GENERALIZED MOMENT-INFORMED FLOW MATCHING FOR TRAJECTORY INFERENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Trajectory inference seeks to reconstruct continuous dynamics from sparse, unpaired snapshots of evolving populations. These observations do not uniquely determine the underlying path process, so modeling assumptions are needed to select plausible dynamics. Collecting more samples per snapshot or observing additional time points can improve reconstruction, but detailed measurements are often costly or destructive. We therefore consider settings in which less costly measurements provide aggregate information in the form of expectations of test functions under intermediate marginal distributions. We build on flow matching, which learns a vector field to reproduce the marginal evolution of a prescribed stochastic interpolant through simulation-free regression. To incorporate the aggregate observations into this framework, we introduce Generalized Moment-Informed Flow Matching (GMI-FM). Our method first learns an endpoint-preserving neural interpolant under the aggregate constraints using an augmented-Lagrangian formulation, then regresses the corresponding vector field against the interpolant's velocities, keeping both stages simulation-free. Our framework also accommodates geometric priors, such as data-dependent metrics, through the interpolant regularization, encouraging learned trajectories to follow the geometry of the data. We provide theoretical grounding for GMI-FM by relating its construction to a Schrödinger bridge formulation with the same aggregate constraints. Experiments on synthetic and real-world datasets demonstrate that GMI-FM achieves better performance on marginal reconstruction metrics than competing methods. These findings highlight the potential of integrating lower-cost aggregate measurements with costly, detailed population snapshots to improve trajectory inference.

1 INTRODUCTION

Many scientific systems are observed as evolving populations, with individual trajectories remaining inaccessible. In single-cell biology, for example, transcriptomic measurements destroy the measured cells, yielding unpaired snapshots of a population at different times. Trajectory inference seeks to reconstruct the evolution underlying these observations, with applications to cellular development, reprogramming, and disease progression (Schiebinger et al., 2019; Huguet et al., 2022).

Population snapshots alone do not generally identify the underlying dynamics: different processes can share the same observed marginals while exhibiting different intermediate distributions and temporal dependence (Weinreb et al., 2018; Tong et al., 2020). Trajectory inference methods therefore introduce inductive biases to select among compatible evolutions. Optimal-transport approaches favor low-cost transport between populations (Schiebinger et al., 2019; Tong et al., 2023a), while Schrödinger bridges select the path law closest to a reference process in relative entropy, subject to the observed marginal constraints (Léonard, 2013). The accuracy of the reconstruction depends on how well these modeling choices describe the unobserved evolution. Additional measurements at intermediate times can restrict this ambiguity by distinguishing processes that agree at the observed snapshots.

In this work, we consider intermediate information in the form of generalized moments, $\mathbb{E}[f_k(X_{\tau_k})] = a_k$, for $k = 1, \ldots, K$, where X_{τ_k} denotes the population state at time τ_k, f_k is a supplied observation function, and a_k is the corresponding observed population average. Such observations arise naturally from experimental measurements that capture selected population properties. In atmospheric monitoring, an instrument measures how much light a population of airborne particles scatters at a

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.

fixed wavelength. Combining this measurement with the particle concentration gives an estimate of the average scattering response per particle (Luoma et al., 2019). In this setting, f describes how much light a particle of a given size is expected to scatter. In cell biology, calibrated bulk RNA-seq measurements can estimate mean marker-gene abundances, corresponding to coordinate observation functions in gene-count space. Fluorescent reporters offer another source of information: with a suitable observation model relating cellular state to reporter signal, population-averaged fluorescence provides a generalized-moment observation without repeated transcriptomic profiling (Park et al., 2022; Ko et al., 2022). When suitable markers distinguish cell types, measured population fractions can likewise be expressed as expectations of indicator functions. These examples illustrate how complementary measurements can constrain intermediate population distributions without requiring those distributions to be observed in full.

The challenge is then to construct a continuous probability path that matches the observed population snapshots while satisfying the generalised moment constraints, and to learn a velocity field that generates trajectories consistent with this marginal evolution. We introduce *Generalized Moment-Informed Flow Matching* (GMI-FM) to address these two components in sequence. First, building on learned-interpolant approaches (Kapuśniak et al., 2024; Petrović et al., 2025), we learn a probability path by augmenting a reference interpolant with a neural correction that preserves the observed endpoint marginals. The correction is optimized to satisfy the intermediate moment constraints through an augmented-Lagrangian objective, while a path regularizer resolves the remaining ambiguity by encoding preferences such as proximity to the reference interpolant or adherence to a prescribed geometry. Once this constrained probability path has been learned, we freeze the interpolant and learn a velocity field that generates its marginal evolution using conditional flow matching (Lipman et al., 2022; Albergo and Vanden-Eijnden, 2022). Both stages operate directly on samples and velocities provided by the interpolant, avoiding numerical trajectory integration during training. Experiments on synthetic, atmospheric, and single-cell datasets show that incorporating aggregate observations in this way can improve intermediate marginal reconstruction, including at times without aggregate supervision. [1]

We summarise our contributions below:

1. We introduce GMI-FM, which learns an endpoint-preserving probability path under generalised moment constraints and trains a corresponding velocity field, with both stages remaining simulation-free.
2. We establish a theoretical connection between GMI-FM and moment-constrained Schrödinger bridges. Through an auxiliary Brownian-bridge perturbation, we derive a relative-entropy interpretation of the interpolant regulariser and characterise its relationship to the constrained bridge solution.
3. On synthetic, atmospheric, and single-cell datasets, we demonstrate that generalised moment information can improve intermediate marginal reconstruction, including with inaccurate estimates of the generalised moments. Comparisons with competing methods, including CSB-FM, support the effectiveness of GMI-FM and highlight the potential of integrating lower-cost aggregate measurements with detailed population snapshots.

2 Preliminaries and Setting

We introduce trajectory inference and review Flow Matching (FM), the scalable, simulation-free training framework underlying our method.

The problem of trajectory inference. Given source and target distributions p_0, p_1 on $\mathbb{R}^d$, we seek a time-dependent vector field $v : [0,1] \times \mathbb{R}^d \to \mathbb{R}^d$ whose ODE transports p_0 to p_1:

$$\dot{X}_t = v(t, X_t), \qquad X_0 \sim p_0, \qquad \mathrm{Law}(X_1) = p_1. \tag{1}$$

In practice, only endpoint samples are available, so p_0 and p_1 may be empirical distributions.

[1] source code: `https://anonymous.4open.science/r/FM_moments_constrained-C43C`.

Conditional Flow Matching (CFM). One approach learns v by minimising a distributional loss between p_1 and samples obtained by integrating the field (Chen et al., 2018; Tong et al., 2020). FM (Lipman et al., 2022; Liu et al., 2022; Albergo and Vanden-Eijnden, 2022) instead uses simulation-free regression, reducing training cost while achieving strong empirical results in trajectory inference (Tong et al., 2023a).

FM relies on a coupling π between p_0 and p_1 and an interpolation map I. For $(X_0, X_1) \sim \pi$, the interpolant $I_t = I(t, X_0, X_1)$ defines a probability path $p_t = \mathrm{Law}(I_t)$, required to satisfy $\mathrm{Law}(I_0) = p_0$ and $\mathrm{Law}(I_1) = p_1$. Averaging the endpoint-conditioned velocities at each state gives the marginal velocity field:

$$\dot{I}_t := \partial_t I(t, X_0, X_1), \qquad u_t(x) = \mathbb{E}[\dot{I}_t \mid I_t = x]. \tag{2}$$

Since u_t is generally unavailable in closed form, CFM regresses a neural vector field directly against $\dot{I}_t$:

$$\mathcal{L}_{\mathrm{CFM}}(\psi) = \mathbb{E}_{(X_0,X_1)\sim\pi,\, t\sim\mathcal{U}(0,1)} \left[\left\| v_\psi(t, I_t) - \dot{I}_t \right\|^2 \right]. \tag{3}$$

The minimiser recovers u_t in the population limit (Lipman et al., 2022).

Common instantiations. The simplest coupling, $\pi = p_0 \otimes p_1$, pairs endpoints independently without geometric preferences. Optimal-transport couplings favour lower-cost endpoint assignments, encouraging straighter paths (Tong et al., 2023a). A common choice is the linear interpolant $I_t = (1-t)X_0 + tX_1$, with velocity $\dot{I}_t = X_1 - X_0$. Stochastic interpolants (Albergo and Vanden-Eijnden, 2022) and nonlinear paths can encode further assumptions about intermediate dynamics (Kapuśniak et al., 2024; Petrović et al., 2025). We provide a more detailed discussion of the related work in Section B.

3 Generalised Moment-Constrained Flow Matching

In this section, we formulate intermediate aggregate observations as generalised moment constraints and cast the learning problem as the selection of a probability path from the corresponding set of admissible path laws under an appropriate regularisation. We then propose a two-stage procedure. First, we identify a probability path that optimises the regularised objective while satisfying the prescribed constraints; second, we learn a velocity field whose induced flow reproduces its marginal evolution.

3.1 Intermediate observations as generalised moments

Let $\Omega = C([0,1]; \mathbb{R}^d)$ denote the trajectory space and let $P^\star \in \mathcal{P}(\Omega)$ denote the unknown data-generating path law. We observe unpaired samples from its endpoint distributions, denoted by p_0 and p_1. In addition, we assume access to aggregate information at intermediate times $0 < \tau_1 < \cdots < \tau_K < 1$. We represent such information through *generalised moments*, namely quantities of the form $\mathbb{E}[f(X_\tau)]$ for a known observation function f.

More formally, let v denote a velocity field and P^v the path law it induces. Our objective is to find a velocity field whose induced path law is consistent with the endpoint distributions and the available intermediate aggregate observations. This leads to the following admissible set.

Definition 1 (Admissible Path Laws)**.** We define the set of admissible path laws as

$$\mathcal{C} = \left\{ P \in \mathcal{P}(\Omega) : \begin{array}{ll} P_0 = p_0, & P_1 = p_1, \\ \mathbb{E}_P[f_k(X_{\tau_k})] = a_k, & k = 1, \dots, K \end{array} \right\}, \tag{4}$$

where $f_k : \mathbb{R}^d \to \mathbb{R}^{d_k}$ is a known observation function and $a_k \in \mathbb{R}^{d_k}$ is the corresponding observed aggregate quantity. We seek a velocity field v such that $P^v \in \mathcal{C}$.

This problem formulation encompasses many real-world settings where measurements are naturally expressed as generalised moments. Depending on the application, the observation function f_k may

encode, for example, a population moment, a bulk spectral signal, a marker abundance, or a population fraction. Table 1 illustrates several such cases.

Table 1: Examples of generalised-moment observations.

Observation	Observation function	Aggregate observation
Classical moments	$x \mapsto x^{\alpha\ 2}, \quad \alpha \in \mathbb{N}_0^d$	Moments of arbitrary order
Bulk spectral measurements (Luoma et al., 2019)	$x \mapsto r(x; \lambda_{\mathrm{opt}})$. A particle's contribution to the optical signal at wavelength λ_{opt}.	Combined emitted or scattered signal from the population, normalised by particle number to obtain a.
Single-cell marker abundance (Semrau et al., 2017)	$x \mapsto x_g$. Count of marker gene g in one cell, expressed in gene-count coordinates.	Mean marker count supplied by independent FISH measurements, after calibration to the chosen count scale.
Population composition (e.g. cell types)	$x \mapsto \mathbf{1}\{c(x) = j\}$. Class-membership indicator, where $c(x)$ is the particle's class.	$\Pr(c(X_\tau) = j)$. Fraction of particles belonging to class j.

3.2 Learning an aggregate-constrained conditional path

At a high level, our strategy to solve this problem separates the construction of a probability path in $\mathcal{C}$ from the learning of a velocity field which generates it. We first fix an endpoint coupling and learn a conditional interpolant whose induced marginals satisfy the prescribed generalised-moment constraints, while remaining as close as possible to a chosen reference interpolant under an appropriate regularisation. Once this constrained path has been identified, we freeze the interpolant and use conditional flow matching to learn a marginal velocity field whose dynamics reproduce its marginal evolution.

3.2.1 Constructing an aggregate-constrained probability path

Our first objective is to construct a probability path whose associated path law belongs to the admissible set $\mathcal{C}$. Following stochastic-interpolant constructions (Albergo and Vanden-Eijnden, 2022), we represent the path law as a mixture of deterministic conditional paths. Because $\mathcal{C}$ imposes fixed endpoint marginals p_0 and p_1, we condition directly on endpoint pairs

$$(X_0, X_1) \sim \pi, \qquad \pi \in \Pi(p_0, p_1),$$

where $\Pi(p_0, p_1)$ denotes the set of couplings between the endpoint distributions. The coupling π specifies how samples from p_0 and p_1 are paired and may, for instance, be chosen as an optimal-transport or entropically regularised optimal-transport coupling.

Given an endpoint pair (x_0, x_1), we consider a deterministic trajectory $I(\cdot, x_0, x_1)$ satisfying $I(0, x_0, x_1) = x_0$ and $I(1, x_0, x_1) = x_1$. The induced path law is

$$P^I = \int \delta_{I(\cdot, x_0, x_1)}\, \pi(dx_0, dx_1). \tag{5}$$

Since π has marginals p_0 and p_1, the endpoint constraints defining $\mathcal{C}$ are satisfied automatically.

A natural starting point is a simple reference interpolant I_{ref}. Throughout much of this work, we consider linear interpolation, i.e. $I_{\mathrm{ref}}(t, x_0, x_1) = (1-t)x_0 + tx_1$. Although this construction satisfies the endpoint constraints, it will not in general satisfy the prescribed intermediate generalised-moment constraints. Moreover, this discrepancy cannot always be resolved by changing the endpoint coupling. For example, for a mean constraint with $f = \mathrm{Id}$, $\mathbb{E}\left[I_{\mathrm{ref}}(\tau, X_0, X_1)\right] = (1-\tau)\mathbb{E}_{p_0}[X] + \tau\mathbb{E}_{p_1}[X]$ is independent of π. Hence, if a different mean is prescribed at time τ, the interpolant itself must be modified.

We therefore fix the coupling π and deform the reference interpolant so that the resulting path law satisfies the aggregate constraints. Since multiple admissible deformations may exist, we favour solutions that remain close to I_{ref} according to a prior regulariser R_{prior}. We parameterise the

deformation by $I_\theta(t, x_0, x_1) = I_{\mathrm{ref}}(t, x_0, x_1) + t(1-t)g_\theta(t, x_0, x_1)$, where g_θ is a neural network (Kapuśniak et al., 2024; Petrović et al., 2025). The factor $t(1-t)$ forces the correction to vanish at both endpoints, so optimisation changes only the intermediate evolution while preserving the endpoint marginals exactly. For convenience, define the random interpolated state $X_t^\theta := I_\theta(t, X_0, X_1)$.

The resulting learning problem is

$$\boxed{\begin{aligned} \theta^\star \in \arg\min_\theta \quad & R_{\mathrm{prior}}(\theta) \\ \text{subject to} \quad & \mathbb{E}\left[f_k\left(X_{\tau_k}^\theta\right)\right] = a_k, \qquad k = 1, \dots, K. \end{aligned}} \tag{6}$$

Assuming these constraints are feasible for the chosen coupling, the resulting interpolant induces a path law in $\mathcal{C}$. In practice, we solve (6) using an augmented-Lagrangian method (Hestenes, 1969; Powell, 1969), combining stochastic-gradient updates of the interpolant parameters with updates of Lagrange multipliers based on the moment residuals. This directly promotes constraint satisfaction without requiring excessively large quadratic-penalty coefficients, which can lead to poorly conditioned optimisation.

The regulariser R_{prior} determines which admissible interpolant is selected when the observations do not uniquely determine the path. For the linear reference, we penalize deviations from its conditional velocity by setting R_{prior} to be:

$$R_{\mathrm{lin}}(\theta) = \alpha \mathbb{E}\left[\|\partial_t I_\theta(t, X_0, X_1) - \partial_t I_{\mathrm{ref}}(t, X_0, X_1)\|^2\right], \qquad t \sim \mathcal{U}(0,1), \quad \alpha > 0, \tag{7}$$

where t is sampled independently of (X_0, X_1).

R_{prior} can also encode alternative geometric prior knowledge, such as the LAND metric considered in Section D.

3.2.2 From the constrained interpolant to a marginal flow

The constrained interpolant defines endpoint-conditioned trajectories, whereas generation requires a velocity field depending only on time and state. Writing $I_t = I_{\hat\theta}(t, X_0, X_1)$ for $(X_0, X_1) \sim \pi$, conditional flow matching learns this field by regressing the trajectory velocity $\dot I_t$ onto the interpolated state:

$$\mathcal{L}_B(\psi; \hat\theta) = \mathbb{E}_{\substack{(X_0,X_1)\sim\pi \\ t\sim\mathcal{U}(0,1)}}\left[\left\| v_\psi(t, I_t) - \dot I_t \right\|^2\right]. \tag{8}$$

The population optimum and its sampling ODE are

$$v^\star(t,x) = \mathbb{E}[\dot I_t \mid I_t = x], \qquad \dot X_t = v^\star(t, X_t), \quad X_0 \sim p_0. \tag{9}$$

Under standard regularity conditions, this flow reproduces the interpolant marginals $p_t^{\hat\theta}$ (Albergo and Vanden-Eijnden, 2022) and therefore inherits their moment constraints: $\mathbb{E}[f_k(X_{\tau_k})] = a_k$ for $k = 1, \dots, K$. Algorithm 1 summarises the two-stage training.

3.3 Connection to Schrödinger bridges

Schrödinger bridges select a stochastic path law by minimising relative entropy to a reference process subject to prescribed observations (Léonard, 2013). We use this formulation to interpret the path preference encoded by R_{lin} and compare the ideal GMI-FM interpolant with a bridge satisfying the same endpoint and generalised moment constraints. Throughout this subsection, $I_{\mathrm{ref}} = \ell$ and $R_{\mathrm{prior}} = R_{\mathrm{lin}}$.

Definition 2 (Schrödinger bridge with generalised moment constraints)**.** Let R be the law of $dX_t = \sigma\, dW_t$, where $\sigma > 0$, $X_0 \sim p_0$, and W is a standard d-dimensional Brownian motion independent of X_0. The moment-constrained Schrödinger bridge solves

$$P_{\mathrm{SB}}^\star \in \arg\min_{P \in \mathcal{C}} \mathrm{KL}(P \| R), \tag{10}$$

where $\mathcal{C}$ is the admissible set defined in (4).

For the SB problem with endpoint constraints alone, the solution is a mixture of Brownian bridges whose endpoint pairs follow the optimal entropic coupling. This provides a natural reference for interpreting our regulariser: when $I_{\mathrm{ref}} = \ell$ and π is this optimal coupling, adding a Brownian-bridge perturbation to the reference interpolant recovers the endpoint-only SB. The following proposition shows that R_{lin} is proportional to the relative entropy between the similarly perturbed learned interpolant and this reference.

Proposition 1 (Relative entropy of the noisy interpolant)**.** *Let $\sigma > 0$ and let B be a standard d-dimensional Brownian bridge independent of $Z = (X_0, X_1) \sim \pi$. Define $Q_\theta^\sigma = \mathrm{Law}((I_\theta(t, Z) + \sigma B_t)_{t\in[0,1]})$ and $Q_{\mathrm{ref}}^\sigma = \mathrm{Law}((I_{\mathrm{ref}}(t, Z) + \sigma B_t)_{t\in[0,1]})$. Assume that $h_\theta(t, Z) = I_\theta(t, Z) - I_{\mathrm{ref}}(t, Z)$ is almost surely absolutely continuous in time, vanishes at both endpoints, and has finite expected energy $A(\theta) = \mathbb{E}_\pi \int_0^1 \|\partial_t h_\theta(t, Z)\|^2 \, dt$. Then*

$$\mathrm{KL}(Q_\theta^\sigma \,\|\, Q_{\mathrm{ref}}^\sigma) = \frac{A(\theta)}{2\sigma^2}. \tag{11}$$

Consequently, the velocity-deviation regulariser satisfies $R_{\mathrm{lin}}(\theta) = 2\alpha\sigma^2 \,\mathrm{KL}(Q_\theta^\sigma \| Q_{\mathrm{ref}}^\sigma)$.

Consequently, minimising R_{lin} favours feasible interpolants whose noisy lifts remain close, in relative entropy, to the Brownian-bridge reference. We now examine how this selection relates to the Schrödinger bridge with the same endpoint and generalised moment constraints. Assume that (10) admits a solution $P^\star$ with finite relative entropy; existence and duality conditions are discussed in Appendix G.1. This comparison clarifies the consequences of the two path constructions. Their reconstruction accuracy depends on how well the reference process and the chosen path family describe the unknown dynamics. To isolate the comparison of conditional paths, we assume a common endpoint coupling, $\pi = P^\star_{0,1}$. Two sources of discrepancy remain: the feasible interpolant family may not represent the constrained bridge exactly, and minimising R_{lin} may select a different interpolant from the one whose noisy lift is closest to that bridge. The following proposition bounds this second discrepancy. It therefore explains when our regulariser selects a good approximation to the constrained SB relative to the best approximation available within the chosen family.

Proposition 2 (Variational comparison with the constrained SB)**.** *Consider the Brownian reference with diffusion coefficient $\sigma > 0$ and a single observation $\mathbb{E}_P[f(X_\tau)] = a$, where $f \in C^2(\mathbb{R}^d; \mathbb{R})$ satisfies $\sup_x \|\nabla^2 f(x)\|_{\mathrm{op}} \le H$. Let $\lambda^\star$ be a finite moment multiplier associated with $P^\star$.*
Take $\pi = P^\star_{0,1}$, $I_{\mathrm{ref}} = \ell$, and $R_{\mathrm{prior}} = R_{\mathrm{lin}}$, and let $\theta^\star$ solve (6)*. Under the assumptions detailed in Appendix G.3,*

$$0 \le \mathrm{KL}(Q_{\theta^\star}^\sigma \| P^\star) - \varepsilon_{\mathrm{app}} \le |\lambda^\star| H d \sigma^2 \tau(1-\tau), \tag{12}$$

where Q_θ^σ is the noisy path law defined in Proposition 1, and $\varepsilon_{\mathrm{app}}$ is the infimum of $\mathrm{KL}(Q_\theta^\sigma \| P^\star)$ over feasible interpolants with finite R_{lin}.

For $f = \mathrm{Id}$, the connection becomes exact after adding the Brownian-bridge perturbation.

Proposition 3 (Exact correspondence for a mean constraint)**.** *Let $P^\star$ be the finite-relative-entropy SB solution with $\mathbb{E}[X_\tau] = a$ and endpoint laws having finite second moments. Set $\pi = P^\star_{0,1}$ and $R_{\mathrm{prior}} = R_{\mathrm{lin}}$. Over all endpoint-preserving finite-energy interpolants, the Stage-A optimum is*

$$I^\star(t, x_0, x_1) = \ell_t(x_0, x_1) + \min\left\{\frac{t}{\tau}, \frac{1-t}{1-\tau}\right\} \Delta, \qquad \Delta = a - \big((1-\tau)\mathbb{E}_{p_0} X + \tau \mathbb{E}_{p_1} X\big).$$

Its noisy lift equals $P^\star$, and its unperturbed marginals satisfy

$$W_2(\mathrm{Law}(I^\star(t, Z)), P_t^\star) \le \sigma\sqrt{d\, t(1-t)}, \qquad Z \sim \pi. \tag{13}$$

The proof is given in Appendix H.3. Mean and covariance constraints are treated in Proposition 5 in the appendix.

Table 2: Synthetic diffusion, velocity-field pushforward. Avg. W_2 averages reconstruction errors at $t = 0.25$ and $t = 0.75$; $R = \sqrt{\|\widehat{\mu}_{0.5} - \mu_{0.5}\|_2^2 + (\widehat{v}_{y,0.5} - v_{y,0.5})^2}$.

Method	Avg. W_2 at $0.25, 0.75 \downarrow$	W_2 at $0.5 \downarrow$	W_2 at $1 \downarrow$	Residual $R \downarrow$
CFM	0.6316 ± 0.0037	0.8851 ± 0.0021	$\mathbf{0.0517 \pm 0.0024}$	0.8367 ± 0.0029
MFM	0.6865 ± 0.0361	0.9234 ± 0.0497	0.0981 ± 0.0268	0.8838 ± 0.0426
CSB-FM	0.2460 ± 0.0144	0.1967 ± 0.0034	0.0803 ± 0.0221	0.0370 ± 0.0167
GMI-FM (lin)	$\mathbf{0.1710 \pm 0.0063}$	$\mathbf{0.1892 \pm 0.0064}$	0.0625 ± 0.0068	$\mathbf{0.0248 \pm 0.0149}$
GMI-FM (LAND)	0.1834 ± 0.0076	0.1956 ± 0.0067	0.0688 ± 0.0099	0.0359 ± 0.0126

4 EXPERIMENTS

We evaluate whether GMI-FM with intermediate aggregate observations improves marginal reconstruction, particularly at times without aggregate constraints. We consider a synthetic diffusion to study mean and variance constraints in a controlled setting, atmospheric particle-size measurements with aggregate optical observations from complementary instruments, and two single-cell datasets with different forms of biological information. In Multi, cell annotations provide cell-type proportions, while in Semrau, calibrated bulk RNA-seq measurements provide estimates of mean marker-gene abundance. We consider two variants: GMI-FM (lin), which uses the velocity-deviation regulariser R_{lin}, and GMI-FM (LAND), which uses the geometric regulariser described in Appendix D. We compare against conditional flow matching with linear interpolants (CFM) (Tong et al., 2023a), Metric Flow Matching (MFM) (Kapuśniak et al., 2024), and CSB-FM, which trains a velocity field from a numerical approximation to the moment-constrained Schrödinger bridge in (10). CFM and MFM provide baselines without aggregate constraints; CSB-FM receives the same observation functions and targets as GMI-FM. Within each benchmark, all methods use the same maximum number of updates for training. Details can be found in Appendix F.6. We measure marginal reconstruction using the empirical W_2 distance between observed populations and samples generated by integrating the learned velocity field. We also report aggregate residuals measuring disagreement between the supplied targets and the corresponding statistics of the generated populations, with the residual definition specified for each benchmark. Unless stated otherwise, results are reported as mean $\pm$ population standard deviation over five optimisation seeds.

4.1 SYNTHETIC NONLINEAR-DRIFT SDE

We consider a two-dimensional Itô process in which the population follows an arched route, contracts vertically near an intermediate bottleneck, and then spreads again (Figure 2). We use this benchmark to test whether midpoint observations of the population's location and spread improve reconstruction at other times. We simulate $N = 200{,}000$ particles from a Gaussian source distribution to construct the reference marginals, and denote the process in normalised time by Z_t, $t \in [0, 1]$. The SDE definition and simulation details are provided in Appendix F. For the generalised-moment constraints, we provide the two-dimensional mean and the vertical variance at $\tau = 0.5$.

Results. We find that midpoint observations improve reconstruction at the unconstrained times $t = 0.25$ and $t = 0.75$ (Table 2). GMI-FM (lin) reduces average W_2 from 0.6316 for CFM to 0.1710, and also improves over CSB-FM, which receives the same observations and achieves 0.2460. GMI-FM (LAND) achieves an average W_2 of 0.1834. The linear prior therefore gives the lowest mean intermediate reconstruction errors on this benchmark.

4.2 ATMOSPHERIC MEASUREMENTS

We reconstruct aerosol size distributions from measurements collected at Hyytiälä, Finland, on June 1, 2012 (Luoma et al., 2019; Kulmala, 2018). We normalise differential mobility particle sizer (DMPS) histograms and represent log diameter in endpoint-standardised coordinates. Five snapshots correspond to $t \in \{0, 0.25, 0.5, 0.75, 1\}$; only the endpoints are used for training, while interior histograms are reserved for evaluation.

At $\tau = 0.5$, a nephelometer measures the bulk scattering coefficient b_{sca} at 550 nm, and a condensation particle counter measures concentration N_{CPC}. Their ratio, $y_\tau := b_{\text{sca}}(\tau)/N_{\text{CPC}}(\tau) \approx \mathbb{E}[k(D_\tau)]$,

Table 3: June 1 aerosol reconstruction from velocity-field pushforwards. Avg. W_2 averages $t = 0.25, 0.75$; $R = |\mathbb{E}_{\hat{p}_{0.5}}[f] - a_{0.5}|$.

Method	Avg. W_2 at $0.25, 0.75 \downarrow$	W_2 at $0.5 \downarrow$	W_2 at $1 \downarrow$	Residual $R \downarrow$
CFM	0.2918 ± 0.0001	0.3414 ± 0.0002	$\mathbf{0.00358 \pm 0.00004}$	0.1908 ± 0.0004
MFM	0.2907 ± 0.0001	0.3393 ± 0.0002	0.00362 ± 0.00006	0.1901 ± 0.0004
CSB-FM	0.2908 ± 0.0028	0.3404 ± 0.0024	0.01551 ± 0.00358	0.0849 ± 0.0168
GMI-FM (lin)	0.2917 ± 0.0001	0.3411 ± 0.0002	0.00410 ± 0.00039	0.0013 ± 0.0005
GMI-FM (LAND)	$\mathbf{0.2873 \pm 0.0002}$	$\mathbf{0.3331 \pm 0.0003}$	0.00362 ± 0.00016	$\mathbf{0.0011 \pm 0.0008}$

estimates the mean scattering cross section per particle. Here D_τ denotes diameter and k is an effective Mie scattering response (Mie, 1908; Fierz-Schmidhauser et al., 2010). Its refractive index and calibration gain are fitted on separate January–February measurements and held fixed. Implementation details are given in Appendix F.

GMI-FM (LAND) reduces W_2 relative to MFM by approximately 1.2% at the unconstrained times and 1.8% at the midpoint (Table 3). Both GMI-FM variants closely match the optical target. GMI-FM (lin) shows little improvement over CFM.

4.3 Single-cell trajectory inference

Single-cell RNA sequencing provides high-dimensional measurements of individual cells, but its destructive nature yields unpaired population snapshots. We evaluate GMI-FM on Multi and Semrau using different forms of aggregate information. In Multi, days 2 and 7 define the endpoints, while days 3 and 4 alternate as constraint and validation days in two reciprocal settings. The unconstrained day supplies the W_2 metric for checkpoint selection. In Semrau (Semrau et al., 2017), which follows mouse embryonic stem-cell differentiation under retinoic acid exposure, we use 0 and 36 hours as endpoints, impose aggregate constraints at 6 hours, and evaluate reconstruction at 12 and 24 hours. Semrau cells are represented by eight principal components and four marker coordinates for Sox2, Tbx3, Nanog, and Ppia, with PCA fitting and scaling based on endpoint training cells. LAND reference pools contain only endpoint cells in both datasets.

The generalised moments describe cell-type composition in Multi and marker-gene abundance in Semrau. For Multi, we train a classifier D_τ on annotated constraint-day cells and freeze it during trajectory learning. Its seven soft outputs define the observation map, and their empirical mean defines the target: $f_\tau(x) = D_\tau(x), \qquad a_\tau = \frac{1}{N_\tau} \sum_{i=1}^{N_\tau} D_\tau(x_i^{(\tau)})$. Thus, the target matches the empirical expectation of the same map used in the constraint loss. For Semrau, we estimate mean Sox2 and Tbx3 counts at 6 hours from accompanying bulk RNA-seq measurements, using a calibration fitted to paired bulk and single-cell measurements and held fixed during trajectory learning. Since marker coordinates are scaled square-root counts, $x_g = \sqrt{n_g}/s_g$, the corresponding moments are $f_g(x) = x_g^2, \qquad a_g = \frac{\widehat{m}_g}{s_g^2}, \qquad g \in \{\mathrm{Sox2}, \mathrm{Tbx3}\}$, where $\widehat{m}_g$ is the calibrated mean count and s_g is the endpoint-derived scale. Calibration details are provided in Appendix F.

Tables 4 and 5 show that GMI-FM (LAND) achieves the lowest mean W_2 across the reported reconstruction metrics in both datasets. In Multi, it reduces reciprocal-average unconstrained-day W_2 from 10.1290 for MFM to 10.0017 and constraint-day W_2 from 10.0707 to 10.0533, while reducing the moment residual from 0.1435 to 0.0789. Both reciprocal-average W_2 improvements hold across all five seeds. In Semrau, it reduces average W_2 at 12 and 24 hours from 1.7111 to 1.6726 and constraint-time W_2 from 2.1161 to 2.0732, corresponding to relative reductions of approximately 2.25% and 2.03%.

4.4 Ablation studies

We examine how reconstruction depends on the information supplied by the generalised moments. We first assess sensitivity to observation error, then compare moment constraints with snapshot supervision when midpoint samples are scarce. Additional studies of complementary and redundant moment information are provided in Appendix C.

Table 4: Multi velocity-field pushforward with empirical moment targets. Mean ± SD over five seeds after averaging reciprocal settings; $R_\tau = \|\mathbb{E}_{\hat{p}_\tau}[D_\tau] - a_\tau\|_2$.

Method	Mean unconstrained-day $W_2 \downarrow$	Mean constraint-day $W_2 \downarrow$	Mean residual $R_\tau \downarrow$
CFM	10.5495 ± 0.0037	10.4755 ± 0.0035	0.1415 ± 0.0006
MFM	10.1290 ± 0.0039	10.0707 ± 0.0027	0.1435 ± 0.0009
CSB-FM	10.9551 ± 0.0037	10.8383 ± 0.0012	0.0870 ± 0.0034
GMI-FM (lin)	10.5551 ± 0.0041	10.4383 ± 0.0039	$\mathbf{0.0448 \pm 0.0023}$
GMI-FM (LAND)	$\mathbf{10.0017 \pm 0.0027}$	$\mathbf{10.0533 \pm 0.0023}$	0.0789 ± 0.0006

Table 5: Semrau, velocity-field pushforward. Avg. W_2 averages 12 and 24 hours; $R_{\text{rel}} = \frac{100}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} |\mathbb{E}_{\hat{p}_\tau}[f_g]/a_g - 1|$.

Method	Avg. W_2 at 12h, 24h $\downarrow$	W_2 at 6h $\downarrow$	W_2 at 36h $\downarrow$	Residual R_{rel} (%) $\downarrow$
CFM	2.0168 ± 0.0016	2.3517 ± 0.0004	1.9636 ± 0.0107	27.3030 ± 0.1291
MFM	1.7111 ± 0.0050	2.1161 ± 0.0032	1.7678 ± 0.0142	28.3955 ± 0.5979
CSB-FM	2.0025 ± 0.0044	2.3433 ± 0.0036	1.9460 ± 0.0051	4.8806 ± 0.8879
GMI-FM (lin)	1.9961 ± 0.0010	2.3343 ± 0.0005	1.9527 ± 0.0061	0.7197 ± 0.6165
GMI-FM (LAND)	$\mathbf{1.6726 \pm 0.0062}$	$\mathbf{2.0732 \pm 0.0020}$	$\mathbf{1.7664 \pm 0.0153}$	$\mathbf{0.6887 \pm 0.3340}$

Sensitivity to observation error. We assess sensitivity to inaccurate aggregate observations using the June 1 atmospheric measurements from Section 4.2. Keeping the calibrated optical response fixed, we compute its reference moment a_{ref} from the observed midpoint histogram and supply perturbed targets

$$a_{r,j} = a_{\text{ref}}(1 + r\xi_j), \qquad \xi_j \sim \mathcal{U}[-1, 1], \qquad j = 1, \ldots, 5.$$

We vary the maximum relative error r up to 0.95 and retrain GMI-FM (LAND), reusing the five noise draws across error levels while fixing the optimisation seed. The unperturbed reference is fitted once. Figure 1 reports reconstruction at the constrained time $t = 0.5$. All five noise draws yield lower W_2 than MFM at the tested error bounds up to 40%. Larger perturbations reveal an asymmetric response: substantial underestimation of the target degrades reconstruction, while positive perturbations retain an advantage over MFM across the tested range. At the largest error bound, the mean W_2 exceeds that of MFM. These results indicate tolerance to moderate observation error in this example, but sufficiently inaccurate targets can reverse the benefit of aggregate information.

Reconstruction from scarce midpoint samples On the arched bridge (Section 4.1), we compare moment constraints with snapshot supervision using $n \in \{4, 16, 64, 256, 1024\}$ midpoint samples. We fix the endpoints and optimisation randomness and use five independent observation sets, nested across sample sizes. GMI-FM and CSB-FM receive only the sample mean and unbiased vertical variance. Three multi-marginal (MM) baselines use the same observations: *MM-CFM (empirical)* trains one velocity field on OT-coupled linear segments between $t = 0, 0.5, 1$; *MM-CFM (Gaussian)* replaces the empirical midpoint with samples from a Gaussian fitted to its mean and full covariance; and *MM-MFM (empirical)* learns LAND segment paths. The latter includes midpoint samples in its local metrics, whereas GMI-FM (LAND) uses only global endpoints. We evaluate velocity rollouts against independent simulations at $t = 0.25, 0.5, 0.75$ (Table 8).

5 Conclusion

We introduced GMI-FM, a simulation-free framework that incorporates generalised moment observations into endpoint-preserving probability paths. Across synthetic and single-cell benchmarks, these observations improve intermediate reconstruction when combined with a suitable path prior. Our results highlight aggregate information as a practical way to reduce the ambiguity of trajectory inference from sparse population snapshots. Further discussion can be found in Appendix A.

BROADER IMPACTS

Our work introduces a foundational method for trajectory inference with primary applications in the natural sciences. Positively, it can advance single-cell biology by improving models of cellular development and disease progression from sparse data. We do not foresee direct malicious applications; however, practitioners should be cautious, as biased data or misspecified constraints could lead to incorrect downstream scientific conclusions.

AI USE STATEMENT

In this work, we used generative AI tools to generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in the writing of proofs, and implement methods.

We have not used generative AI tools to propose or refine hypotheses, design or provide feedback on research methodology or experiments, assist with translation, clean and reformat datasets, support qualitative and thematic data analysis, or interpret results.

Additionally, we used generative AI tools to create or modify scientific figures or images, create or edit software code, and draft parts of a research paper.

We have reviewed all AI-assisted work. All mathematical derivations have been carefully checked and human-verified, as has AI-generated software by the authors. The final version of the manuscript is human written.

We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

Aurélien Alfonsi, Rafaël Coyaud, Virginie Ehrlacher, and Damiano Lombardi. Approximation of optimal transport problems with marginal moments constraints. *Mathematics of Computation*, 90 (328):689–737, 2021. doi: 10.1090/mcom/3568.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.

Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pages 146–158, 1975.

Rahel Fierz-Schmidhauser, Paul Zieger, Martin Gysel, Lukas Kammermann, Peter F DeCarlo, Urs Baltensperger, and Ernest Weingartner. Measured and predicted aerosol light scattering enhancement factors at the high alpine site jungfraujoch. *Atmospheric Chemistry and Physics*, 10 (5):2319–2333, 2010.

Magnus R Hestenes. Multiplier and gradient methods. *Journal of optimization theory and applications*, 4(5):303–320, 1969.

Guillaume Huguet, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.

Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957. doi: 10.1103/PhysRev.106.620.

Kacper Kapuśniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Avishek J Bose, and Francesco Di Giovanni. Metric flow matching for smooth interpolations on the data manifold. *Advances in Neural Information Processing Systems*, 37:135011–135042, 2024.

Jina Ko, Martin Wilkovitsch, Juhyun Oh, Rainer H Kohler, Evangelia Bolli, Mikael J Pittet, Claudio Vinegoni, David B Sykes, Hannes Mikula, Ralph Weissleder, et al. Spatiotemporal multiplexed immunofluorescence imaging of living cells and tissues with bioorthogonal cycling of fluorescent probes. *Nature biotechnology*, 40(11):1654–1662, 2022.

Markku Kulmala. Measurement of particle number size distribution at hyytiälä. EBAS, NILU, 2018. URL `https://doi.org/10.48597/KB6J-XYGA`. Dataset. Measurement period: 1 January 2012–1 January 2013.

Christian Léonard. A survey of the schr\" odinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and Ricky TQ Chen. Generalized schrödinger bridge matching. In *International Conference on Learning Representations*, volume 2024, pages 14527–14552, 2024.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

Krista Luoma, Aki Virkkula, Pasi Aalto, Tuukka Petäjä, and Markku Kulmala. Over a 10-year record of aerosol optical properties at smear ii. *Atmospheric Chemistry and Physics*, 19(17):11363–11382, 2019.

Gustav Mie. Beiträge zur optik trüber medien, speziell kolloidaler metallösungen. *Annalen der physik*, 330(3):377–445, 1908.

Olga Movilla Miangolarra, Asmaa Eldesoukey, and Tryphon T Georgiou. Inferring potential landscapes: A schrödinger bridge approach to maximum caliber. *Physical Review Research*, 6(3): 033070, 2024.

Gwanghyun Park, Minkyung Shin, Wonyoung Lee, Akitsu Hotta, Taeko Kobayashi, and Yoichi Kosodo. Direct visualization of the transition status during neural differentiation by dual-fluorescent reporter human pluripotent stem cells. *Stem cell reports*, 17(9):1903–1913, 2022.

Katarina Petrović, Lazar Atanackovic, Viggo Moro, Kacper Kapuśniak, Ismail Ilkan Ceylan, Michael M Bronstein, Joey Bose, and Alexander Tong. Curly flow matching for learning non-gradient field dynamics. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

Michael JD Powell. A method for nonlinear constraints in minimization problems. *Optimization*, pages 283–298, 1969.

Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.

Stefan Semrau, Johanna E Goldmann, Magali Soumillon, Tarjei S Mikkelsen, Rudolf Jaenisch, and Alexander Van Oudenaarden. Dynamics of lineage commitment revealed by single-cell transcriptomics of differentiating embryonic stem cells. *Nature communications*, 8(1):1096, 2017.

Manu Setty, Vaidotas Kiseliovas, Jacob Levine, Adam Gayoso, Linas Mazutis, and Dana Pe'er. Characterization of cell fate probabilities in single-cell data with palantir. *Nature Biotechnology*, 37(4):451–460, 2019. doi: 10.1038/s41587-019-0068-4.

Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pages 9526–9536. PMLR, 2020.

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023a.

Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schr\" odinger bridges via score and flow matching. *arXiv preprint arXiv:2307.03672*, 2023b.

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.

A Discussion

Moment information and path priors. Prior effectiveness depends on the dynamics: the linear prior performs best on our synthetic bridge, while LAND gives the strongest reconstruction on the real datasets. There, adding generalised moment constraints improves reconstruction at unobserved times over MFM, supporting the complementary roles of geometry and aggregate information. Since moments do not determine the full marginal evolution, $R_{\mathrm{prior}}(\theta)$ remains essential for selecting among interpolants whose path laws lie in $\mathcal{C}$. These choices are modular: differentiable, endpoint-preserving reference interpolants and geometric regularisers can be incorporated while retaining the same moment constraints and velocity-learning stage. CSB-FM's large aggregate residuals arise mainly during velocity learning: the numerical bridge closely satisfies the constraints, but the learned field fails to preserve them. Noisy training velocities, which can become large near endpoints and observation times, hinder regression. Accumulated rollout errors can then distort marginals and violate constraints even as the training loss decreases.

Runtime. Compared with classical CFM, GMI-FM adds Stage A to learn an interpolant satisfying the supplied moments. Stage B then requires additional neural function evaluations and time derivatives of this frozen interpolant to construct the regression targets, whereas CFM uses an analytic linear path. These computations account for the fitting overhead in Table 9, which excludes setup and sampling. With LAND, GMI-FM and MFM share the same two-stage structure and have comparable Stage-B runtimes. Moment constraints therefore build on the computation already required for geometric path learning, although enforcing them can increase Stage-A and total fitting time. CSB-FM incurs the cost of a numerical bridge solve. For nonlinear observations in high dimension, as in Multi, fitting the moment multipliers requires repeated numerical integration and conditional sampling, which can dominate runtime. This cost is much smaller when analytic Gaussian moment calculations are available, as in Semrau.

Limitations. Moment errors can bias reconstruction. The atmospheric sensitivity study (Appendix C.3) shows tolerance to moderate optical-target perturbations, while larger errors can reverse gains over MFM. More accurate targets need not improve marginal W_2 monotonically, as reconstruction also depends on the path prior and interpolant family. Accounting for uncertainty in targets and calibrated observation functions remains an important extension. Gains also depend on the information supplied: on the synthetic bridge, midpoint mean and vertical variance constrain complementary bottleneck features. Less informative observables may yield smaller gains despite accurate matching, while still enforcing scientifically meaningful properties. With exact population observations and constraint satisfaction, the selected and ground-truth path laws belong to $\mathcal{C}$ and agree on the supplied observables. This excludes incompatible evolutions but neither guarantees a smaller distance to the full ground-truth path law nor identifies individual trajectories. In practice, optimisation and velocity regression make constraint satisfaction approximate.

B Related Work

Trajectory inference from population snapshots. Trajectory inference from unpaired population snapshots has been widely studied in single-cell biology and related scientific domains, where individual trajectories are typically unavailable. Early optimal-transport approaches infer couplings

Table 6: Moment complementarity, velocity-field pushforward. Avg. W_2 averages $t = 0.25$ and $t = 0.75$. Results are aggregated over five seeds.

Method	Avg. W_2	W_2 at 0.5	W_2 at 1	R_μ	R_v	R
CFM	0.6316 ± 0.0037	0.8851 ± 0.0021	0.0517 ± 0.0024	0.8000 ± 0.0035	0.2449 ± 0.0048	0.8367 ± 0.0029
GMI-FM (mean only)	0.2774 ± 0.0067	0.3840 ± 0.0077	0.0587 ± 0.0041	0.0197 ± 0.0115	0.2514 ± 0.0090	0.2524 ± 0.0098
GMI-FM (vertical variance only)	0.5857 ± 0.0052	0.8054 ± 0.0078	0.0552 ± 0.0025	0.7840 ± 0.0083	0.0070 ± 0.0016	0.7840 ± 0.0083
GMI-FM (mean + vertical variance)	0.1710 ± 0.0063	0.1892 ± 0.0064	0.0625 ± 0.0068	0.0247 ± 0.0149	0.0019 ± 0.0019	0.0248 ± 0.0149

between temporally adjacent populations and use these couplings to reconstruct developmental progressions (Schiebinger et al., 2019; Setty et al., 2019). Subsequent neural approaches parameterize continuous-time dynamics directly, for instance with neural ODEs and continuous normalizing flows, and regularize the resulting dynamics through kinetic-energy, density, or manifold-based penalties (Chen et al., 2018; Tong et al., 2020; Huguet et al., 2022). These methods highlight the central challenge of snapshot trajectory inference: endpoint or sparsely observed marginal distributions alone do not determine a unique dynamics, so additional inductive biases are required (Weinreb et al., 2018; Tong et al., 2020).

Flow matching and interpolant design. Flow matching provides a simulation-free framework for learning continuous-time generative dynamics by regressing vector fields along prescribed conditional probability paths (Lipman et al., 2022; Liu et al., 2022; Albergo and Vanden-Eijnden, 2022). In trajectory inference, the choice of coupling and interpolant determines the implied probability path between observed marginals. OT-CFM uses optimal-transport couplings to favor low-action trajectories (Tong et al., 2023a), while geometric variants modify the interpolant itself: Metric Flow Matching learns approximate geodesic interpolants under data-dependent Riemannian metrics, encouraging paths to remain near the data manifold (Kapuśniak et al., 2024). Other recent approaches alter the reference dynamics, for example by using non-zero-drift Schrödinger bridge references to encode non-gradient or cyclic behavior (Petrović et al., 2025). Our work also intervenes at the conditional-path level, but uses intermediate moment information rather than only geometry, endpoint couplings, or reference drift to select among endpoint-preserving interpolants.

Entropy projections and constrained bridges. Moment constraints on probability measures are classical in maximum-entropy inference and information projection (Jaynes, 1957; Csiszár, 1975), and have also been studied in optimal transport (Alfonsi et al., 2021). On path space, Movilla Miangolarra et al. (2024) combine Schrödinger bridges with maximum caliber to incorporate ensemble-path constraints. Generalized Schrödinger Bridge Matching addresses task-specific state costs through conditional stochastic control and variational approximations (Liu et al., 2024). These works establish that additional observables and path costs fit within existing bridge frameworks. Simulation-free score and flow matching learns classical bridges from entropic endpoint couplings and analytic conditional targets (Tong et al., 2023b). Building on these entropy-projection and bridge-matching results, we formulate conditional-path learning for partial marginal observations and characterise its relation to the entropic reference.

C Additional Experiments

C.1 Complementarity of moment information

We assess whether different observations provide complementary information using the arched synthetic diffusion in Section 4.1. We compare GMI-FM (lin) constrained by the midpoint mean, the vertical variance, or both, keeping the remaining configuration fixed. We evaluate both midpoint residuals from velocity rollouts, including quantities absent from the training constraints:

$$R_\mu = \|\widehat{\mu}_{0.5} - \mu_{0.5}\|_2, \qquad R_v = |\widehat{v}_{y,0.5} - v_{y,0.5}|, \qquad R = \sqrt{R_\mu^2 + R_v^2}.$$

Table 6 shows that each single-moment model fits its supplied observation while retaining a large error in the unobserved quantity. Combining the observations reduces midpoint W_2 from 0.3840 for mean-only to 0.1892, and average W_2 at the unconstrained times from 0.2774 to 0.1710. The two observations therefore contribute complementary information about the population's location and spread.

Table 7: Redundant-mean control. Mean ± population standard deviation over three seeds; W_2 evaluates velocity rollouts.

Method	W_2 at 0.25	W_2 at 0.5	W_2 at 0.75	R_μ	$D_{0.5}$
CFM	0.3284 ± 0.0058	0.3818 ± 0.0023	0.2253 ± 0.0012	0.00476 ± 0.00119	0.00000 ± 0.00000
GMI-FM (lin)	0.3284 ± 0.0057	0.3816 ± 0.0022	0.2252 ± 0.0010	0.00534 ± 0.00228	0.00327 ± 0.00068

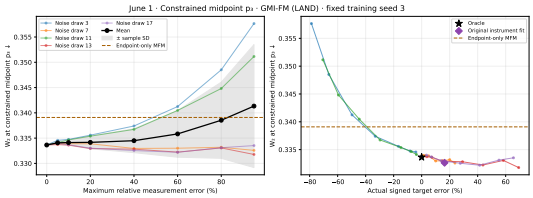

Figure 1: Observation-error sensitivity at the constrained midpoint. Empirical W_2 versus maximum relative error (left) and realised signed error (right). Colours distinguish noise draws; the black curve and band show mean ± sample standard deviation with the optimisation seed fixed. Dashed lines indicate MFM.

C.2 Redundant moment information

We test whether supplying a moment already satisfied by the reference path alters reconstruction. On the non-arched synthetic diffusion (Section 4.1, $A = 0$), we provide only the reference midpoint mean $a_{\mathrm{ref}} = \mathbb{E}_\pi[I_{\mathrm{ref}}(0.5, X_0, X_1)]$. We initialise the correction g_θ at zero and keep it trainable. Suppressing endpoint arguments, we measure the rollout mean residual and interpolant deformation as

$$R_\mu = \|\widehat{\mu}_{0.5}^{\mathrm{roll}} - a_{\mathrm{ref}}\|_2, \qquad D_{0.5} = \left(\mathbb{E}_\pi \|I_\theta(0.5) - I_{\mathrm{ref}}(0.5)\|_2^2\right)^{1/2}.$$

Table 7 shows essentially unchanged reconstruction, with paired W_2 differences below 3.3×10^{-4} across seeds and evaluated times. The small nonzero interpolant deformation is consistent with minibatch fluctuations in the constraint estimates.

C.3 Observation-error sensitivity at the constrained midpoint

We report the results in Figure 1.

C.4 Reconstruction from scarce midpoint samples

We report the results in Table 8.

C.5 Runtime

We report the results in Table 9.

D Geometric prior based on LAND

We use the data-dependent Riemannian metric associated with the locally adaptive normal distribution (LAND), following Kapuśniak et al. (2024). This metric makes the cost of motion depend on the local distribution of reference samples, allowing the path regulariser to reflect the geometry of the observed data.

Let $\mathcal{Z} = \{z_i\}_{i=1}^N \subset \mathbb{R}^d$ be a fixed reference set. At a position x, each reference point receives a weight

$$w_i(x) = \exp\left(-\frac{\|x - z_i\|_2^2}{2\gamma_{\mathrm{LAND}}^2}\right), \tag{14}$$

Table 8: Reconstruction from scarce midpoint observations. Empirical W_2 is reported as mean $\pm$ population standard deviation over five observation draws. Lowest means are bold.

Method	$n = 4$	$n = 16$	$n = 64$	$n = 256$	$n = 1024$
			Avg. W_2 at $t = 0.25$ and $t = 0.75$		
MM-CFM (empirical)	0.3751 ± 0.0284	0.2561 ± 0.0190	0.2508 ± 0.0085	0.2402 ± 0.0066	0.2400 ± 0.0092
MM-CFM (Gaussian)	0.3328 ± 0.0246	0.2500 ± 0.0263	0.2561 ± 0.0137	0.2531 ± 0.0035	0.2582 ± 0.0048
MM-MFM (empirical)	0.4482 ± 0.0430	0.3611 ± 0.0288	0.3221 ± 0.0446	0.2963 ± 0.0160	0.3144 ± 0.0215
CSB-FM	0.2782 ± 0.0295	0.2421 ± 0.0231	0.2499 ± 0.0128	0.2455 ± 0.0052	0.2511 ± 0.0018
GMI-FM (lin)	$\mathbf{0.2640} \pm 0.0682$	$\mathbf{0.2193} \pm 0.0168$	$\mathbf{0.1903} \pm 0.0093$	$\mathbf{0.1832} \pm 0.0038$	$\mathbf{0.1797} \pm 0.0018$
GMI-FM (LAND)	0.2849 ± 0.0746	0.2384 ± 0.0259	0.2040 ± 0.0130	0.1976 ± 0.0061	0.1921 ± 0.0014
			W_2 at the midpoint $t = 0.5$		
MM-CFM (empirical)	0.5245 ± 0.0744	0.2793 ± 0.0099	0.1674 ± 0.0296	0.1128 ± 0.0102	$\mathbf{0.1025} \pm 0.0134$
MM-CFM (Gaussian)	0.4738 ± 0.0779	0.2397 ± 0.0241	0.2134 ± 0.0239	0.1938 ± 0.0053	0.1943 ± 0.0108
MM-MFM (empirical)	0.4808 ± 0.0858	0.2467 ± 0.0142	$\mathbf{0.1565} \pm 0.0242$	$\mathbf{0.1094} \pm 0.0084$	0.1035 ± 0.0080
CSB-FM	$\mathbf{0.3035} \pm 0.0904$	$\mathbf{0.2243} \pm 0.0304$	0.2064 ± 0.0145	0.1911 ± 0.0040	0.1950 ± 0.0044
GMI-FM (lin)	0.3168 ± 0.1096	0.2469 ± 0.0334	0.2114 ± 0.0181	0.1963 ± 0.0063	0.1929 ± 0.0067
GMI-FM (LAND)	0.3129 ± 0.1063	0.2445 ± 0.0363	0.2129 ± 0.0172	0.1996 ± 0.0060	0.1967 ± 0.0075

Table 9: Fitting time in seconds (seed 3). Multi averages reciprocal settings. Total is the sum of path and velocity fitting times, computed before rounding; setup and sampling are excluded.

Method	Path fit	Velocity fit	Total
Arched bridge			
CFM	—	4.353	4.353
MFM	1.602	10.374	11.976
CSB-FM	2.281	7.925	10.206
GMI-FM (lin)	3.418	12.373	15.791
GMI-FM (LAND)	4.597	9.997	14.594
June 1 atmosphere			
CFM	—	16.692	16.692
MFM	8.660	49.687	58.346
CSB-FM	5.606	49.091	54.697
GMI-FM (lin)	60.513	58.219	118.732
GMI-FM (LAND)	179.252	54.571	233.823
Multi			
CFM	—	40.315	40.315
MFM	14.211	44.379	58.590
CSB-FM	558.883	83.587	642.471
GMI-FM (lin)	60.492	58.449	118.940
GMI-FM (LAND)	61.150	35.567	96.717
Semrau			
CFM	—	28.677	28.677
MFM	60.781	46.109	106.889
CSB-FM	0.176	25.729	25.905
GMI-FM (lin)	12.393	35.363	47.756
GMI-FM (LAND)	172.585	42.226	214.812

where $\gamma_{\mathrm{LAND}} > 0$ controls the size of the neighbourhood contributing to the metric. The LAND metric is the diagonal matrix

$$M_{\mathrm{LAND}}(x) = \operatorname{diag}\big(m_1(x), \ldots, m_d(x)\big),$$
$$m_j(x) = \left[\rho_{\mathrm{LAND}} + \sum_{i=1}^{N} w_i(x)(z_{i,j} - x_j)^2\right]^{-1}, \qquad j = 1, \ldots, d, \tag{15}$$

with $\rho_{\mathrm{LAND}} > 0$ ensuring that the matrix remains finite and positive definite. The kernel weights are used without normalisation.

The denominator measures the local spread of reference samples along each coordinate. Greater spread lowers the cost of motion along that coordinate. Far from the reference samples, the kernel contributions vanish and the metric approaches $\rho_{\mathrm{LAND}}^{-1} I_d$. A small stabiliser therefore makes motion through regions with little reference support costly.

For an endpoint pair $Z = (X_0, X_1) \sim \pi$, write $\dot{I}_\theta(t, Z) = \partial_t I_\theta(t, X_0, X_1)$. We define the corresponding path regulariser as

$$R_{\mathrm{LAND}}(\theta) = \mathbb{E}_{Z\sim\pi} \int_0^1 \dot{I}_\theta(t, Z)^\top M_{\mathrm{LAND}}(I_\theta(t, Z))\, \dot{I}_\theta(t, Z)\, \mathrm{d}t. \tag{16}$$

This is the expected kinetic energy measured in the LAND metric. In GMI-FM (LAND), we set $R_{\mathrm{prior}} = R_{\mathrm{LAND}}$ in (6). The metric selects among interpolants satisfying the generalised moment constraints by favouring motion supported by the reference geometry. After fitting the interpolant, the velocity field is learned through the same conditional flow-matching stage.

For GMI-FM (LAND), the reference set is constructed from endpoint samples and held fixed during optimisation. The bandwidth γ_{LAND}, stabiliser ρ_{LAND}, and reference-set size N determine the resulting geometry and are specified with the experimental settings.

E Algorithm

Algorithm 1 Generalized Moment-Informed Flow Matching

Require: Coupling $\pi \in \Pi(p_0, p_1)$, reference interpolant I_{ref}, regulariser R_{base}, observations $\{(\tau_k, f_k, a_k)\}_{k=1}^K$, prescribed residual scalings, hyperparameters $\{(\eta_k, \rho_k, \nu_{\max,k})\}_{k=1}^K$, batch size B, step counts (N_A, N_B)

Define $I_\theta(t, x_0, x_1) = I_{\mathrm{ref}}(t, x_0, x_1) + t(1-t) g_\theta(t, x_0, x_1)$
Initialise parameters θ, ψ and multipliers $\nu_k \leftarrow \mathbf{0} \in \mathbb{R}^{d_k}$ for $k = 1, \dots, K$
Stage A: learn the moment-constrained interpolant
for $s = 1, \dots, N_A$ **do**
    Sample $(x_0^{(i)}, x_1^{(i)}) \sim \pi$ and independent $t^{(i)} \sim \mathcal{U}(0,1)$ for $i = 1, \dots, B$
    Compute $\widehat{R}_{\mathrm{base}}(\theta)$ using $\dot{I}_\theta = \partial_t I_\theta$
    Compute the raw residuals $\widehat{c}_k(\theta)$ using (20) with $w_i = 1/B$
    Apply the prescribed scalings to obtain $r_k(\theta)$ for $k = 1, \dots, K$
    Take a gradient step on θ using the minibatch objective $\widehat{\mathcal{L}}_A(\theta, \nu)$ from (21)
    Recompute the scaled residuals $r_k(\theta)$ at the updated parameters
    Update each multiplier using (22), with detached residuals and bounds $\nu_{\max,k}$
end for
Stage B: learn the marginal velocity field
Set $\hat{\theta} \leftarrow \theta$ and freeze $\hat{\theta}$
for $s = 1, \dots, N_B$ **do**
    Sample $(x_0^{(i)}, x_1^{(i)}) \sim \pi$ and independent $t^{(i)} \sim \mathcal{U}(0,1)$ for $i = 1, \dots, B$
    Compute $I_{\hat{\theta}}(t^{(i)}, x_0^{(i)}, x_1^{(i)})$ and $\partial_t I_{\hat{\theta}}(t^{(i)}, x_0^{(i)}, x_1^{(i)})$
    Take a gradient step on ψ using the minibatch objective $\widehat{\mathcal{L}}_B(\psi; \hat{\theta})$ from (8)
end for
return Learned interpolant $I_{\hat{\theta}}$ and velocity field v_ψ

F Implementation Details

This section describes the construction of the observation functions, the core training procedure, and the rule used to initialise the constraint hyperparameters.

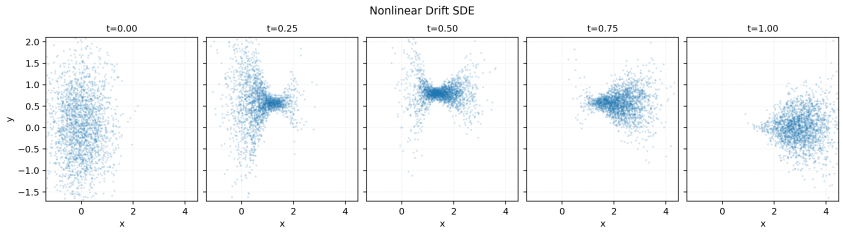

Figure 2: Marginal evolution of the synthetic diffusion. The population follows a vertically displaced route and concentrates near the bottleneck before spreading again.

F.1 Data preparation and observation models

Synthetic diffusion. We generate the synthetic dataset using a two-dimensional diffusion with an arched displacement and a localised bottleneck. In physical time $s \in [0, T]$, the process satisfies

$$\mathrm{d}X_s = v_x \,\mathrm{d}s + \sigma_x \,\mathrm{d}W_s^x, \tag{17}$$

$$\mathrm{d}Y_s = \big[h'(s) - \kappa\,\phi(X_s)(Y_s - h(s))\big]\,\mathrm{d}s + \sigma_y\big(1 - \delta\,\phi(X_s)\big)\,\mathrm{d}W_s^y, \tag{18}$$

where W^x and W^y are independent standard Brownian motions, and

$$\phi(x) = \exp\!\left(-\frac{(x-c)^2}{2w^2}\right), \qquad h(s) = A\sin\!\left(\frac{\pi s}{T}\right).$$

The horizontal coordinate evolves with constant drift v_x and diffusion coefficient σ_x. The function h introduces a vertical displacement of amplitude A, which reaches its maximum at $s = T/2$ and vanishes at the endpoints. The gate ϕ localises the restoring drift towards $h(s)$ around $x = c$, with width w and strength controlled by κ. The factor $1 - \delta\phi(X_s)$ also reduces vertical diffusion in this region. Together, these terms produce a population that follows an arched route and contracts near the bottleneck before spreading again. The non-arched control uses $A = 0$.

We initialise the process from a Gaussian source distribution and simulate $N = 200{,}000$ particles to construct the reference marginals. We normalise time as $t = s/T$ and denote the resulting process by $Z_t = (X_{Tt}, Y_{Tt})$. All methods receive samples from the endpoint populations. For the main synthetic experiment, we compute the sample mean and vertical variance at $\tau = 0.5$ and supply these aggregate quantities to GMI-FM and CSB-FM. The marginals at $t = 0.25$ and $t = 0.75$ provide reconstruction targets without aggregate constraints.

Atmospheric measurements. We normalise each DMPS histogram to obtain a probability distribution over particle diameter. We then transform diameter to logarithmic coordinates and standardise using the endpoint distributions. The observation function combines a Mie-scattering response with an approximate inlet transmission function. Its refractive-index and gain parameters are calibrated using the separate January–February measurements and subsequently held fixed. At the constraint time, the nephelometer scattering coefficient divided by the CPC particle concentration supplies the target mean scattering cross section. The observation function and target are both divided by the same scale $s_{\mathrm{opt}} = 0.01\,\mu\mathrm{m}^2$.

Multi: composition observations. We use the supplied 100-dimensional PCA representation and standardise its coordinates using the endpoint populations. For each constraint day, we fit a classifier D_τ to the annotated cells from that day, using a stratified training and validation split. The classifier is trained with class-weighted cross-entropy and held fixed during interpolant and velocity training. Its seven soft outputs define the observation function $f(x) = D_\tau(x)$. We compute the target by averaging these outputs over all observed constraint-day cells:

$$a_\tau = \frac{1}{N_\tau}\sum_{i=1}^{N_\tau} D_\tau(x_i^{(\tau)}).$$

The target is therefore the empirical expectation of the same fixed map used in the constraint loss. The classifier forms part of the supplied observation model, and the composition residual measures agreement between the generated classifier-output averages and this empirical target.

We use the supplied 100-dimensional PCA representation and standardise its coordinates using the endpoint populations. For each constraint day, we fit a classifier D_τ to the annotated cells from that day, using a stratified training and validation split. The classifier is trained with class-weighted cross-entropy and held fixed during interpolant and velocity training. Its seven soft outputs define the observation function $f(x) = D_\tau(x)$, while the target is the empirical vector of annotated cell-type proportions.

The classifier is therefore part of the supplied observation model. Its outputs provide differentiable surrogates for class-membership indicators; their population averages need not equal the annotated proportions exactly. Consequently, the composition residual measures agreement with the supplied classifier-based constraint.

Semrau: marker observations and calibration. We represent each cell using eight principal components of normalised log-expression and four additional marker coordinates for Sox2, Tbx3, Nanog, and Ppia. Gene selection, PCA fitting, and coordinate scaling use endpoint training cells. For a marker gene g, the additional coordinate is

$$x_g = \frac{\sqrt{n_g}}{s_g},$$

where n_g is its raw UMI count and $s_g > 0$ is an endpoint-derived scale.

To map bulk expression to mean single-cell counts, we consider triplets of calibration times $\mathcal{T}_g = \{s_1, s_2, s_3\}$ with paired bulk and single-cell measurements. Let $b_g(s)$ denote the library-normalised bulk abundance and $\overline{n}_g(s)$ the mean single-cell UMI count. For each triplet, we fit a gene-specific proportional calibration:

$$\widehat{c}_g = \arg\min_{c \geq 0} \sum_{s \in \mathcal{T}_g} \big(c\, b_g(s) - \overline{n}_g(s)\big)^2 = \frac{\sum_{s \in \mathcal{T}_g} b_g(s)\overline{n}_g(s)}{\sum_{s \in \mathcal{T}_g} b_g(s)^2}, \tag{19}$$

assuming a nonzero denominator. Applying the fitted gain to the bulk measurement at the observation time gives $\widehat{m}_g = \widehat{c}_g b_g(s_{\mathrm{obs}})$. The generalised-moment observation in model coordinates is therefore

$$f_g(x) = x_g^2, \qquad a_g = \frac{\widehat{m}_g}{s_g^2}.$$

We impose these observations for Sox2 and Tbx3 at 6 hours. The calibration and resulting targets are fixed during trajectory learning and shared across optimisation seeds.

F.2 Endpoint coupling and neural parameterisation

Endpoint pairs are drawn from a fixed coupling between the empirical source and target populations. We use global optimal-transport couplings with squared Euclidean cost, or their entropically regularised counterparts. A discrete coupling is represented by endpoint pairs and their associated masses. Expectations are evaluated either by sampling pairs according to these masses or by summing over the weighted coupling support.

The basic interpolant parameterisation is

$$I_\theta(t, x_0, x_1) = (1-t)x_0 + tx_1 + t(1-t)g_\theta(t, x_0, x_1).$$

The correction network receives the time and both endpoints, while the velocity network receives the time and current position. Their input dimensions are therefore $2d+1$ and $d+1$, respectively, and both output vectors in $\mathbb{R}^d$. Time derivatives of the interpolant are computed by automatic differentiation. For LAND regularisation, the metric reference samples are drawn from the endpoint populations and held fixed.

F.3 Interpolant and velocity training

Constraint residuals. For an observation (τ_k, f_k, a_k), the empirical residual is

$$\widehat{c}_k(\theta) = \sum_{i=1}^{B} w_i f_k\Big(I_\theta(\tau_k, x_0^{(i)}, x_1^{(i)})\Big) - a_k, \qquad \sum_{i=1}^{B} w_i = 1. \tag{20}$$

For independently sampled coupling pairs, $w_i = 1/B$; for integration over a discrete coupling, w_i is its normalised support mass.

We distinguish raw residuals from the scaled residuals used for optimisation. Scaling can account for the units of an observation or the dimension of a moment block. In particular, dividing a block by the square root of its dimension prevents its quadratic penalty from growing solely because it contains more coordinates. Reported observation errors are evaluated in the raw or explicitly stated relative units.

Stage A: augmented-Lagrangian optimisation. Writing $r_j(\theta)$ for a scaled constraint residual, the augmented-Lagrangian objective has the form

$$\mathcal{L}_A(\theta, \nu) = R_{\mathrm{base}}(\theta) + \sum_j \eta_j \left[\nu_j^\top r_j(\theta) + \frac{\rho_j}{2} \|r_j(\theta)\|_2^2 \right], \tag{21}$$

where R_{base} is the implemented path regulariser. For groups containing several residual blocks, the terms are combined using the prescribed block sums or averages. The same convention is used during hyperparameter calibration.

The multipliers are initialised at zero. Primal optimisation updates the interpolant parameters, while dual updates take the form

$$\nu_j \leftarrow \mathrm{clip}(\nu_j + \rho_j r_j, -\nu_{\max,j}, \nu_{\max,j}), \tag{22}$$

with componentwise clipping and detached residuals. The numerical schedule determines how many primal updates are performed between dual updates. In stochastic training, the objective is estimated from coupled endpoint batches. For the atmospheric benchmark, we instead use weighted endpoint quadrature and time quadrature, with L-BFGS inner optimisation before multiplier updates.

Stage B: conditional velocity regression. After Stage A, we freeze the selected interpolant $I_{\hat{\theta}}$ and train the velocity field against its conditional velocities:

$$\mathcal{L}_B(\psi) = \mathbb{E}_{\pi,t} \left[\left\| v_\psi(t, I_{\hat{\theta}}(t, X_0, X_1)) - \partial_t I_{\hat{\theta}}(t, X_0, X_1) \right\|_2^2 \right].$$

Where constant Gaussian input noise is used, the regression position is replaced by $I_{\hat{\theta}}(t, X_0, X_1) + \sigma_B \epsilon$, with $\epsilon \sim \mathcal{N}(0, I_d)$, while the conditional velocity target is unchanged. This corresponds to smoothing the conditional position distributions.

Both training objectives can be evaluated directly from endpoint pairs, interpolant values, and time derivatives. Numerical integration of the learned velocity field is used to generate marginal samples for evaluation.

F.4 Initialisation of constraint hyperparameters

We initialise the relative constraint strength by comparing gradients of the path regulariser and the quadratic constraint terms. Geometric-prior parameters and residual normalisation are fixed before this calculation, since both affect the gradient scales.

Let $Q_j(\theta)$ denote the quadratic term for constraint group j with its penalty coefficient removed, including the actual residual scaling and block averaging. At a calibration state, we compute

$$g_0 = \nabla_\theta R_{\mathrm{base}}, \qquad g_j = \nabla_\theta Q_j.$$

With zero multipliers, the contribution of group j to the constraint gradient is $\eta_j \rho_j g_j$. We therefore calibrate the effective strength $\kappa_j = \eta_j \rho_j$ through

$$\kappa_{j,\mathrm{bal}} = \frac{\|g_0\|_2}{\|g_j\|_2}, \qquad \kappa_{j,0} = s \ \mathrm{median}_{\text{calibration seeds}} \ \kappa_{j,\mathrm{bal}}, \tag{23}$$

using nondegenerate gradient measurements and a conservative factor s, typically 0.7.

Calibration uses a fixed endpoint-coupling sample, or the complete weighted coupling support, and fixed time quadrature. When the initial geometric gradient is strongly transient, we measure the ratio after a short geometry-only warmup. The same warmup is included in the subsequent training run. If

several constraint groups are active, their strengths share a common gradient budget, and we check the norm of their combined gradient rather than balancing each group independently against the full base gradient.

Gradient balance provides an initial proposal. Frequent multiplier updates can accumulate a substantial dual force, so the proposal must also account for the dual-update schedule. Dividing the initial balance by the planned number of active dual updates provides a conservative starting scale when persistent residuals would otherwise overdrive optimisation. This is a heuristic adjustment, not a convergence guarantee.

We confirm the proposal with a small multiplicative sweep. Diagnostics include contraction and stability of the active residuals, the base loss, multiplier clipping, and the relative norms of the base and constraint gradients. We tune $\eta_j \rho_j$ and the effective multiplier bound $\eta_j \nu_{\max,j}$ jointly, avoiding independent searches over redundant augmented-Lagrangian scales.

F.5 Neural architectures

The path-correction and velocity networks are multilayer perceptrons. The path network receives the concatenated input $(t, x_0, x_1) \in \mathbb{R}^{2d+1}$ and outputs a correction in $\mathbb{R}^d$. The velocity network receives $(t, x) \in \mathbb{R}^{d+1}$ and outputs a velocity in $\mathbb{R}^d$. Hidden-layer widths and activation functions are given in Table 10. The output layers are linear.

Table 10: Base neural-network architectures. Lists specify the width of each hidden layer. Semrau uses different path-network sizes for the linear and LAND priors, with a common velocity architecture.

Benchmark	d	Path hidden layers	Velocity hidden layers	Activation
Synthetic	2	$[128, 128]$	$[128, 128]$	SiLU
Atmospheric	1	$[128, 128]$	$[128, 128, 128]$	SiLU
Multi	100	$[256, 256, 256]$	$[256, 256, 256]$	SELU
Semrau, linear prior	12	$[96, 96]$	$[128, 128, 128]$	SiLU
Semrau, LAND prior	12	$[48]$	$[128, 128, 128]$	SiLU

Within each benchmark, the compared methods use the same velocity-network architecture. CFM requires only the velocity network, while methods with learned interpolants additionally use a path network during Stage A.

F.6 Training budgets and stopping

We specify a maximum computational budget separately for interpolant learning and velocity regression. When early stopping is enabled, training terminates when either its stopping criterion is satisfied or the maximum budget is exhausted. Fixed-budget runs execute the prescribed number of updates. The checkpoint retained for evaluation is selected according to the benchmark's checkpoint rule; its update index can therefore differ from the number of updates actually executed.

Stage A: interpolant learning. The path-training budgets are summarised in Table 11. For stochastic optimisation, budgets count primal parameter updates. For the atmospheric benchmark, Stage A instead uses nested optimisation: each outer iteration performs an L-BFGS inner solve before updating the multipliers. Its budget is consequently expressed in outer iterations. LAND training additionally uses a geometry-only warmup, whose computation is included in the Stage-A cost.

During Stage A, we monitor the active observation residuals, the path regulariser, and the multiplier magnitudes. These diagnostics assess constraint satisfaction and optimisation stability. They are distinguished from the criterion used to select a checkpoint.

Stage B: velocity regression. With the interpolant or teacher held fixed, we train the velocity network with Adam, using batches of 256 endpoint pairs and an initial learning rate of 10^{-3}. The maximum Stage-B budget is shared across the compared methods within each benchmark.

For patience-based early stopping, we separately track the best validation score and a reference score for significant improvements. A decrease greater than δ relative to this reference resets the patience

counter and updates the reference. Training stops after P consecutive checks without a reset, subject to a minimum training duration, or upon reaching the maximum budget. Checkpoint selection retains the lowest validation score among eligible checkpoints, including improvements smaller than δ.

Table 11: Training budgets for the main experiments. Stage-A entries apply to learned interpolants; CFM has no Stage A, and CSB-FM uses a numerical bridge solver. Stage-B caps count optimiser updates.

Benchmark	Stage-A budget	Stage-B cap	Stage-B checkpoint rule
Synthetic	600 updates for MFM; 600 for GMI-FM	9,600	Final iterate
Atmospheric	Up to 240 outer iterations	30,000	Final iterate
Multi	Up to 2,400 updates	30,000	Minimum validation W_2, with early stopping
Semrau	6,000 updates	18,000	Minimum validation W_2 within the fixed budget

For Multi, we evaluate W_2 against the unconstrained-day validation marginal every 300 updates, with a minimum training duration of 1,200 updates, patience $P = 6$, and threshold $\delta = 10^{-3}$. All evaluated checkpoints are eligible for selection, starting at update 300; the minimum training duration restricts stopping only. After three checks without a significant improvement, the learning rate is reduced once per plateau, with at most two reductions: $10^{-3} \to 3 \times 10^{-4} \to 10^{-4}$. Learning-rate reductions do not reset the patience counter.

For Semrau, all 18,000 velocity updates are executed. We maintain an exponential moving average of the network parameters with decay 0.99 and evaluate it every 600 updates using W_2 against the reserved 36-hour validation population. Checkpoints become eligible at update 3,000, and the eligible checkpoint with the lowest validation score is retained. This is checkpoint selection within a fixed training budget.

The synthetic and atmospheric velocity fits retain their final iterates after the prescribed budgets. For comparisons involving early stopping or intermediate checkpoint selection, we distinguish the executed update count from the retained checkpoint's update index.

We freeze the selected interpolant and train the velocity network with Adam, using batches of 256 endpoint pairs and an initial learning rate of 10^{-3}. The maximum Stage-B budget is shared across the compared methods within each benchmark.

For runs using patience-based early stopping, we evaluate a specified validation W_2 at regular intervals. An improvement larger than a threshold δ resets the patience counter. Training stops after P consecutive checks without such an improvement, subject to a minimum training duration. The retained checkpoint minimises the validation score among the eligible checkpoints, including improvements smaller than δ.

Table 12: Training budgets for the main experiments. Stage-A entries concern learned interpolants: CFM has no Stage A, while CSB-FM constructs its teacher using a numerical bridge solver. Stage-B budgets count optimiser updates.

Benchmark	Stage-A budget	Stage-B budget	Stage-B checkpoint rule
Synthetic	600 updates for MFM; 600 for GMI-FM	9,600	Final iterate
Atmospheric	Up to 240 outer iterations	30,000	Final iterate
Multi	Up to 2,400 updates	30,000	Minimum validation W_2, with early stopping
Semrau	6,000 updates	18,000	Minimum validation W_2 within the fixed budget

For Multi velocity training, validation is performed every 300 updates, with a minimum duration of 1,200 updates, patience $P = 6$, and improvement threshold $\delta = 10^{-3}$. The learning rate can be reduced after validation plateaus.

For Semrau, all 18,000 velocity updates are executed. We maintain an exponential moving average of the network parameters with decay 0.99 and evaluate it every 600 updates. Checkpoints become eligible for selection after 3,000 updates, and the eligible checkpoint with the lowest validation W_2 is retained. This is checkpoint selection within a fixed training budget.

The synthetic and atmospheric velocity fits retain their final iterates after the prescribed budgets. For comparisons involving early stopping or intermediate checkpoint selection, we distinguish the executed update count from the retained checkpoint's update index.

F.7 Evaluation and reproducibility

We distinguish observation residuals of the direct interpolant from those of samples generated by the learned velocity field. Stage B leaves the direct interpolant unchanged, but finite regression accuracy and numerical integration can produce a difference between its marginals and the velocity rollouts.

Marginal reconstruction is measured using empirical balanced W_2: the square root of the optimal transport cost with squared Euclidean ground cost in the stated model coordinates. Unless otherwise specified, we report the mean and population standard deviation over seeds $\{3, 7, 11, 13, 17\}$. When a score averages several times or reciprocal settings, we first compute that average within each seed and then aggregate across seeds.

The implementation uses PyTorch for neural models and automatic differentiation, with NumPy, SciPy, and optimal-transport solvers for data processing and evaluation. Random seeds control network initialisation and sampling. The numerical precision, network sizes, optimiser settings, training budgets, and integration settings are recorded with each experimental configuration.

G Generalised Moment-Constrained Schrödinger Bridges

G.1 Entropy projection and the optimal law

Population assumptions. Throughout, R is the probability law of Brownian motion with diffusion coefficient σI_d and initial law p_0. We assume that the constraint set in Eq. (10) contains a law of finite relative entropy. For bounded continuous features, this guarantees a unique minimiser: the constraints are weakly closed, relative-entropy sublevel sets are tight, and relative entropy is lower semicontinuous and strictly convex on densities. The exponential representation additionally assumes attainment of the entropy dual with finite feature multipliers and endpoint potentials for which the expectations used below are well defined. Dual attainment does not follow merely from boundedness of f. For unbounded moment features, the corresponding integrability, normalisability, and dual-attainment conditions are required. The explicit mean case in Appendix H.3 can also be established directly by Gaussian translation.

These are assumptions on population laws. An empirical terminal measure is singular with respect to the terminal law of a nondegenerate Brownian reference, so literal finite-entropy statements for R should not be applied to a finite sample measure. Discrete coupling computations serve as numerical approximations. The shared-coupling identity of Proposition 1, in contrast, applies to empirical couplings as well.

Reduction to three positions. For one observation time, write $\gamma = P_{0,\tau,1}$ and let $P^{x,z,y}$ denote a regular conditional path law given $(X_0, X_\tau, X_1) = (x, z, y)$. The chain rule for relative entropy yields

$$\begin{aligned} \mathrm{KL}(P\|R) = \mathrm{KL}(\gamma\|R_{0,\tau,1}) \\ + \int \mathrm{KL}(P^{x,z,y}\|R^{x,z,y})\,\gamma(dx, dz, dy). \end{aligned} \tag{24}$$

All the constraints depend only on γ. For any feasible γ, the second term is minimised at zero by retaining the reference conditional law. Therefore Eq. (10) is equivalent to

$$\min_{\gamma} \mathrm{KL}(\gamma\|R_{0,\tau,1}) \quad \text{s.t.} \quad \gamma_0 = p_0, \quad \gamma_1 = p_1, \quad \int f(z)\,\gamma(dx, dz, dy) = a, \tag{25}$$

where the subscripts $0, 1$ denote the initial and terminal coordinates. Under R, conditioning on the three positions leaves independent Brownian bridges on $[0, \tau]$ and $[\tau, 1]$. This proves the concatenated-bridge representation without requiring a finite multiplier.

Exponential form and optimality. Apply the entropy-projection duality to Eq. (25) (Csiszár, 1975; Léonard, 2013). We use the sign convention that the Lagrangian contains $-\lambda^\top(\mathbb{E}_P f(X_\tau) - a)$. At a dual optimum, its first-order condition gives

$$\gamma^\star(dx, dz, dy) = \varphi_\lambda(x) e^{\lambda^\top f(z)} \psi_\lambda(y) p_0(dx) k_\tau(x, z) k_{1-\tau}(z, y)\, dz\, dy, \tag{26}$$

with normalisation absorbed into the endpoint factors. Conversely, any normalised law of this form satisfying the constraints is the unique finite-entropy minimiser. Indeed, for any other feasible P for which the terms are finite,

$$\begin{aligned} \mathrm{KL}(P\|R) - \mathrm{KL}(P^\star\|R) &= \mathrm{KL}(P\|P^\star) + \mathbb{E}_P \log \frac{dP^\star}{dR} - \mathbb{E}_{P^\star} \log \frac{dP^\star}{dR} \\ &= \mathrm{KL}(P\|P^\star) \geq 0. \end{aligned} \tag{27}$$

The last equality uses the common endpoint marginals and feature expectation.

Endpoint coupling and conditional tilt. The optimal coupling for a fixed multiplier is

$$\pi^\lambda \in \operatorname*{arg\,min}_{\pi \in \Pi(p_0, p_1)} \left\{ \mathrm{KL}(\pi\|R_{0,1}) - \int \log Z_\lambda(x, y)\, \pi(dx, dy) \right\}. \tag{28}$$

One must determine the endpoint scaling factors for the modified kernel $K_\lambda = k_1 Z_\lambda$. Simply tilting an already solved endpoint-only bridge will generally change its endpoint marginals.

For each multiplier for which the endpoint-constrained problem is well defined, denote its solution by

$$P^\lambda \in \operatorname*{arg\,min}_{P:\, P_0 = p_0,\, P_1 = p_1} \{\mathrm{KL}(P\|R) - \lambda^\top \mathbb{E}_P f(X_\tau)\}. \tag{29}$$

The desired multiplier satisfies

$$m(\lambda) := \mathbb{E}_{P^\lambda} f(X_\tau) = \int \nabla_\lambda \log Z_\lambda(x, y)\, \pi^\lambda(dx, dy) = a. \tag{30}$$

The derivative here is taken at fixed endpoints. Comparing optimality inequalities at two multipliers gives $(\lambda_2 - \lambda_1)^\top (m(\lambda_2) - m(\lambda_1)) \geq 0$. In the scalar case this permits monotone root finding, with the endpoint coupling updated for each multiplier. Boundary targets require separate treatment. For instance, if $f \leq 1$ and $a = 1$, the constraint forces $f(X_\tau) = 1$ almost surely; if this level set has zero reference probability, no finite-entropy feasible law exists.

Several observation times. The same argument applies to the skeleton $(X_0, X_{\tau_1}, \ldots, X_{\tau_K}, X_1)$. Under dual attainment,

$$\frac{dP^\star}{dR} = \varphi_\lambda(X_0) \exp\left(\sum_{k=1}^{K} \lambda_k^\top f_k(X_{\tau_k}) \right) \psi_\lambda(X_1). \tag{31}$$

Conditional on the endpoints, the intermediate positions have a jointly tilted Gaussian bridge law. Before tilting, their means are $\ell_{\tau_k}(x, y)$ and covariance blocks are $\sigma^2(\min\{\tau_i, \tau_j\} - \tau_i \tau_j) I_d$. They must be sampled jointly or through the appropriate conditional laws. Given the full skeleton, all successive subbridges remain conditionally independent reference bridges.

G.2 Sampling and conditional flow matching

Sampling the intermediate position. Assume that λ satisfying Eq. (30) and the corresponding coupling π^λ are available. First draw $(x, y) \sim \pi^\lambda$. For bounded features, take a known bound $M_\lambda \geq \sup_z \lambda^\top f(z)$, repeatedly propose $z \sim G_{xy}$, and accept with probability

$$\exp\big(\lambda^\top f(z) - M_\lambda\big). \tag{32}$$

The accepted position has law $q_\lambda(\cdot \mid x, y)$. The endpoints are held fixed throughout this rejection loop; resampling them after a rejection would reweight π^λ. The acceptance probability is $Z_\lambda(x, y)e^{-M_\lambda}$, so strong tilts can make this procedure expensive. The sampler itself does not require Z_λ, but constructing the modified endpoint coupling does.

For mean and raw second-moment features, write the tilt as $\lambda^\top f(z) = b^\top z + z^\top H z$, with H symmetric, and set $c = \sigma^2\tau(1-\tau)$ and $m = \ell_\tau(x, y)$. If $S^{-1} = c^{-1}I_d - 2H$ is positive definite, completing the square gives

$$q_\lambda(\cdot \mid x, y) = \mathcal{N}(\mu, S), \qquad \mu = S(c^{-1}m + b), \qquad \log Z_\lambda = \frac{1}{2}\log\frac{\det S}{c^d} + \frac{1}{2}\mu^\top S^{-1}\mu - \frac{\|m\|^2}{2c}. \tag{33}$$

Thus neither conditional rejection sampling nor numerical integration for the normaliser is necessary in this case. Computing the coupling and moment multipliers remains a numerical task in general.

Sampling a bridge position and its velocity target. For a sampled time t, take

$$(s, r, A, B) = \begin{cases} (0, \tau, x, z), & 0 < t < \tau, \\ (\tau, 1, z, y), & \tau < t < 1. \end{cases} \tag{34}$$

Draw $\xi \sim \mathcal{N}(0, I_d)$ independently and define

$$X_t = m_t + \sqrt{c_t}\,\xi, \qquad m_t = \frac{r-t}{r-s}A + \frac{t-s}{r-s}B, \qquad c_t = \sigma^2\frac{(t-s)(r-t)}{r-s}. \tag{35}$$

At $t = \tau$, return the sampled middle position z. These are exact one-time samples from $P^\star$. To draw an entire trajectory, sample the correlated Brownian subbridges; independent Gaussian draws at different times do not produce their joint path law.

The conditional Gaussian density in Eq. (35) solves the continuity equation with velocity

$$\begin{aligned} v_{\mathrm{cond}}(t, w \mid A, B) &= \dot{m}_t + \frac{\dot{c}_t}{2c_t}(w - m_t) \\ &= \frac{B-A}{r-s} + \frac{s+r-2t}{2(t-s)(r-t)}(w - m_t) \\ &= \frac{1}{2}\left(\frac{B-w}{r-t} + \frac{w-A}{t-s}\right). \end{aligned} \tag{36}$$

This is a probability-flow velocity, not a derivative of a Brownian sample path. The forward SDE drift of the conditional bridge is instead $(B-w)/(r-t)$, with diffusion coefficient σI_d. Averaging the conditional velocity at a fixed location gives

$$v^\star(t, w) = \mathbb{E}[v_{\mathrm{cond}}(t, X_t \mid A, B) \mid X_t = w]. \tag{37}$$

Mixtures of the conditional continuity equations show that this field transports $P_t^\star$. Conditional flow matching regresses against Eq. (36) to learn this expectation, as in Gaussian bridge matching (Lipman et al., 2022; Tong et al., 2023b). Under well-posed-flow assumptions its ODE reproduces the marginals. Recovering the stochastic path law requires retaining the diffusion, for example by learning the forward drift instead.

Time weighting. The conditional velocity variance diverges near the anchors $0, \tau, 1$: its Gaussian contribution has mean squared norm $d\sigma^2(s+r-2t)^2/[4(r-s)(t-s)(r-t)]$. Consequently, an unweighted continuous-time squared regression loss can be infinite. With finite second moments of the skeleton, the weighted objective

$$\mathbb{E}\left[w(t)\,\|v_\psi(t, X_t) - v_{\mathrm{cond}}(t, X_t \mid A, B)\|^2\right], \qquad w(t) = (t-s)(r-t), \tag{38}$$

has an integrable target. Since $w(t) > 0$ in every segment interior and depends only on time, its pointwise population minimiser is still $v^\star$. This provides a direct training construction without simulating an ODE during training.

G.3 Assumptions and proof for the SB comparison

We detail the assumptions and notation used in Proposition 2.

Reference process and constrained bridge. Let p_0, p_1 have finite second moments, and let R be the law of

$$dX_t = \sigma\, dW_t, \qquad X_0 \sim p_0, \qquad \sigma > 0,$$

where W is a standard d-dimensional Brownian motion independent of X_0. Fix $\tau \in (0,1)$, $a \in \mathbb{R}$, and $f \in C^2(\mathbb{R}^d;\mathbb{R})$ with $\sup_x \|\nabla^2 f(x)\|_{\mathrm{op}} \le H < \infty$.

Assume that the corresponding constrained SB has a minimiser $P^\star$ with finite relative entropy and $\mathbb{E}_{P^\star}[|f(X_\tau)|] < \infty$. Assume also that the entropy dual is attained with a finite moment multiplier $\lambda^\star$. We use the sign convention in which the Lagrangian contains $-\lambda(\mathbb{E}_P[f(X_\tau)] - a)$.

Set $\pi = P^\star_{0,1}$. Writing $P^{\star,z}$ and R^z for the conditional path laws given the endpoint pair $Z = z$, the dual representation from Appendix G.1 gives

$$\frac{dP^{\star,z}}{dR^z}(\omega) = \frac{\exp(\lambda^\star f(\omega_\tau))}{\mathcal{Z}_{\lambda^\star}(z)}, \qquad \mathcal{Z}_{\lambda^\star}(z) = \mathbb{E}_{R^z}[\exp(\lambda^\star f(X_\tau))]\,.$$

The normalisers are assumed positive and finite π-almost surely. The finite entropy and feature expectation of $P^\star$ imply that $\log \mathcal{Z}_{\lambda^\star} \in L^1(\pi)$.

Interpolant family. Consider the parameterisation in section 3 with the linear reference $\ell_t(x_0, x_1) = (1-t)x_0 + tx_1$. Assume that its paths are almost surely absolutely continuous in time. Using R_{lin} from (7), with $\alpha > 0$, define

$$\Theta_a = \{\theta \in \Theta : \mathbb{E}_\pi[f(I_\theta(\tau, Z))] = a, \quad R_{\mathrm{lin}}(\theta) < \infty\}\,.$$

Assume that Θ_a is nonempty and that a minimiser

$$\theta^\star \in \operatorname*{arg\,min}_{\theta \in \Theta_a} R_{\mathrm{lin}}(\theta)$$

exists.

For $\theta \in \Theta_a$, let

$$Q^\sigma_\theta = \mathrm{Law}\big((I_\theta(t, Z) + \sigma B_t)_{t\in[0,1]}\big), \qquad \varepsilon_{\mathrm{app}} = \inf_{\theta \in \Theta_a} \mathrm{KL}(Q^\sigma_\theta \| P^\star),$$

where $Z \sim \pi$ and B is an independent standard d-dimensional Brownian bridge.

Proof of Proposition 2. For each feasible θ, define the change in the feature expectation caused by the Brownian perturbation:

$$b_\theta = \mathbb{E}_{Q^\sigma_\theta}[f(X_\tau)] - a.$$

Since the deterministic interpolant satisfies the moment constraint and $B_\tau \sim \mathcal{N}(0, \tau(1-\tau)I_d)$, Taylor's theorem and the Hessian bound give

$$|b_\theta| \le \frac{H}{2} d\sigma^2 \tau(1-\tau).$$

The first-order Taylor term has zero expectation because the perturbation is independent and centred.

All the path laws Q^σ_θ and $P^\star$ have the same endpoint coupling π. The entropy chain rule, the conditional exponential representation of $P^\star$, and the Cameron–Martin identity therefore yield

$$\mathrm{KL}(Q^\sigma_\theta \| P^\star) = \frac{R_{\mathrm{lin}}(\theta)}{2\alpha\sigma^2} - \lambda^\star b_\theta + C,$$

where $C = \mathbb{E}_\pi[\log \mathcal{Z}_{\lambda^\star}(Z)] - \lambda^\star a$ is finite and independent of θ.

By optimality of $\theta^\star$, for every $\theta \in \Theta_a$,

$$\begin{aligned} &\mathrm{KL}(Q^\sigma_{\theta^\star} \| P^\star) - \mathrm{KL}(Q^\sigma_\theta \| P^\star) \\ &\quad \le \lambda^\star(b_\theta - b_{\theta^\star}) \\ &\quad \le |\lambda^\star| H d\sigma^2 \tau(1-\tau). \end{aligned}$$

Taking the infimum over $\theta \in \Theta_a$ proves the upper bound. The lower bound follows directly from the definition of $\varepsilon_{\mathrm{app}}$. □

H Interpretation of the Regularization

H.1 Relative entropy of the auxiliary noisy path

Proof of Proposition 1. Condition on $Z = (X_0, X_1) \sim \pi$. The difference $h_\theta(\cdot, Z) = I_\theta(\cdot, Z) - I_{\mathrm{ref}}(\cdot, Z)$ lies in the Cameron–Martin space $H_0^1([0,1]; \mathbb{R}^d)$ of a pinned Brownian bridge. Its squared Cameron–Martin norm for covariance scaled by σ^2 is $\sigma^{-2} \int_0^1 \|\partial_t h_\theta(t, Z)\|^2 dt$. The relative entropy between the two conditional Gaussian translations is one half of that norm. Both path laws have the same endpoint distribution π, so the entropy chain rule gives

$$\mathrm{KL}(Q_\theta^\sigma \| Q_{\mathrm{ref}}^\sigma) = \int \mathrm{KL}(Q_\theta^{\sigma,xy} \| Q_{\mathrm{ref}}^{\sigma,xy})\, \pi(dx, dy) = \frac{A(\theta)}{2\sigma^2}. \tag{39}$$

This is an equality for the path laws themselves: Z can be recovered as the path endpoints because the added Brownian bridge vanishes at 0 and 1.

If $I_{\mathrm{ref}} = \ell$, then $Q_{\mathrm{ref}}^{\sigma,xy} = R^{xy}$. Disintegrating relative to R therefore gives

$$\mathrm{KL}(Q_\theta^\sigma \| R) = \mathrm{KL}(\pi \| R_{0,1}) + \int \mathrm{KL}(Q_\theta^{\sigma,xy} \| R^{xy})\, \pi(dx, dy), \tag{40}$$

□

The distinction between the two reference laws is material. A Monge OT coupling, or a discrete empirical coupling, can have infinite entropy relative to $R_{0,1}$. Adding pinned noise leaves that coupling unchanged, so it does not make the resulting law a finite-entropy competitor for the population Brownian problem. Equation (11) remains valid because both of its laws use exactly the same coupling. The deterministic interpolant itself has absolutely continuous trajectories and is singular relative to a nondegenerate Brownian path law; the finite path-space KL comparison concerns its auxiliary noisy lift.

H.2 Moment constraints and variational approximation

Centre constraints versus noisy constraints. For a feature block at time τ_k, the observation under the noisy lift is

$$\mathbb{E}_{Q_\theta^\sigma} f_k(X_{\tau_k}) = \mathbb{E}_{\pi,\xi} f_k\Big(I_\theta(\tau_k, X_0, X_1) + \sigma\sqrt{\tau_k(1-\tau_k)}\,\xi\Big). \tag{41}$$

This equals the implemented observation for affine f_k. For nonlinear features it instead evaluates a Gaussian-smoothed feature at the centre. If f_k is L_k-Lipschitz, the discrepancy is at most

$$\big\|\mathbb{E}_{Q_\theta^\sigma} f_k(X_{\tau_k}) - \mathbb{E}_\pi f_k(I_\theta(\tau_k, X_0, X_1))\big\| \le L_k \sigma \sqrt{d\,\tau_k(1-\tau_k)}. \tag{42}$$

For second moments, the difference is exactly $\sigma^2 \tau_k(1-\tau_k) I_d$; the same addition holds for covariance because the noise is independent and centred. These relations describe the constraint discrepancy, not convergence to the optimal bridge as σ decreases.

An exact projection statement for the noisy formulation. Suppose that the lifted law has finite entropy relative to R. The identity in Eq. (27) gives

$$\mathrm{KL}(Q_\theta^\sigma \| P^\star) = \mathrm{KL}(Q_\theta^\sigma \| R) - \mathrm{KL}(P^\star \| R). \tag{43}$$

At a fixed coupling, minimising the Euclidean prior with $\beta = 0$ is therefore equivalent to minimising $\mathrm{KL}(Q_\theta^\sigma \| P^\star)$ over its feasible lifted family. This family need not contain $P^\star$, or even contain a feasible member for every observation. It preserves conditional Brownian covariance, whereas a general exponential tilt does not.

If a feasible lifted law has excess entropy at most ε, the same identity and data processing imply

$$\mathrm{KL}(Q_\theta^\sigma \| P^\star) \le \varepsilon, \qquad \|Q_{\theta,t}^\sigma - P_t^\star\|_{\mathrm{TV}} \le \sqrt{\varepsilon/2}. \tag{44}$$

Here ε includes the best approximation error within the family and any remaining optimisation error, measured in entropy units. Feasibility alone does not bound it. Coupling error is also visible through

$$\mathrm{KL}(Q_\theta^\sigma \| P^\star) = \mathrm{KL}(\pi \| \pi^\star) + \int \mathrm{KL}(Q_\theta^{\sigma,xy} \| P^{\star,xy})\, \pi(dx, dy), \qquad \pi^\star = P^\star_{0,1}. \tag{45}$$

Using $\pi = \pi^\star$ removes the first term but not the conditional-law restriction.

For a candidate Q with the correct endpoints but residuals $r_k = \mathbb{E}_Q f_k(X_{\tau_k}) - a_k$, the corresponding identity is

$$\mathrm{KL}(Q \| P^\star) = \mathrm{KL}(Q \| R) - \mathrm{KL}(P^\star \| R) - \sum_{k=1}^{K} \lambda_k^\top r_k. \tag{46}$$

Thus one cannot apply the feasible projection statement to the implemented centre constraints without accounting for their noisy residuals. Similarly, the acceleration and metric energies introduce additional prior choices beyond the Brownian-relative-entropy objective.

Role of the augmented Lagrangian. At a feasible stationary point of Eq. (21),

$$c_k(\theta) = 0, \qquad \nabla_\theta R_{\mathrm{prior}}(\theta) + \sum_{k=1}^{K} \eta_k Dc_k(\theta)^\top \nu_k = 0. \tag{47}$$

The quadratic penalty has zero gradient at feasibility. These are the first-order conditions for the constrained interpolant problem, so the variational interpretation concerns the ideal constrained optimum targeted by Stage A. It does not establish that stochastic, clipped multiplier updates reach a global feasible optimum in the neural parametrisation.

H.3 Conditional means and the exact mean-constraint case

Proposition 4 (Ideal recovery for a mean constraint)**.** *Consider one constraint with linear f, $\mathbb{E}[f(X_\tau)] = a$ and endpoint laws with finite second moments for which the Brownian bridge problem has finite value. Let P^* be the solution to the optimisation problem Eq.* (10)*. With $\pi = P^\star_{0,1}$, $I_{\mathrm{ref}} = \ell$, $\alpha > 0$, minimising Eq.* (6) *over all endpoint-preserving finite-energy interpolants yields $I^\star(t,x,y) = \mathbb{E}_{P^\star}[X_t \mid X_0 = x, X_1 = y]$. Its noisy lift equals $P^\star$, and*

$$W_2(\mathrm{Law}(I^\star(t, X_0, X_1)), P_t^\star) \le \sigma \sqrt{d\, t(1-t)}. \tag{48}$$

Conditional mean for a general single-time tilt. Define

$$b_\tau(t) = \begin{cases} t/\tau, & 0 \le t \le \tau, \\ (1-t)/(1-\tau), & \tau \le t \le 1, \end{cases} \qquad \bar{z}_\lambda(x,y) = \int z\, q_\lambda(dz \mid x, y). \tag{49}$$

Whenever these first moments exist, conditional expectation of the two subbridge means yields

$$\mathbb{E}_{P^\star}[X_t \mid X_0 = x, X_1 = y] = \ell_t(x,y) + b_\tau(t)\big(\bar{z}_\lambda(x,y) - \ell_\tau(x,y)\big). \tag{50}$$

The resulting function has finite endpoint limits. With finite endpoint and intermediate second moments, smoothing the tent at τ approximates the correction in expected finite-energy norm. This is a representational connection. For nonlinear f, evaluating f at this conditional mean generally does not preserve the constrained expectation of $f(X_\tau)$.

Proof of Proposition 4. Let

$$\Delta = a - \big((1-\tau)\mathbb{E}_{p_0} X + \tau \mathbb{E}_{p_1} X\big), \qquad c = \sigma^2 \tau(1-\tau). \tag{51}$$

For $f(z) = z$, Gaussian tilting gives

$$q_\lambda(\cdot \mid x, y) = \mathcal{N}(\ell_\tau(x,y) + c\lambda, cI_d), \qquad Z_\lambda(x,y) = \exp\Big(\lambda^\top \ell_\tau(x,y) + \frac{c}{2}\|\lambda\|^2\Big). \tag{52}$$

The logarithm of Z_λ is the sum of a function of x, a function of y, and a constant. Its integral is therefore constant over $\Pi(p_0, p_1)$, so Eq. (28) has the same minimiser as the endpoint-only entropic problem: $\pi^\star = \pi^{\mathrm{SB}}$. The mean constraint sets $\lambda = \Delta/c$.

The linear tilt shifts the conditional mean of the entire Gaussian bridge by $b_\tau(t)\Delta$ and leaves its covariance unchanged. Consequently,

$$I^\star(t,x,y) = \ell_t(x,y) + b_\tau(t)\Delta, \qquad P^{\star,xy} = \mathrm{Law}(I^\star(\cdot,x,y) + \sigma B). \tag{53}$$

To verify that this is also the deterministic Stage-A optimum, let $h = I - \ell$ be any feasible finite-energy correction and put $\bar{h}_t = \mathbb{E}_\pi h_t$. It satisfies $\bar{h}_0 = \bar{h}_1 = 0$ and $\bar{h}_\tau = \Delta$. Jensen's and Cauchy–Schwarz inequalities give

$$\mathbb{E}_\pi \int_0^1 \|\dot{h}_t\|^2 dt \geq \int_0^1 \|\dot{\bar{h}}_t\|^2 dt \geq \frac{\|\Delta\|^2}{\tau} + \frac{\|\Delta\|^2}{1-\tau}. \tag{54}$$

Equality is attained by the endpoint-independent correction $h_t = b_\tau(t)\Delta$. This proves the optimum for $\alpha > 0$ and $\beta = 0$. With the correct coupling its noisy lift is $P^\star$ by Eq. (53). Finally, coupling the centre and noisy path through the same endpoints and Brownian bridge gives

$$W_2^2(\mathrm{Law}(I^\star(t,X_0,X_1)), P_t^\star) \leq \mathbb{E}\|\sigma B_t\|^2 = d\sigma^2 t(1-t). \tag{55}$$

□

The corner in b_τ is admissible for the finite-energy problem. A smooth neural parametrisation approximates it; a positive acceleration penalty selects a different, smoother optimum in general. Even in this exact mean case the two multiplier conventions differ: for one AL block, stationarity gives $\eta\nu = -2\alpha\sigma^2\lambda$. No generic identification of the neural optimisation multiplier with the entropy-dual multiplier is implied.

Transfer through Stage B. Let p_t^θ be the deterministic interpolant marginals and $u_t^\theta = \mathbb{E}[\dot{I}_\theta \mid I_\theta = \cdot]$. If I_θ has finite expected kinetic energy, conditional-expectation orthogonality yields

$$\mathcal{L}_B(v;\theta) = \mathcal{L}_B(u^\theta;\theta) + \int_0^1 \|v_t - u_t^\theta\|^2_{L^2(p_t^\theta)} dt. \tag{56}$$

This identifies the teacher's marginal velocity. It does not change the teacher path. To make the resulting error separation explicit, assume that u^θ and v generate unique flows, the u^θ-flow has marginals p_t^θ, and v_t is $L(t)$-Lipschitz with $L \in L^1[0,1]$. Coupling both flows by their initial state and applying Gronwall's inequality gives

$$W_2(\widehat{p}_t, P_t^\star) \leq W_2(p_t^\theta, P_t^\star) + e^{\int_0^t L(s)ds}\sqrt{t\,\varepsilon_B(t)} + \delta_{\mathrm{num}}(t), \tag{57}$$

where $\varepsilon_B(t) = \int_0^t \|v_s - u_s^\theta\|^2_{L^2(p_s^\theta)} ds$, and the numerical rollout has RMS pathwise error at most $\delta_{\mathrm{num}}(t)$ relative to the exact v-flow. Specifically, the synchronous RMS separation is bounded by $\int_0^t \exp(\int_s^t L(r)dr)\|v_s - u_s^\theta\|_{L^2(p_s^\theta)} ds$; Cauchy–Schwarz and the triangle inequality give Eq. (57). Proposition 4 bounds the first term in its ideal mean-constraint setting. In general, that term remains a separate error due to path selection.

H.4 Proposition for the Covariance

Proposition 5 (Ideal comparison for mean and covariance constraints)**.** *Let p_0, p_1 be endpoint laws on $\mathbb{R}^d$ with finite second moments. Fix $\sigma > 0$, and let R be the law of*

$$dX_t = \sigma\, dW_t, \qquad X_0 \sim p_0, \qquad t \in [0,1],$$

where W is a standard d-dimensional Brownian motion independent of X_0.
Fix $\tau \in (0,1)$, $\mu \in \mathbb{R}^d$, and $\Sigma \succ 0$. Let $P^\star$ be a finite-relative-entropy solution of Eq. (10)*, with endpoint laws of finite second moments and constraints*

$$\mathbb{E}_{P^\star}[X_\tau] = \mu, \qquad \mathrm{Cov}_{P^\star}(X_\tau) = \Sigma.$$

Set $\pi = P^\star_{0,1}$, $m(x,y) = \ell_\tau(x,y)$, and $c = \sigma^2\tau(1-\tau)$. Assume

$$M := \mathrm{Cov}_\pi(m(X_0,X_1)) \succ 0.$$

Define

$$b_\tau(t) = \begin{cases} t/\tau, & 0 \leq t \leq \tau, \\ (1-t)/(1-\tau), & \tau \leq t \leq 1, \end{cases} \qquad \kappa = \max\left\{1, \sqrt{\frac{\lambda_{\max}(\Sigma)}{\lambda_{\min}(M)}}\right\}.$$

With $I_{\mathrm{ref}} = \ell$, $\alpha > 0$, and $\beta = 0$, minimising Eq. (6) under these mean and covariance constraints over all endpoint-preserving finite-energy interpolants yields a unique interpolant $I^\star_{\mathrm{A}}$, up to π-almost-sure equality.
Writing

$$\bar{I}^\star(t,x,y) = \mathbb{E}_{P^\star}[X_t \mid X_0 = x, X_1 = y], \qquad p_t^{\mathrm{A}} = \mathrm{Law}(I^\star_{\mathrm{A}}(t, X_0, X_1)),$$

where $(X_0, X_1) \sim \pi$, we have, for every $t \in [0,1]$,

$$\left\| I^\star_{\mathrm{A}}(t,\cdot) - \bar{I}^\star(t,\cdot) \right\|_{L^2(\pi)} \le \frac{c\, b_\tau(t)}{2} \sqrt{\mathrm{tr}(M^{-1})}. \tag{58}$$

Moreover,

$$W_2(p_t^{\mathrm{A}}, P_t^\star) \le \left[\frac{c^2 b_\tau(t)^2}{4} \mathrm{tr}(M^{-1}) + \kappa d \sigma^2 t(1-t) \right]^{1/2}. \tag{59}$$

Proof. Write $\bar{m} = \mathbb{E}_\pi m$ and define

$$C = M^{1/2} \Sigma M^{1/2}, \qquad K = M^{-1/2} C^{1/2} M^{-1/2},$$
$$T = M^{-1/2} \left[\left(C + \frac{c^2}{4} I_d \right)^{1/2} - \frac{c}{2} I_d \right] M^{-1/2}.$$

All matrix square roots are symmetric positive definite. The matrices $K, T \succ 0$ satisfy

$$KMK = \Sigma, \qquad TMT + cT = \Sigma. \tag{60}$$

We first identify the optimal deterministic interpolant. For an admissible I, put $z(x,y) = I(\tau, x, y)$. Cauchy–Schwarz on $[0,\tau]$ and $[\tau,1]$ gives

$$\mathbb{E}_\pi \int_0^1 \|\partial_t (I_t - \ell_t)\|^2 \, dt \ge \frac{\mathbb{E}_\pi \|z - m\|^2}{\tau(1-\tau)}.$$

Equality holds precisely for

$$I_t = \ell_t + b_\tau(t)(z - m).$$

Since $\alpha > 0$ and $\beta = 0$, the interpolant problem therefore reduces to minimising $\mathbb{E}_\pi \|z - m\|^2$ subject to $\mathbb{E}_\pi z = \mu$ and $\mathrm{Cov}_\pi(z) = \Sigma$.

Set $u = z - \mu$ and $v = m - \bar{m}$. Every feasible z satisfies

$$\mathbb{E}_\pi \|z - m\|^2 = \|\mu - \bar{m}\|^2 + \mathrm{tr}(\Sigma) + \mathrm{tr}(M) - 2\mathbb{E}_\pi[u^\top v].$$

Using $KMK = \Sigma$, we obtain

$$0 \le \mathbb{E}_\pi[(u - Kv)^\top K^{-1} (u - Kv)] = 2\,\mathrm{tr}(KM) - 2\mathbb{E}_\pi[u^\top v].$$

Equality holds if and only if $u = Kv$ almost surely. Consequently,

$$\begin{aligned} z_{\mathrm{A}}(x,y) &= \mu + K(m(x,y) - \bar{m}), \\ I^\star_{\mathrm{A}}(t,x,y) &= \ell_t(x,y) + b_\tau(t)(z_{\mathrm{A}}(x,y) - m(x,y)). \end{aligned} \tag{61}$$

We next identify the conditional midpoint law of $P^\star$. Consider

$$q^{xy} = \mathcal{N}(z_{\mathrm{SB}}(x,y), cT), \qquad z_{\mathrm{SB}}(x,y) = \mu + T(m(x,y) - \bar{m}).$$

Under the endpoint coupling π, this kernel has aggregate mean μ and covariance $TMT + cT = \Sigma$.

Let $G^{xy} = \mathcal{N}(m(x,y), cI_d)$ be the reference conditional midpoint law. Completing the square gives

$$\log \frac{dq^{xy}}{dG^{xy}}(z) = h^\top z + z^\top H z - \chi(x,y),$$

where

$$h = c^{-1}(T^{-1}\mu - \bar{m}), \qquad H = \frac{I_d - T^{-1}}{2c}.$$

The prescribed mean and covariance fix the aggregate expectations of z and $zz^\top$. Hence, for any competing conditional midpoint kernel $\widetilde{q}^{xy}$ with these aggregate moments and finite conditional relative entropy,

$$\begin{aligned}\mathbb{E}_\pi \,\mathrm{KL}(\widetilde{q}^{xy}\|G^{xy}) - \mathbb{E}_\pi \,\mathrm{KL}(q^{xy}\|G^{xy}) \\ = \mathbb{E}_\pi \,\mathrm{KL}(\widetilde{q}^{xy}\|q^{xy}) \geq 0.\end{aligned}$$

Entropy disintegration retains the reference Brownian bridges conditional on (X_0, X_τ, X_1). Since $P^\star$ also minimises entropy when its endpoint coupling π is held fixed, its conditional midpoint kernel is q^{xy}. The conditional means of the two subbridges give

$$\bar{I}^\star(t,x,y) = \ell_t(x,y) + b_\tau(t)(z_{\mathrm{SB}}(x,y) - m(x,y)).$$

To compare the conditional means, set

$$\begin{aligned}D &= M^{1/2}(K-T)M^{1/2} \\ &= C^{1/2} - \left(C + \frac{c^2}{4}I_d\right)^{1/2} + \frac{c}{2}I_d.\end{aligned}$$

Applying

$$0 \leq \sqrt{s} - \sqrt{s + c^2/4} + c/2 \leq c/2$$

to each eigenvalue $s > 0$ of C yields $0 \preceq D \preceq (c/2)I_d$. It follows that

$$\begin{aligned}\mathbb{E}_\pi\|(K-T)(m-\bar{m})\|^2 &= \mathrm{tr}(M^{-1}D^2) \\ &\leq \frac{c^2}{4}\,\mathrm{tr}(M^{-1}).\end{aligned}$$

Multiplying by $b_\tau(t)^2$ proves Eq. (58).

Finally, conditioning the reference subbridges on the midpoint gives

$$\mathrm{Cov}_{P^\star}(X_t \mid X_0, X_1) = v_t I_d + b_\tau(t)^2 cT,$$

where

$$v_t = \sigma^2 t(1-t) - b_\tau(t)^2 c \geq 0.$$

Couple $I^\star_{\mathrm{A}}(t, X_0, X_1)$ and X_t using the same endpoints. The conditional bias–variance identity then gives

$$W_2^2(p_t^{\mathrm{A}}, P_t^\star) \leq \frac{c^2 b_\tau(t)^2}{4}\,\mathrm{tr}(M^{-1}) + dv_t + b_\tau(t)^2 c\,\mathrm{tr}(T).$$

Since $TMT \preceq \Sigma$,

$$\lambda_{\min}(M)T^2 \preceq \Sigma, \qquad \lambda_{\max}(T) \leq \sqrt{\frac{\lambda_{\max}(\Sigma)}{\lambda_{\min}(M)}} \leq \kappa.$$

Consequently,

$$\begin{aligned}dv_t + b_\tau(t)^2 c\,\mathrm{tr}(T) &\leq d\kappa\big(v_t + b_\tau(t)^2 c\big) \\ &= d\kappa\sigma^2 t(1-t),\end{aligned}$$

which proves Eq. (59).

At $t = \tau$, the same argument yields the sharper bound

$$W_2^2(p_\tau^{\mathrm{A}}, P_\tau^\star) \leq \frac{c^2}{4}\,\mathrm{tr}(M^{-1}) + c\,\mathrm{tr}(T).$$

□