled model and averaging their readouts gives an unbiased estimate of that model's label-transition probabilities. If there are m independent sampling groups, each with a weighted readout in $[0, 1]$ whose expectation is the label probability, then*

$$\Pr\left(\max_{1 \le f \le F} |\widehat{p}_f(x) - p_f(x)| \ge \epsilon\right) \le 2F \exp(-2m\epsilon^2). \tag{26}$$

Proof. Convex averaging preserves nonnegativity and $\sum_f p_f(x) = 1$, because both the weights $q(\cdot \mid x)$ and each readout $r(\cdot)$ are normalized. For group j, the weighted readout $Y_{j,f} = \sum_k q(k \mid x) r_f(D(Z_{N,j}^{x,k}))$ lies in $[0, 1]$ and has expectation $p_f(x)$, so its sample mean is unbiased. Hoeffding's inequality for independent groups (Hoeffding, 1963) and a union bound over labels give Equation (26).

For antithetic sampling, replace each $r_f(D(Z_{N,j}^{x,k}))$ by the average of the positive- and negative-noise readouts. Both noise sequences have the correct Gaussian law, so each path retains the model's marginal law and the group's range and expectation are preserved. Independent seeds across groups give the same bound, with $m = M/2$ for M paired paths per condition. □

For the source set S_c of a clone, let $p_c = |S_c|^{-1} \sum_{x \in S_c} p(x)$. Finite averaging preserves probability conservation and unbiasedness. If every group simulates all source cells of the clone, then $Y_{j,c,f} = |S_c|^{-1} \sum_{x \in S_c} Y_{j,x,f} \in [0, 1]$. Under independence across groups, Eq. (26) also applies to the clone average. These results quantify readout and sampling errors for a fixed generative model. Independent descendant labels evaluate biological validity, while path identification from snapshots has been studied separately (Weinreb et al., 2018; Lavenant et al., 2024).

C.20 Structural Guarantee of Rooted Parent-Pointer Decoding

The cell readout preserves probabilities; the Tree readout preserves trees via ordered parent pointers.

Proposition C.22 (Ordered parent pointers form a tree)**.** *Let the valid nodes be* $\{0, \dots, n-1\}$*, with* $n \geq 1$ *and root* 0*. For each* $i > 0$*, choose exactly one parent* $p(i) \in \{0, \dots, i-1\}$ *and output the undirected edge* $\{i, p(i)\}$*. The resulting graph is a tree with* $n - 1$ *edges.*

Proof. The edge $\{i, p(i)\}$ has maximum index i, so the $n - 1$ edges are distinct. Repeated parent selection strictly decreases the index and must reach the sole parentless node 0; hence the graph is connected. A connected graph with n vertices and $n - 1$ edges is a tree, including $n = 1$. □

With a valid-node mask and a fixed node order, choosing the highest-scoring allowed parent satisfies every condition of the proposition, converting the continuous transport outputs into valid discrete objects with a guaranteed tree structure.

D Supplementary Experiments and Ablations

We provide the remaining result tables, distribution comparisons, and reference-diagnostic details. Section 6.3 and the subsequent ablations interpret the results; Appendix F specifies task settings, data access, and baseline adaptations.

D.1 Supplementary Tree Generation

Figure 6 arranges deep, bushy, and hub-dominated target families by column, with observed trees followed by RDSB and HyRef-SB samples. The comparison uses the rooted tree task of Appendix F.3: source nodes are retained, additional nodes occupy masked slots, and conditional transport is followed by structural decoding. The parent-pointer construction guarantees a connected acyclic tree output (Proposition C.22).

The metrics in Table 2 assess complementary properties. Degree discrepancies concern local branching, spectral discrepancies concern global connectivity, and paired Edge F1 uses fixed node correspondence. These criteria distinguish distributional structure from recovery of individual edges. Both methods share the frozen representation, structural decoder, and graph supervision described in Appendix F.4.

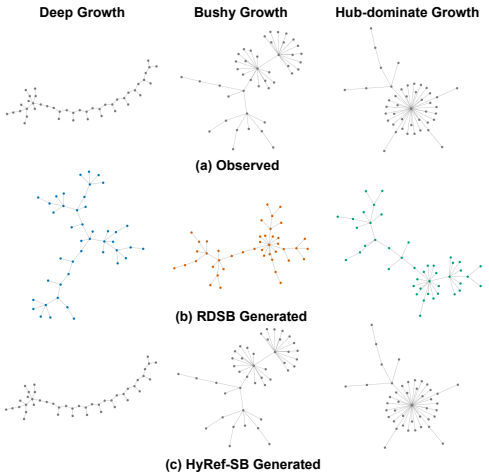

Figure 6. Tree generation. Columns show deep, bushy, and hub-dominated growth (left to right). Rows show (a) observed target trees, (b) RDSB samples, and (c) HyRef-SB samples.

D.2 Supplementary cellular results

Tables 9, 10, and 11 supplement the main cellular results using the common protocol of Appendix F. Intermediate observations are reserved for evaluation. Conditional paths appear in Figures 3 and 4; Figures 7 and 8 compare observed and generated populations.

D.2.1 Additional basic cell tasks

Table 9. Mouse Hematopoiesis distribution errors (↓). MMD^2 values use units of 10^{-2}; bold indicates the lowest error for each evaluation time and metric.

	Intermediate: $t = 0.5$			Terminal: $t = 1$		
Method	MMD^2	W_1	W_2	MMD^2	W_1	W_2
FreeBridge	9.435	0.436	0.508	2.132	0.154	0.209
Departures-PC	12.286	0.417	0.514	5.863	0.365	0.429
BranchSBM	11.466	0.393	**0.474**	4.687	0.226	0.257
HyRef-SB	**8.355**	**0.387**	0.476	**1.084**	**0.150**	**0.196**

Table 10. Veres errors (↓): held-out times t_1–t_6 and endpoint t_7. MMD^2 in 10^{-2}; W_2 uses exact empirical optimal transport. Bold marks the lowest error per metric and time.

	Held-out intermediate times						Terminal
Method	t_1	t_2	t_3	t_4	t_5	t_6	t_7
				Veres: MMD^2 ($\times 10^{-2}$)			
FreeBridge	13.674	14.798	12.404	**8.749**	7.287	4.910	3.907
Departures-PC	15.459	18.618	14.334	10.505	**6.605**	**2.612**	5.650
BranchSBM	18.239	24.043	19.781	18.209	14.107	9.525	5.621
HyRef-SB	**11.842**	**12.547**	**10.121**	8.833	8.803	4.759	**2.173**
				Veres: W_1			
FreeBridge	12.521	13.937	12.828	12.527	12.045	10.891	10.288
Departures-PC	12.522	13.939	12.094	11.841	10.995	10.347	13.981
BranchSBM	12.506	13.796	11.805	11.842	11.099	10.619	9.956
HyRef-SB	**10.120**	**7.490**	**11.069**	**11.014**	**10.254**	**8.135**	**6.408**
				Veres: W_2 (exact)			
FreeBridge	13.540	15.073	13.585	13.523	13.131	11.934	11.366
Departures-PC	13.608	15.102	12.828	12.745	11.703	10.906	14.739
BranchSBM	13.500	14.810	12.440	12.674	11.941	11.504	10.837
HyRef-SB	**12.379**	**14.750**	**11.307**	**12.469**	**11.077**	**9.295**	**8.355**

D.2.2 Advanced cell tasks

Table 11. Tahoe and Gastrulation distribution errors (↓). MMD^2 values are in units of 10^{-2}; intermediate scores average evaluation times, and W_2 uses exact empirical optimal transport. Bold marks the lowest error among all compared methods within each task and metric.

	Terminal			Intermediate mean		
Method	MMD^2	W_1	W_2	MMD^2	W_1	W_2
			Tahoe			
FreeBridge	8.588	2.330	20.133	11.874	1.370	22.028
Departures-PC	4.492	**0.675**	**14.407**	4.329	0.649	14.183
BranchSBM	4.409	4.098	14.627	4.041	4.599	15.053
HyRef-SB	**4.030**	0.676	16.002	**2.120**	**0.483**	**10.868**

Continued on next page

Table 11. Tahoe and Gastrulation distribution errors (continued).

Method	Terminal			Intermediate mean		
	MMD^2	W_1	W_2	MMD^2	W_1	W_2
			Gastrulation			
FreeBridge	8.193	2.251	17.221	5.103	2.191	18.570
Departures-PC	2.298	1.208	17.298	4.936	1.539	15.887
BranchSBM	1.157	3.921	15.118	1.266	2.738	15.162
HyRef-SB	**0.735**	**0.639**	**10.124**	**0.709**	**0.409**	**12.347**

D.3 Qualitative cellular comparisons

Figure 7 and Figure 8 compare held-out and terminal population coverage under the same projection. The main text discusses the trends together with the conditional paths.

D.4 Reference-calibration diagnostic for Proposition 5.3

This controlled experiment connects reference regression to observable marginal error and the energy of a fitted residual control. It uses five seeds (801–805), two teacher widths, and five reference checkpoints: 50 configurations in total. The teacher and noise level are fixed within each width, and only the fitted reference changes across checkpoints. Appendix F.7 specifies the full recorded configuration and selection rule.

D.4.1 Reference marginals and control at a common endpoint tolerance

Let $j \in \{0, 2, 6, 16, 64\}$ denote the number of reference-regression updates, and $s \in \{0.35, 0.60\}$ the teacher's y-coordinate standard deviation. Write D_{ref} for the reported mean reference-marginal SW_2, D_N for the controlled endpoint SW_2 using N integration steps, and K_N for the normalized intrinsic residual energy. All SW_2 values use the standardized chart $(x/0.5, y/s)$.

Table 12 reports means and sample standard deviations across all five seeds. At each reference checkpoint, selection minimizes validation energy among candidates with validation $\text{SW}_2 \le 0.085$; independent testing uses the common tolerance 0.12. The tolerance defines a shared feasibility criterion, allowing realized endpoint errors to vary. All selected controls satisfy the 128-step test criterion, with held-out endpoint SW_2 ranging from 0.059 to 0.101.

Table 12. Reference calibration at matched endpoint tolerance: mean ± sample SD over five seeds. D_{ref} is a mean marginal chart SW_2; D_N is an endpoint chart SW_2. The 256-step columns independently retest the same frozen controls without any additional checkpoint selection.

		Reference	128-step test		256-step retest	
s	j	D_{ref}	K_{128}	D_{128}	K_{256}	D_{256}
0.35	0	0.919 ± 0.009	4.550 ± 0.134	0.077 ± 0.014	4.547 ± 0.181	0.076 ± 0.012
0.35	2	0.495 ± 0.011	1.278 ± 0.062	0.073 ± 0.008	1.277 ± 0.065	0.073 ± 0.014
0.35	6	0.268 ± 0.006	0.232 ± 0.024	0.086 ± 0.017	0.234 ± 0.028	0.082 ± 0.007
0.35	16	0.248 ± 0.010	0.180 ± 0.010	0.085 ± 0.012	0.181 ± 0.010	0.080 ± 0.010
0.35	64	0.248 ± 0.010	0.180 ± 0.007	0.085 ± 0.015	0.181 ± 0.009	0.082 ± 0.008
0.60	0	0.914 ± 0.012	7.243 ± 0.222	0.077 ± 0.012	7.267 ± 0.247	0.076 ± 0.010
0.60	2	0.327 ± 0.022	1.023 ± 0.101	0.077 ± 0.012	1.041 ± 0.115	0.080 ± 0.010
0.60	6	0.232 ± 0.012	0.1672 ± 0.0593	0.075 ± 0.012	0.1695 ± 0.0618	0.078 ± 0.008
0.60	16	0.232 ± 0.012	0.1687 ± 0.0554	0.079 ± 0.013	0.1710 ± 0.0579	0.076 ± 0.007
0.60	64	0.232 ± 0.012	0.1791 ± 0.0603	0.079 ± 0.011	0.1814 ± 0.0626	0.080 ± 0.006

Figure 5 and Section 6.5 interpret the reference and energy trends. Across seeds, Spearman correlations between regression error and reference-marginal SW_2 range from 0.9 to 1.0 for $s = 0.35$ and 0.7 to 1.0 for $s = 0.60$. At both widths, mean trained energy falls sharply during early reference calibration and reaches a similar low-energy range at later checkpoints.

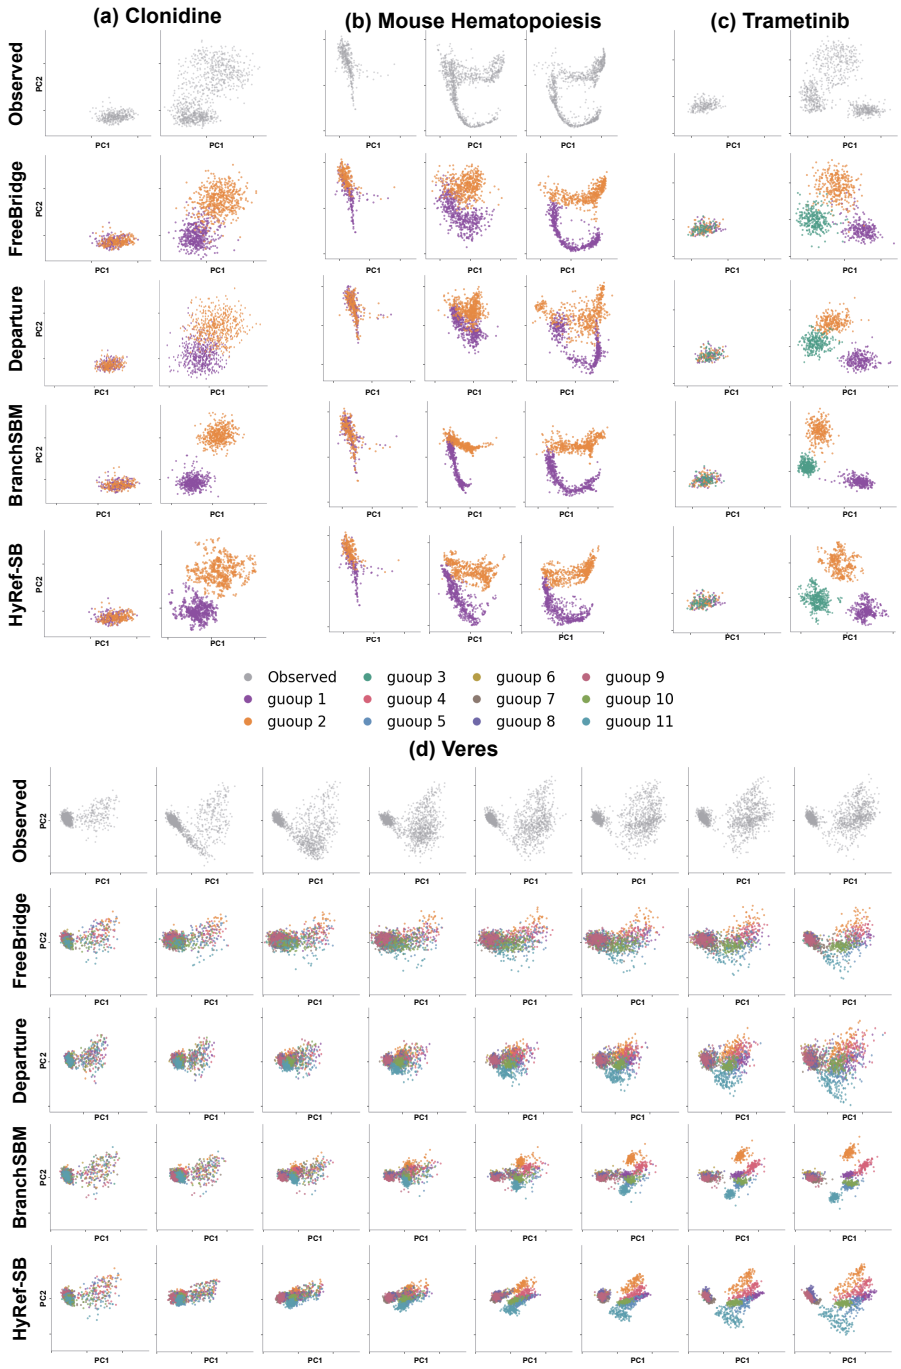

Figure 7. Basic cellular distribution comparisons: (a,c) control and treatment; (b) $t = 0, 0.5, 1$; (d) t_0–t_7. Rows: Observed, FreeBridge, Departures-PC, BranchSBM, and HyRef-SB. PC 1–PC 2 projections use gray for observations and colors for generated terminal conditions.

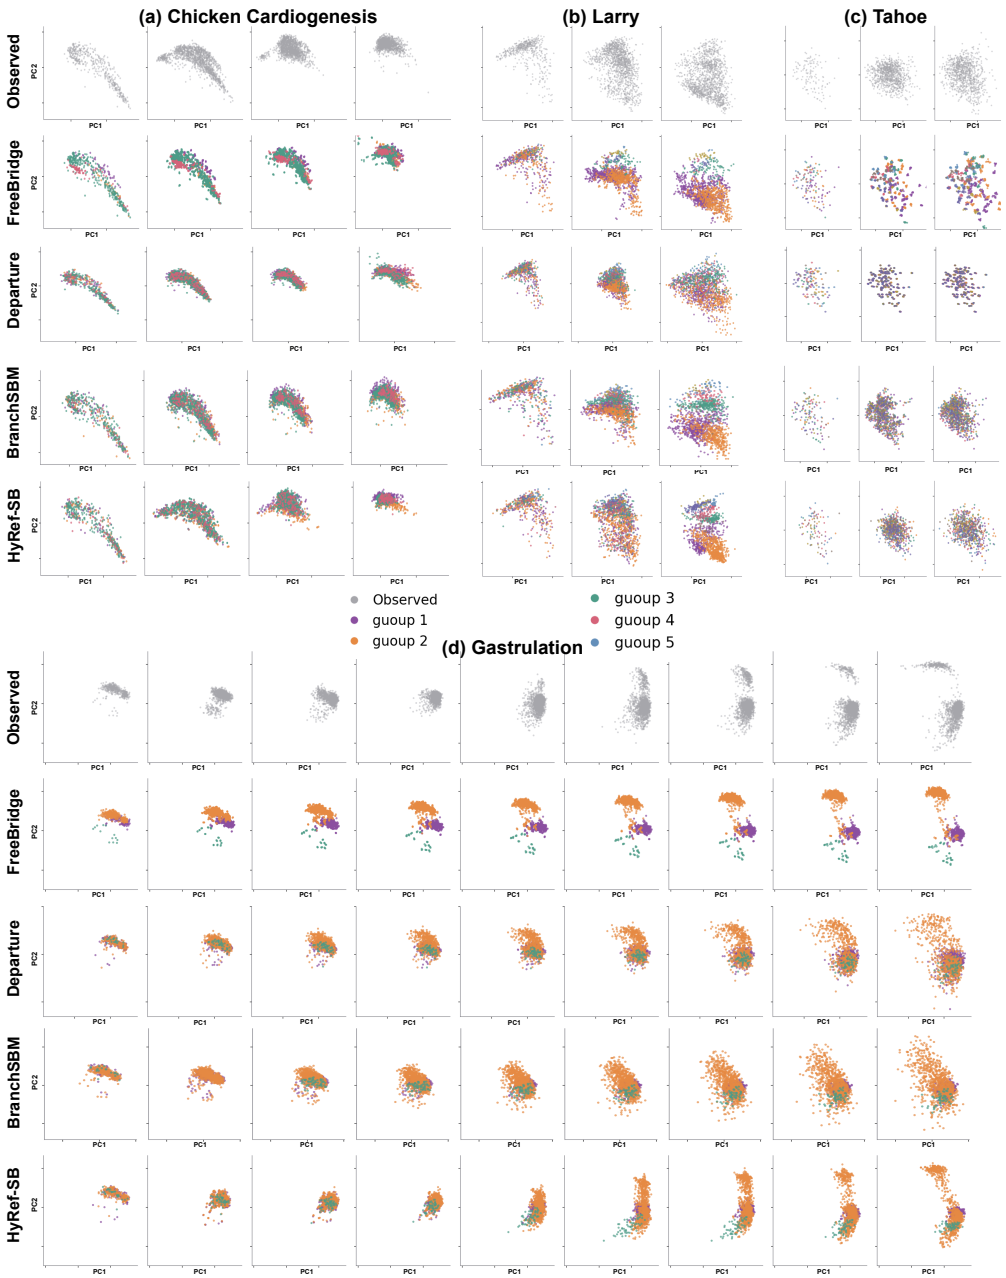

Figure 8. Advanced cellular comparisons in PC 1–PC 2. (a) $t = 0, 0.3, 0.6, 1$; (b) $t = 0, 0.5, 1$; (c) dose order $0, 0.5, 1$; (d) $t = 0, 0.125, \ldots, 1$. Rows: Observed, FreeBridge, Departures-PC, BranchSBM, and HyRef-SB. Gray: observations; colors: generated terminal conditions.

D.4.2 Fisher information and the feasible witness

The diagnostic supplies analytical values of $\mathcal{I}_{\mathrm{T}}$ and of the comparison energy, rather than estimating a score for the cellular benchmarks. The witness from Appendix C.11 has control

$$u^{\dagger} = b_{*} - b_{\mathrm{ref}} + \frac{\sigma^2}{2}\nabla_g \log p_t, \qquad K^{\dagger} = \frac{1}{2\sigma^2}\int_0^1 \mathbb{E}_{\varrho_t}\|u^{\dagger}\|_g^2\, dt.$$

Its density is defined relative to $d\mathrm{vol}_g$, not chart Lebesgue measure. Table 13 separates this analytical witness from the trained control. In this diagnostic the exported exact values satisfy $K^{\dagger} = \varepsilon_{\mathrm{reg}}/(2\sigma^2) + \sigma^2\mathcal{I}_{\mathrm{T}}/8$; this equality is specific to the diagnostic, whereas Proposition 5.3 uses the more general Minkowski upper bound B_{σ}.

Table 13. Fisher contribution and witness diagnostics. $K_{\text{floor}} = \sigma^2 \mathcal{I}_{\text{T}}/8$ with $\sigma = 0.35$. $D^\dagger$ is the witness endpoint chart SW_2, and D_{samp} is the reported teacher-sampling SW_2 floor. These quantities do not vary with the reference checkpoint. Budget and witness energies are summarized below across five seeds as mean $\pm$ sample standard deviation.

s	$\mathcal{I}_{\text{T}}$	K_{floor}	$D^\dagger$	D_{samp}
0.35	14.274	0.219	0.030 ± 0.006	0.027 ± 0.005
0.60	11.996	0.184	0.029 ± 0.004	0.027 ± 0.005

s	j	B_σ	$K^\dagger$ (exact)	$K^\dagger$ (Monte Carlo)
0.35	0	7.569 ± 0.000	5.433 ± 0.000	5.432 ± 0.042
0.35	2	2.466 ± 0.070	1.435 ± 0.049	1.434 ± 0.048
0.35	6	0.526 ± 0.022	0.285 ± 0.008	0.285 ± 0.007
0.35	16	0.225 ± 0.001	0.219 ± 0.000	0.219 ± 0.001
0.35	64	0.219 ± 0.000	0.219 ± 0.000	0.219 ± 0.001
0.60	0	11.051 ± 0.000	8.569 ± 0.000	8.566 ± 0.156
0.60	2	1.392 ± 0.236	0.751 ± 0.146	0.753 ± 0.147
0.60	6	0.231 ± 0.017	0.187 ± 0.002	0.188 ± 0.003
0.60	16	0.184 ± 0.000	0.184 ± 0.000	0.185 ± 0.003
0.60	64	0.184 ± 0.000	0.184 ± 0.000	0.185 ± 0.003

The analytical witness obeys the feasible budget throughout. Monte Carlo energy estimates agree with the exact values, and witness endpoint discrepancy is near the teacher-sampling floor. Section 6.5 distinguishes this witness from the trained control and explains the Fisher contribution.

D.4.3 Frozen-control retest

The separate 256-step retest uses 8192 samples per configuration and a new random stream, with all controls frozen and no additional selection. Every configuration remains below the common endpoint threshold. The energy reduction and low late-stage means persist at both widths, with small mean differences between the two grids relative to the reported seed variability. This retest checks robustness jointly to finer integration and independent stochastic sampling.

E Stagewise Algorithms

Algorithm 1 HyRef-SB training pipeline

Input: Training/validation endpoints, task geometry, condition priors, and training-clone labels for LARRY.
Output: Frozen encoder/decoder, reference drift b_{ref}, control $u_{\hat{\theta}}$, and conditional sampler.

1: **Geometry calibration:** Fit E, D, terminal prototypes, π, and the initial joint law from training endpoints.
2: **Reference construction:** Specify the geometric teacher by Algorithm 2; fit its velocity field.
3: Freeze the fitted drift as b_{ref} and initialize the residual control to zero output.
4: Calibrate and freeze loss scales; fix validation endpoints and validation random numbers.
5: **for** each bridge optimization update **do**
6: Sample from the training initial joint law and unroll all stochastic paths by exponential stochastic steps.
7: Update θ using Equation (5), augmented by the clone-fate term (30) on LARRY.
8: **end for**
9: **Selection and readout:** Apply the endpoint-validation criterion in Appendix F.9; generate independent test paths.

Algorithm 2 Geometry calibration and frozen reference

Input: Training endpoints $\mathcal{X}_0, \{\mathcal{X}_1^k\}$ and task-specific geometry/teacher configuration.
Output: $E, D, \pi, \{c_k\}, b_{\mathrm{ref}}$ and initialization of the residual control.

Fit $m, s, Q, \alpha_{\mathrm{cal}}$ from training endpoints and define the calibrated encoder E and decoder D.
Freeze terminal prototypes, condition priors, and the initial joint law; construct the support pool from training endpoints.
Construct batch couplings from encoded training endpoints; compute the geodesic baseline and endpoint-preserving teacher paths.
Minimize $\mathcal{L}_{\mathrm{teach}}(\eta)$ in Equation (3), with fixed Γ_k, χ_k, and teacher-loss coefficients.
Hold the specified teacher and pairing rule unchanged; sample t, z_0, z_1, k.
Compute the full path derivative $\dot{\psi}_\eta^k$ and warm up b_ω by teacher velocity regression.
Set $b_{\mathrm{ref}} \leftarrow b_{\hat{\omega}}$ and freeze its parameters.
Initialize u_θ to have zero output.
Retain the derivatives of $b_{\mathrm{ref}}(t, z)$ with respect to the state z.

Algorithm 3 One stochastic bridge-control update

Input: Frozen $E, D, b_{\mathrm{ref}}, \pi, s_j$, current parameters θ, step count N, and $\sigma > 0$.
Output: Updated control parameters θ, with the reference and calibration quantities frozen.

Set $h = 1/N$ and sample condition-specific sources Z_0^k from the training initial joint law.
for $n = 0, \ldots, N-1$ **do**
    Set $t_n \leftarrow nh$ and $b_n^k \leftarrow b_{\mathrm{ref}}^k(t_n, Z_n^k) + u_\theta^k(t_n, Z_n^k)$.
    Choose $F_n F_n^\top = g_{Z_n^k}^{-1}$ and sample standard Gaussian noise ξ_n.
    $Z_{n+1}^k \leftarrow \mathrm{Exp}_{Z_n^k}(hb_n^k + \sigma\sqrt{h}F_n\xi_n)$.
    Apply the configured numerical safeguards and record boundary diagnostics over the monitored range.
    Accumulate $h\|u_\theta^k(t_n, Z_n^k)\|_g^2/(2\sigma^2)$ and $hV_{\mathcal{U}}(t_n, Z_n^k, k)$.
end for
Decode all endpoints and compare them with independent conditional target batches to obtain joint endpoint scores $\widehat{e}_k$.
$A \leftarrow \{k : \widehat{e}_k = \max_j \widehat{e}_j\}$; $w_k \leftarrow (1-\alpha)\pi_k + \alpha \mathbf{1}_{k\in A}/|A|$.
Aggregate conditional energy $\widehat{\mathcal{K}}$ and support cost $\widehat{\mathcal{V}}$ using π_k.
Set $\mathcal{L}_{\mathrm{aux}} = 0$; on LARRY, simulate a training-clone batch and set $\mathcal{L}_{\mathrm{aux}} = 0.5\mathcal{L}_{\mathrm{fate}}/s_{\mathrm{fate}}$ using Appendix F.8.
Backpropagate through $\sum_k w_k\widehat{e}_k + \beta\widehat{\mathcal{K}} + \lambda_V\widehat{\mathcal{V}} + \mathcal{L}_{\mathrm{aux}}$, with w as subgradient weights.
Update θ; select checkpoints by the task's endpoint-validation criterion (Appendix F.9).

Algorithm 4 Conditional sampling and source readout

Input: Source x, optional condition k, frozen model, sample count M, and optional frozen q, r.
Output: Conditional path samples and, when requested, source fate probabilities $\widehat{p}(x)$.

Set $z_0 \leftarrow E(x)$; fix k for conditional sampling, or weight conditions by $q(k \mid x)$ for source readout.
for each required condition k **do**
    Generate M paths with positive noise by exponential stochastic steps, using random numbers independent of training and validation.
    Decode $Y_m^k \leftarrow D(Z_{N,m}^k)$ and retain all samples.
    if source fate readout is required **then**
        Compute $r_f(Y_m^k)$ using the frozen training-terminal KDE.
    end if
end for
When requested, set $\widehat{p}_f(x) \leftarrow \sum_k q(k \mid x) M^{-1} \sum_m r_f(Y_m^k)$.
For clone readout, average over the clone's source cells and compare with the corresponding test descendant proportions.

F Experimental Settings and Evaluation Metrics

F.1 Tasks and Data

Clonidine and Trametinib use 50 PCs of Tahoe-100M A549 cells (Zhang et al., 2026) for transport from controls to two and three post-treatment subpopulations, respectively (Tang et al., 2026).

Mouse Hematopoiesis evaluates progenitor differentiation into mature fates (Weinreb et al., 2020) in two-dimensional expression coordinates at normalized times $t = 0.5, 1$. Veres uses 30 PCs and 11

terminal conditions for in vitro pancreatic differentiation of human pluripotent stem cells toward islet β-like and other cells (Veres et al., 2019); t_1–t_6 are intermediate times and t_7 the endpoint.

Gastrulation evaluates lineage diversification during germ-layer formation in a mouse gastrulation and early-organogenesis atlas (Pijuan-Sala et al., 2019). Chicken Cardiogenesis follows myocardial, endocardial, epicardial, and other cardiac-associated populations during embryonic chicken heart development (Mantri et al., 2021). LARRY retains the link between hematopoietic expression states and lineage barcodes (Weinreb et al., 2020). It uses 30 PCs and 9 terminal conditions to evaluate population recovery, and clone descendants to evaluate source-conditioned fate prediction.

The Tahoe dose-response task uses 8 drug conditions and a 50-PC subset from the same atlas (Zhang et al., 2026). Paths index dose order, with intermediate doses held out for evaluation and a separate support reference per drug. Intermediate evaluation positions are 0.5 for LARRY, $0.125, 0.25, \ldots, 0.875$ for Gastrulation, $0.3, 0.6$ for Chicken Cardiogenesis, and 0.5 for Tahoe.

F.2 Data Splits and Supervision

Cell endpoints use an $80/10/10$ training/validation/test split with seed 804, by clone identity for LARRY and stable sample identity for other tasks. Standardization, project-fitted principal component analysis (PCA), dimension-preserving calibration, terminal prototypes, support pools, endpoint pairing, teacher fitting, and reference fitting use training endpoints.

All cell experiments use **no intermediate-observation supervision**: intermediate times and doses are reserved for post-training evaluation. Hyperparameters and checkpoints use endpoint validation, and paths evolve continuously from the source across evaluation times. On LARRY, HyRef-SB, FreeBridge, Departures-PC, BranchSBM, and scDiffEq additionally use the same source–terminal lineage associations and descendant-fate labels from training clones. Validation-terminal information guides model selection, and test-clone descendants are used only for evaluation. Appendix F.8 defines the shared auxiliary supervision.

To align LARRY fate labels and evaluation samples, we retain mature fates with at least 50 terminal cells and 5 clones, giving 9 classes: Neutrophil, Monocyte, Basophil, Mast, Megakaryocyte, Erythroid, Eosinophil, Lymphoid, and Ccr7_DC. Test clones contain at least 1 source cell and 2 terminal descendants. The resulting 61 clones form the common evaluation set for Prior, scDiffEq, and the final HyRef-SB model. Training uses these 9 terminal conditions. A fixed class mapping groups them into five broad classes: granulocytes (Neutrophil, Eosinophil, and Basophil), monocytes/dendritic cells (Monocyte and Ccr7_DC), mast cells (Mast), megakaryocyte/erythroid cells (Megakaryocyte and Erythroid), and lymphoid cells (Lymphoid).

F.3 Network Parameterization and Task Structure

This subsection specifies conditional fields and task readouts; Appendix F.9 gives geometry, losses, budgets, and checkpoint criteria, and Appendix F.8 the LARRY objective.

The drift backbone and shared conditional residual are multilayer perceptrons (MLPs) with three SiLU hidden layers (widths in Table 14) and linear D-dimensional origin-tangent outputs. SiLU is $x\,\mathrm{sigmoid}(x)$; dropout, batch normalization, and layer normalization are disabled. Normalized time $t \in [0,1]$ has the fixed 9-dimensional representation

$$\phi(t) = \big[t, \{\sin(j\pi t)\}_{j=1}^{4}, \{\cos(j\pi t)\}_{j=1}^{4}\big].$$

The backbone input is $[r(z), \phi(t)]$, and the conditional residual input is $[r(z), \phi(t), c_k(z)]$. The $2D + 2$ conditional features derive from training terminal statistics, without a learned branch-identity embedding. Cellular teacher corrections use two SiLU hidden layers of width 128, taking both endpoints in origin tangent coordinates, the same time features, and a 16-dimensional condition embedding to output a D-dimensional correction.

Inputs are normalized by a common scalar scale. For the pooled training source and terminal samples $\{x_i\}_{i=1}^{n}$, set $m = n^{-1}\sum_i x_i$ and $s = \max\{[\sum_i \|x_i - m\|_2^2/(nD)]^{1/2}, 10^{-8}\}$. All cellular tasks use the training-endpoint rule in Appendix B.2 for the frozen coefficient $\eta = \alpha_{\mathrm{cal}}$. The network takes $r(z) = \mathrm{Log}_0\, z/\eta$, and the combined tangent output is multiplied by η before parallel transport. The field notation in Appendix C.13 incorporates this output factor. Let $\bar{v}_0$ be the mean training-source tangent coordinate, p a terminal prototype in tangent coordinates, and ς the coordinatewise

standard deviation of its population. The prototype descriptor is

$$\left[\frac{p-\bar{v}_0}{\|p-\bar{v}_0\|_2},\ \frac{\|p-\bar{v}_0\|_2}{\eta},\ \log\frac{\varsigma}{\eta},\ \log\pi_k\right].$$

Standard deviations have a floor of 10^{-4}, the direction denominator is protected by 10^{-8}, and π_k is the fixed training terminal prior. Clonidine and LARRY use one mean prototype per condition. Other tasks use farthest-point initialization and approximate medoid clustering under intrinsic distance, with up to two prototypes per condition and up to four for Veres. For multiple prototypes, $c_k(z)$ is a weighted sum of descriptors, with weights proportional to $w_{kj}\exp[-d_g(z,p_{kj}^{\mathcal{M}})^2/(2\ell^2)]$. Here $p_{kj}^{\mathcal{M}}=\operatorname{Exp}_0 p_{kj}$, w_{kj} is the cluster's sample fraction within condition k, and ℓ is the median distance from training terminal samples to $\operatorname{Exp}_0\bar{v}_0$, with a floor of 10^{-4}. Prototypes and scales are fitted and frozen before training; attention remains state-dependent.

The temporal gate for the conditional residual is $a_k(t,z)=\operatorname{sigmoid}((t-\tau)/T_g)$. The two drug tasks use $(\tau,T_g)=(0.2,0.1)$; Mouse Hematopoiesis, LARRY, Gastrulation, and Chicken Cardiogenesis use $(0.4,0.12)$. Tahoe uses $a_k\equiv 1$. Veres uses the state-dependent gate

$$a_k(t,z)=0.1+0.9\operatorname{sigmoid}\left(\frac{t-\tau_k+0.15\tanh g_\psi([r(z),c_k(z)])}{T_{g,k}}\right).$$

The network g_ψ uses one width-32 SiLU layer and a zero-initialized scalar linear output. Sigmoid parameterizations constrain τ_k and $T_{g,k}$ to $(0.05,0.95)$ and $(0.03,0.4)$, initialized at 0.4 and 0.12, respectively. The linear adapter is $0.2B_kC_kr$, where $C_k\in\mathbb{R}^{4\times D}$ and $B_k\in\mathbb{R}^{D\times 4}$. Veres and Trametinib instead use $0.2[U_{2,k}\operatorname{SiLU}(U_{1,k}r+b_{1,k})+b_{2,k}]$, a $D\to 32\to D$ network with biases. Mouse Hematopoiesis uses no adapter, and the remaining tasks use the rank-4 linear form. Each task uses at most one adapter type, with zero-initialized output matrices. Both residual-control MLP output layers are zeroed after freezing the reference.

Navigation uses target condition proportions $(0.34,0.33,0.33)$ on $\mathbb{B}_1^2$. The source is centered at the origin with tangent standard deviation 0.12. Target centers have geodesic radius 1.85, direction angles $(-1,0,1)$, and tangent standard deviation 0.11. The two obstacles have disk-coordinate centers $(0.25,-0.11)$ and $(0.28,0.16)$, with intrinsic radii 0.28 and 0.24. Stage budgets and objective coefficients are summarized in Appendix F.9.

Tree generates a 44-node target tree from 16 source nodes, represented with 48 masked slots. Hyperbolic vector quantization encodes nodes, transport acts on the masked product space of nodes, and ordered parent pointers decode adjacency. Geometry, budgets, and sampling appear in Appendix F.9.

Each ablation preserves the main experiment's split, supervision, evaluation coordinates, and optimization budget, changing only the studied component. The scale-matched Euclidean control preserves the dimension, center, scale, and full orthogonal basis, and matches the hyperbolic metric at the origin. The Brownian-reference ablation zeros the reference drift; the single-prototype variant uses one terminal prototype per condition; the linear adapter replaces the nonlinear module; and the no-adapter variant retains the shared conditional residual, prototypes, and condition labels. The decaying teacher anchor reaches zero within the first 25% of the bridge training budget.

The representation–dynamics experiment uses 30-dimensional LARRY and 50-dimensional Gastrulation inputs, each with an 8-dimensional hyperbolic factor. Variants share splits, terminal definitions, preprocessing, supervision, and budgets (Appendix F.9); population testing uses four Monte Carlo replicas. LARRY variants share the clone-fate loss (30). Teachers, references, and bridges are trained independently under matching task settings, varying only geometry.

The hybrid model $H\to E$ retains the full-dimensional calibration E and decoder D of the hyperbolic setting and uses Euclidean geometric operations in the same representation coordinates. $E\to E$ uses both Euclidean representation and Euclidean dynamics. The complete model uses the intrinsic stochastic update in Equation (6). Before numerical safeguards, the hybrid update and control cost are

$$z_{n+1}=z_n+h(b_{\mathrm{ref}}+u)(t_n,z_n)+\sigma\sqrt{h}\,\xi_n,\qquad K_E=\mathbb{E}\sum_n\frac{h\|u(t_n,z_n)\|_2^2}{2\sigma^2}.$$

Geometry enters the teacher baseline, drift features, distances, noise, control norm, and integration. The hybrid retains radius safeguards for its hyperbolic decoder's domain.

F.4 Baselines and a Common Comparison Protocol

Task-adapted FreeBridge, Departures-PC (continuous-PC Departures), and BranchSBM share HyRef-SB's data versions, seed 804, splits, inputs, training endpoints, and conditions. All follow Appendix F.2, including clone-disjoint LARRY splits and intermediate-observation holdouts, and generate continuous paths from their sources.

For LARRY distribution and fate comparisons, these methods and scDiffEq share HyRef-SB's training-clone associations, descendant labels, and auxiliary fate functional, including clone averaging, class mapping, and training-only scaling (Appendix F.8). Prior uses training terminal proportions.

Methods share evaluation coordinates, kernels, bandwidth rules, Wasserstein definitions, sampling, and time aggregation, with method-specific transport parameterizations and budgets. Checkpoints use **endpoint validation**, including nine-class validation-clone fate accuracy for LARRY (Appendix F.9).

RDSB (Thornton et al., 2022) is adapted to Navigation and Tree with task-specific training supervision. On Tree, both methods share the pretrained, frozen hyperbolic vector-quantized representation, structural decoder, and graph-structure supervision; RDSB trains no separate decoder. Structural metrics use decoded transported nodes. On Navigation, both methods receive obstacle information through state/safety training costs. Each RDSB forward and backward network has three hidden layers of width 512 with attention. The HyRef-SB transport fields use widths 256 for Navigation and 128 for Tree, also with attention (Table 14). RDSB uses wider directional networks, with unmatched capacity and optimization. The shared Tree representation, decoder, and supervision fix these components; comparisons concern complete transport systems, without isolating reference learning, architecture, or optimization.

FreeBridge learns a Euclidean controlled diffusion penalizing drift energy and squared nearest-neighbor distance to training-endpoint support (Wang et al., 2026). Training uses at most 200 epochs; evaluation uses the selected checkpoint's exponential moving average parameters.

Departures-PC (Chi et al., 2026) uses minibatch OT endpoint pairing, Brownian bridge interpolants, and endpoint prediction with mean squared error over the input PCs. Its stochastic sampler follows the common evaluation protocol.

BranchSBM uses branch-specific velocities, growth fields, and training-endpoint state costs (Tang et al., 2026). Its four-stage procedure includes the shared fate term during LARRY joint optimization.

scDiffEq (Vinyard et al., 2025) uses the same training lineage and fate information and auxiliary fate supervision under the common clone-disjoint split. Its checkpoint is selected by nine-class terminal-fate accuracy on validation clones. Evaluation uses the same 61 held-out clones as Prior and HyRef-SB, the fine-to-broad mapping in Appendix F.2, and clone-level accuracy from Appendix F.5.2.

F.5 Evaluation Metrics

Let $X = \{x_i\}_{i=1}^n$ and $Y = \{y_j\}_{j=1}^m$ denote generated and observed samples in common task coordinates. We report population, clone-fate, and structural/safety metrics separately.

F.5.1 Population Distribution Matching

We report squared MMD using the empirical V-statistic with diagonal terms (Gretton et al., 2012):

$$\widehat{\mathrm{MMD}}_k^2(X,Y) = \frac{1}{n^2}\sum_{i,i'} k(x_i,x_{i'}) + \frac{1}{m^2}\sum_{j,j'} k(y_j,y_{j'}) - \frac{2}{nm}\sum_{i,j} k(x_i,y_j). \tag{27}$$

Drug tasks and Mouse Hematopoiesis sum five Gaussian kernels; advanced tasks average three scales:

$$k_{\mathrm{drug}}(x,y) = \sum_{s\in\{0.01,0.1,1,10,100\}} \exp\left[-\frac{\|x-y\|_2^2}{2s^2}\right],$$
$$k_{\mathrm{local}}(x,y) = \frac{1}{3}\sum_{a\in\{0.5,1,2\}} \exp\left[-\frac{\|x-y\|_2^2}{2(a\sigma_Y)^2}\right].$$

The common scale σ_Y is the median positive pairwise distance among at most 1024 reference points. Advanced-task MMD^2 uses at most 4096 samples. The two drug tasks are evaluated in

all 50 PCs, Mouse Hematopoiesis uses its two-dimensional coordinates, and the advanced tasks use their respective preprocessed coordinates. MMD^2 is compared within tasks, since kernel normalization and evaluation coordinate scales differ.

For empirical Wasserstein distances, let $a_i = 1/n$, $b_j = 1/m$, and $\Pi(a,b) = \{\gamma \geq 0 : \gamma \mathbf{1} = a, \gamma^\top \mathbf{1} = b\}$. We define

$$W_p(\widehat{\mu}_X, \widehat{\mu}_Y) = \left[\min_{\gamma \in \Pi(a,b)} \sum_{i,j} \gamma_{ij} \|x_i - y_j\|_2^p \right]^{1/p}, \qquad p \in \{1,2\}.$$

Coordinates are fixed across methods: W_1 uses Euclidean costs and W_2 squared Euclidean costs followed by a square root. Clonidine and Trametinib use the first two PCs for W_1, W_2; Mouse Hematopoiesis uses its given two-dimensional coordinates; Veres uses all 30 PCs for W_1; and the advanced tasks use their prescribed preprocessed coordinates for W_1. Veres and advanced tasks also report W_2 using exact empirical OT, without entropy regularization. Methods share sampling rules for each metric. Training uses a separate debiased Sinkhorn discrepancy (Appendix C.18).

Terminal metrics compare terminal distributions; intermediate metrics average times equally. Drug tasks use equal branch mixtures; other generated populations use data priors or fixed source-conditioned weights. Population MMD^2 includes both within- and cross-condition kernel terms.

F.5.2 Clone Fate Prediction

For each test clone c, the ground-truth fine-class probabilities are its terminal descendant proportions, and predicted probabilities average the source-cell readouts in Eq. (25). The class mapping in Appendix F.2 aggregates these into five broad-class probability vectors $\bar{p}_c, \widehat{\bar{p}}_c$. For $N_C = 61$ test clones and $G = 5$ broad classes, dominant-fate accuracy is

$$\text{Top-1} = \frac{1}{N_C} \sum_c \mathbf{1}\left[\arg\max_{1 \leq g \leq G} \widehat{\bar{p}}_{cg} = \arg\max_{1 \leq g \leq G} \bar{p}_{cg} \right].$$

Each clone receives equal weight. Top-1 compares the broad class with the largest descendant proportion, resolving ties by a fixed class order; probabilities are aggregated before taking the maximum. Prior uses the training terminal prior, and HyRef-SB uses the generated-endpoint readout. Table 5 reports clone-level accuracy. The scDiffEq re-evaluation follows the matched-clone comparison protocol in Appendix F.4.

F.5.3 Navigation: Path Safety and Length

For a generated path $\{z_{i,n}\}_{n=0}^N$, let o_j, r_j be obstacle centers and intrinsic radii. A path is marked as colliding if it enters any obstacle at any interior grid point:

$$I_i = \mathbf{1}\left[\min_{1 \leq n < N,\, j} \{d_g(z_{i,n}, o_j) - r_j\} < 0 \right].$$

For B paths, collision rate is $B^{-1} \sum_i I_i$ and safety rate is its complement. Path length is

$$\text{PathLength} = \frac{1}{B} \sum_{i=1}^{B} \sum_{n=0}^{N-1} d_g(z_{i,n}, z_{i,n+1}).$$

Distances use the hyperbolic metric; safety refers to obstacle avoidance on the interior grid. Table 1 reports safety, collision rate, and length. RDSB safety is the complement of its path-collision rate, and collision percentages are converted to per mille.

F.5.4 Tree: Structural Distributions and Paired Edge Quality

All structural metrics use valid nodes. Adjacency matrices are thresholded at 0.5; undirected edges use the upper triangle without self-loops. Paired metrics retain target node indices.

Degree MMD^2 uses normalized 48-bin histograms of node degrees on $[0, 48]$ to describe local branching connectivity. Spectral MMD^2 uses normalized 20-bin histograms on $[0, 2]$ of the eigenvalues of the normalized graph Laplacian $L_A = I - D_A^{-1/2} A D_A^{-1/2}$, capturing global connectivity.

The inverse square root of a zero degree is set to zero. Both metrics compare generated and reference graph collections using the unit-bandwidth RBF kernel $k(h, h') = \exp(-\|h - h'\|_2^2/2)$ and Eq. (27). The overall collections have condition masses $(0.34, 0.33, 0.33)$. These MMD^2 metrics compare unpaired graph collections on separate numerical scales.

Edge F1 aggregates true-positive (TP), false-positive (FP), and false-negative (FN) edge counts over the valid node pairs across the entire collection of paired graphs:

$$\mathrm{EdgeF1} = \frac{2\mathrm{TP}}{2\mathrm{TP} + \mathrm{FP} + \mathrm{FN}}.$$

This metric pools edges with fixed node correspondence alongside degree/spectral distribution metrics.

F.6 Reporting conventions

Benchmarks use seed 804; the reference-calibration diagnostic (Appendix F.7) uses seeds 801–805. Repeated sampling from a trained model is distinguished from independent retraining. Tables display MMD^2 on the indicated $\times 10^{-2}$ scale. Methods share within-task evaluation coordinates, metric definitions, and times. Bold marks the best available value among methods reporting a metric: lower distribution errors and collisions, or higher safety, accuracy, and Edge F1.

F.7 Configuration of the reference-calibration diagnostic

Section 6.5 and Appendix D.4 use the following diagnostic settings; benchmarks retain seed 804.

The state space is $\mathbb{H}^2$ of curvature -1, expressed in coordinates $(x, y) \in \mathbb{R}^2$ with $g = \mathrm{diag}(e^{-2y}, 1)$ and $d\mathrm{vol}_g = e^{-y}\,dx\,dy$. The two teacher configurations have y-coordinate standard deviations $s = 0.35$ and 0.60. Noise intensity is fixed at $\sigma = 0.35$. SGD fits a constant horizontal reference drift $b_{\mathrm{ref},j} = a_j \partial_x$, with checkpoints after $j = 0, 2, 6, 16, 64$ regression updates. Each configuration records a_j, $\varepsilon_{\mathrm{reg}}$, $\mathcal{I}_{\mathrm{T}}$, B_σ, and exact witness energy. Fisher information and the witness are analytical quantities for these prescribed teachers.

Distribution comparisons use sliced 2-Wasserstein distance in the standardized chart $S_s(x, y) = (x/0.5, y/s)$, fixed across checkpoints within each width. Reference quality is compared within widths because y scaling depends on s. Chart SW_2, intrinsic Wasserstein distance, and marginal KL in Proposition 5.3 measure distinct discrepancies. With $Z_n = (X_n, Y_n)$ and $u_n = (u_n^x, u_n^y)$, residual energy uses the theoretical metric normalization:

$$K_N = \mathbb{E} \sum_{n=0}^{N-1} \frac{h}{2\sigma^2} \left[e^{-2Y_n} (u_n^x)^2 + (u_n^y)^2 \right], \qquad h = 1/N.$$

We report mean reference-marginal chart SW_2 as D_{ref}, alongside exact and Monte Carlo witness energies, the teacher-sampling SW_2 floor, and trained-control metrics.

The experiment uses seeds 801–805, batch size 768, 32 training integration steps, 360 optimization iterations per recorded endpoint-penalty setting, and penalty values $10, 50, 250, 1000$. The history logs candidate validation metrics every 60 iterations and also includes the zero-control candidate. For each seed, width, and reference checkpoint, the selected control has the smallest validation energy among candidates satisfying validation $\mathrm{SW}_2 \le 0.085$. Held-out testing uses 128 steps and 8192 samples with a common endpoint tolerance $\mathrm{SW}_2 \le 0.12$. All 50 selected controls satisfy both criteria; energy comparisons concern these controls at a common tolerance.

The controls remain frozen for a separate 256-step retest with 8192 samples per configuration and stream offset 4,000,000 (original: 1,000,000). The retest columns report these independent results.

Means and sample standard deviations use all five seeds per (s, j) group, with SD denominator $5 - 1$. Fisher information is analytically constant within widths. Reference-marginal Spearman correlations use five checkpoints for each seed and width.

F.8 LARRY Fate Objective

For a training clone c, let n_{cf} count its terminal descendants in fine fate $f \in \{1, \ldots, 9\}$ and set $y_{cf} = n_{cf} / \sum_{j=1}^{9} n_{cj}$. The target y_c is a probability vector. Let $A \in \{0, 1\}^{9 \times 5}$ encode the fixed

class mapping in Appendix F.2, with one nonzero entry per row. Given a minibatch $\mathcal{B}_C$ of training clones, sample source cells $x_{c,1}, \ldots, x_{c,S}$ within each clone and form

$$\widehat{p}_{\theta,c} = \frac{1}{S}\sum_{s=1}^{S} \widehat{p}_\theta(\cdot \mid x_{c,s}), \qquad H_\tau(y) = (1-\tau)y + \tau e_{\arg\max_f y_f}. \tag{28}$$

Here $\widehat{p}_\theta$ is the Monte Carlo generated-endpoint readout of Equation (25), and e_f is the corresponding one-hot vector; ties use a fixed class order. Each source-cell prediction averages M stochastic paths per terminal condition before the clone average. The terminal KDE is fitted to training endpoints and frozen before bridge training. Its input derivatives propagate the fate loss through generated endpoints to the transport parameters.

With $\mathrm{CE}(y,p) = -\sum_f y_f \log p_f$ and $\tau = 0.5$, the shared fate functional is

$$\mathcal{L}_{\mathrm{fate}}(\theta) = \frac{1}{|\mathcal{B}_C|}\sum_{c\in\mathcal{B}_C}\Big[\mathrm{CE}(H_{0.5}(y_c), \widehat{p}_{\theta,c}) + 0.1\|y_c - \widehat{p}_{\theta,c}\|_2^2 + \mathrm{CE}(H_{0.5}(y_c A), \widehat{p}_{\theta,c} A)\Big]. \tag{29}$$

Clones receive equal weight. Cross-entropy sums over classes, and the fine-class squared-error term sums its nine squared differences without division by the number of classes. The broad target is formed by aggregation with A before hard–soft mixing. The complete HyRef-SB LARRY objective is

$$\mathcal{J}_{\mathrm{LARRY}}(\theta) = \underbrace{\mathcal{R}_\alpha(e) + \beta\mathcal{K}(\theta) + \lambda_V\mathcal{V}(\theta)}_{\mathcal{J}(\theta) \text{ in Equation (5)}} + \frac{0.5}{s_{\mathrm{fate}}}\mathcal{L}_{\mathrm{fate}}(\theta). \tag{30}$$

The endpoint weights are $\lambda_{\mathrm{MMD}} = 0.2$ and $\lambda_{\mathrm{Sinkhorn}} = 0.3$, with $\lambda_V = 0.01$, $\sigma = 0.05/\sqrt{30}$, and $\beta = 0.01\sigma^2$. After reference fitting, four training-only probe batches determine

$$s_{\mathrm{fate}} = \max\left\{1, \frac{1}{4}\sum_{\ell=1}^{4}\mathcal{L}_{\mathrm{fate}}^{(\ell)}\right\}.$$

This frozen scale jointly normalizes all three fate-loss terms. Supervision uses terminal descendants and their source–terminal lineage associations under Appendix F.2.

Each fate minibatch contains at most eight training clones, sampled uniformly without replacement. For each clone, $S = 2$ source cells are sampled with replacement; each source generates $M = 2$ antithetic paths for each of the nine conditions using 64 integration steps. Clone size and generated path count do not change the clone's loss weight. For a trainable comparison method m, the LARRY adaptation uses the combined objective

$$\mathcal{J}_{\mathrm{LARRY}}^{(m)} = \mathcal{J}_{\mathrm{transport}}^{(m)} + \frac{0.5}{s_{\mathrm{fate}}^{(m)}}\mathcal{L}_{\mathrm{fate}}^{(m)}.$$

The methods in Appendix F.4 retain their transport objectives and use this common fate functional and training-only scale rule. At evaluation, HyRef-SB generates eight paths per source cell and terminal condition with 128 steps, averaging predictions over all source cells in each test clone.

F.9 Geometry and Optimization Settings

Reference drifts regress endpoint-preserving teacher velocities; cellular teachers use training endpoints and task-specific geometric objectives. Residual training freezes reference parameters while retaining state derivatives; the Brownian ablation uses zero drift.

Cells use Equation (5), augmented by Equation (30) for LARRY. Endpoint coefficients are $(0.2, 0.3)$ for MMD^2 and debiased Sinkhorn divergence; data losses are divided by fixed training-calibrated scales. Path-support weights are 0 for Clonidine/Trametinib and 0.01 otherwise.

Control weight is $\beta = 0.01\sigma^2$ with orthonormal noise $\sigma = 0.05/\sqrt{D}$. Training uses batch size 64, 64 updates/epoch, 64 integration steps, and two antithetic paths; evaluation uses 128 steps. AdamW uses decay 10^{-5} and a cosine schedule with minimum rate 0.1 of its initial value.

Table 14. Geometry and vector fields. Cells use $D = d_H + d_E$; Tree dimensions are per node in its masked product. Networks list hidden width × depth. Tree fits curvature from $c_{\text{init}} = 1$ before final transport, then freezes it throughout that stage.

Task	K	D	d_H	d_E	(c, w_E)	Vector field
Mouse Hematopoiesis	2	2	2	0	$(1, –)$	128×3
LARRY	9	30	8	22	$(1, 1)$	256×3
Tahoe	8	50	8	42	$(1, 1)$	256×3
Veres	11	30	8	22	$(1, 1)$	256×3
Trametinib	3	50	8	42	$(1, 1)$	256×3
Gastrulation	3	50	8	42	$(1, 1)$	256×3
Chicken Cardiogenesis	4	50	8	42	$(1, 1)$	256×3
Clonidine	2	50	8	42	$(1, 1)$	256×3
Tree	3	16	16	0	fitted, –	128×4
Navigation	3	2	2	0	$(1, –)$	256×3

Table 15. Cellular optimization. T/W/B: teacher fitting/reference-velocity warm-up/residual-control training. Entries give maximum epochs (early-stopping patience). Teacher coefficients follow Equation (3); E/F denote the endpoint-distribution/fate validation criteria below.

Task	$(\lambda_{\text{T,kin}}, \lambda_{\text{T,sup}})$	T / W / B	W / B learning rates	Select
Mouse	$(1, 0.2)$	$80(15)/180(35)/240(40)$	$3\times10^{-4}/10^{-4}$	E
LARRY	$(1, 1)$	$80(15)/180(25)/240(40)$	$3\times10^{-4}/10^{-4}$	F
Tahoe	$(1, 1)$	$80(15)/180(15)/240(25)$	$10^{-4}/5\times10^{-5}$	E
Veres	$(1, 1)$	$80(15)/240(40)/240(40)$	$3\times10^{-4}/5\times10^{-5}$	E
Trametinib	$(1, 1)$	$80(15)/240(40)/240(40)$	$3\times10^{-4}/5\times10^{-5}$	E
Gastrulation	$(1, 1)$	$80(15)/180(25)/240(40)$	$3\times10^{-4}/10^{-4}$	E
Chicken	$(1, 1)$	$80(15)/180(25)/240(40)$	$3\times10^{-4}/10^{-4}$	E
Clonidine	$(1, 1)$	$80(15)/180(25)/240(40)$	$3\times10^{-4}/10^{-4}$	E

Teacher fitting uses learning rate 3×10^{-4} and Equation (3), with kinetic factor $1/2$ and averaging over times, endpoint pairs, and fixed χ_k. Trametinib uses endpoint-risk parameter $\alpha = 0.15$; Tahoe activates conditional corrections initially. Cells use Appendix B.2 calibration.

Tree uses 16 hyperbolic dimensions per node, at most 400 representation-pretraining epochs and $150/300/10/150$ teacher/drift/mass/final epochs, with respective rates 10^{-3} and $10^{-3}/3\times10^{-4}/5\times10^{-4}/10^{-4}$. Batch size 64 gives 63 updates/epoch. Navigation uses at most $150/300/100/220$ teacher/drift/mass/safety epochs with rates $10^{-3}/3 \times 10^{-4}/5 \times 10^{-4}/2 \times 10^{-4}$ and 19 updates/epoch at batch size 256. Its teacher kinetic/obstacle-state/boundary weights are $1/4/0.3$; final endpoint/safety/action weights are $1/1/0.02$. Tree teacher kinetic/boundary weights are $1/0.1$. Both use AdamW decay 10^{-5}; Tree/Navigation evaluate with $(\sigma, M) = (0.05, 8)$ and $(0.08, 16)$, respectively, using 128 integration steps.

Cellular distribution checkpoints (E) minimize fixed endpoint-validation risk $\mathcal{R}_\alpha(e_{\text{val}})$. LARRY (F), including baseline adaptations, maximizes original nine-class terminal-fate accuracy on validation clones; selected checkpoints yield five-class test metrics. Teacher/reference stages use their losses on validation endpoint pairs. Tree uses terminal structural validation; Navigation prioritizes validation safety with terminal-goal matching. Selection excludes observed intermediate times, intermediate doses, and test outcomes.