
Generalized Flow Matching for Transition Dynamics Modeling

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Simulating transition dynamics between metastable states is a fundamental chal-
2 lenge in dynamical systems and stochastic processes with wide real-world appli-
3 cations in understanding protein folding, chemical reactions and neural activities.
4 However, the computational challenge often lies on sampling exponentially many
5 paths in which only a small fraction ends in the target metastable state due to
6 existence of high energy barriers. To amortize the cost, we propose a data-driven
7 approach to warm-up the simulation by learning nonlinear interpolations from local
8 dynamics. Specifically, we infer a potential energy function from local dynamics
9 data. To find plausible paths between two metastable states, we formulate a gener-
10 alized flow matching framework that learns a vector field to sample propable paths
11 between the two marginal densities under the learned energy function. Furthermore,
12 we iteratively refine the model by assigning importance weights to the sampled
13 paths and buffering more likely paths for training. We validate the effectiveness of
14 the proposed method to sample probable paths on both synthetic and real-world
15 molecular systems.

16 1 Introduction

17 Transition dynamics simulation aims to sample transition paths between two metastable states, which
18 is a fundamental challenge in dynamical systems, stochastic processes, and molecular simulations,
19 with broad applications [5, 30, 2, 21, 9, 10]. The key computational obstacle lies in the rarity of
20 transition events such that it requires long-run molecular dynamics (MD) simulations to go over high
21 energy barriers. In addition, there exists an infinite amount of paths that do not end in the target state.

22 To address these limitations, early work leverage Markov chain Monte Carlo (MCMC) approaches to
23 mix the path distributions [7]. However, MCMC-based simulation in such high-dimensional spaces
24 suffer from slow mixing time. Alternatively, later work formulate it as a path integral control problem
25 such that an external control is learned to guide the stochastic process for certain terminal conditions
26 (i.e. arrive at the target in finite time) [12, 31, 6, 24]. Nevertheless, the path integral control method
27 is known as a shooting method with high variance (i.e. the probability to hit the target is very low).
28 Recently, a variational formulation of learning Doob’s h-transform is proposed, which leads to a
29 collocation method that optimizes a family of tractable probability paths [8].

30 Nevertheless, all previous work require an extensive amount of expensive energy evaluations to sample
31 from the path distributions. In this paper, we study the possibility to find a low-cost approximation of
32 the transition paths distribution. In many studies of rare events in molecular systems, local short-run
33 molecular simulations are used to provide information about transitions [28]. Our approach builds
34 on the idea that even with very short-run local molecular simulations, we can initialize transition
35 paths more effectively than common methods like linear interpolation or heuristic-based method [26].

36 Specifically, we aim to learn a “potential energy” from these short-run trajectories which indicates
 37 the likelihood of the states despite they are by no means well-equilibrated.

38 Building on the recent progress of Schrödinger bridge and flow matching models [15, 19, 13], we
 39 formulate the problem as a generalized flow matching problem which introduces an additional
 40 potential energy as constraint than the normal flow matching setup [14]. We solve the generalized
 41 flow matching problem by learning a vector field to moving from the distribution of one metastable
 42 state to another that minimizes the overall transportation cost. To further improve the quality of the
 43 learned path, we apply importance sampling to resample transition paths reweighted by the path
 44 probability induced by the true energy function.

45 We validate the effectiveness of the proposed method on both synthetic data and real-world molecular
 46 systems. The results demonstrate that our method can sample high-quality transition paths (close to
 47 saddle points) with a significantly less energy evaluations to generate local dynamics data.

48 2 Background

49 2.1 Score & Flow Matching

50 Score-based generative models or diffusion models [27, 11] build a generative process from tractable
 51 prior distribution $p_T(x)$ (e.g. Gaussian) to complex data distribution $p_0(x) \approx p_{\text{data}}(x)$ where $x \in \mathbb{R}^d$.
 52 It can be learned by reversing a forward stochastic process from $p_0(x)$ to $p_T(x)$, as follows:

$$dx_t = f(x_t, t)dt + g(t)dW_t \quad (1)$$

$$dx_t = [f(x_t, t) - g^2(t)\nabla_x \log p_t(x_t)] dt + g(t)d\bar{W}_t \quad (2)$$

53 where $f : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ is the drift, $g : \mathbb{R} \rightarrow \mathbb{R}$ is the scalar diffusion coefficient, W_t, \bar{W}_t are
 54 d -dimensional Wiener process and the score network $s_\theta(x_t, t)$ is learned to match the conditional
 55 score function $\nabla_x \log p_t(x_t|x_0)$:

$$\min_s \mathbb{E}_{\mathcal{U}(t)} \mathbb{E}_{p_0} \mathbb{E}_{p_{t|0}} [\|s_\theta(x_t, t) - \nabla_x \log p_t(x_t|x_0)\|^2] \quad (3)$$

56 Flow matching [14, 1, 17] generalizes the idea of score matching by extending it to a general family of
 57 Gaussian conditional probability paths $p_t(x_t|x_0) = \mathcal{N}(x_t|\mu_t(x_0), \sigma_t(x_0)^2 I)$ and regress the vector
 58 field $v_\theta(x_t, t)$ to the conditional vector field $u_t(x_t|x_0) = \frac{\sigma_t'(x_0)}{\sigma_t(x_0)}(x_t - \mu_t(x_0)) + \mu_t'(x_0)$ corresponding
 59 to the Gaussian path $p_t(x_t|x_0)$:

$$\min_v \mathbb{E}_{\mathcal{U}(t)} \mathbb{E}_{p_0} \mathbb{E}_{p_{t|0}} [\|v_\theta(x_t, t) - u_t(x_t|x_0)\|^2] \quad (4)$$

60 2.2 Generalized Schrödinger Bridge Matching

61 One key ingredient behind the success of score-based generative and flow matching models is the
 62 simulation-free forward process (diffusion paths or Gaussian paths) that can be evaluated analytically.
 63 However, there is a broader class of problems with nonlinear drift or constraint that require simulation
 64 of the forward process. The generalized Schrödinger bridge problem is more general such that in
 65 addition to learn a stochastic process with minimal kinetic energy that connects $p_0(x)$ and $p_T(x)$, it
 66 further minimizes the potential energy $V(\cdot)$ along the path as [15]:

$$\min_{f, p_t} \int_0^T \int \left(\frac{1}{2} \|f_\theta(x_t)\|^2 + V(x_t) \right) p_t(x_t) dx dt, \quad (5)$$

$$\text{s.t. } \partial_t p_t(x_t) = -\nabla \cdot (p_t(x_t) f_\theta(x_t, t)) + \frac{1}{2} g(t) \Delta p_t(x_t), p_0 = \mu_0, p_T = \mu_T \quad (6)$$

67 To solve the problem in Eq. (5), [15] propose to approximate the true marginal distribution p_t of the
 68 stochastic process with Gaussian density conditioned on both end points $p_{t|0,T}(x_t) = \mathcal{N}(x_t|\mu_t, \sigma_t^2 I)$
 69 where μ_t and σ_t can be parameterized by neural networks or spline. Similar to flow matching,
 70 approximating the marginals using Gaussian paths allows for simulation-free sampling and thus easy
 71 optimization within the the objective Eq. (5). In the end, f_θ is learned by matching against the vector
 72 field corresponding to the Gaussian path similar to Eq. (4). It is worth noting that a similar general
 73 framework has also been proposed in [18, 19].

74 3 Generalized Flow Matching for Transition Dynamics Modeling

75 3.1 Short-Run Molecular Dynamics

76 Molecular dynamics simulation follows Newton’s equation of motion such that $M\ddot{x}_t = -\nabla_x U(x_t)$,
 77 where M is the mass of the particles, $U : \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}$ is the potential energy function and $x =$
 78 $(x_1, \dots, x_N) \in \mathbb{R}^{N \times 3}$ is one state of the molecular system. Nevertheless, one common goal of
 79 interest to run molecular simulation is to sample from the thermal equilibrium (known as the NVT
 80 ensemble) with a heat bath at a fixed temperature, which is taken into account by running the following
 81 Langevin equation.

$$\begin{pmatrix} dx_t \\ dv_t \end{pmatrix} = \begin{pmatrix} v_t \\ -M^{-1}\nabla_x U(x_t) - \gamma v_t \end{pmatrix} dt + \begin{pmatrix} 0 & 0 \\ 0 & M^{-1/2} \sqrt{2\gamma k_B T_p} \end{pmatrix} dW_t \quad (7)$$

82 where γ is the friction coefficient, k_B is the Boltzmann constant and temperature T_p . The stationary
 83 distribution of running this Markov process is the thermal equilibrium $p(x) \propto e^{-U(x)/k_B T_p}$. However,
 84 with the goal of learning transition paths with a limited number of potential energy evaluations, we
 85 do not attempt to sample from the equilibrium.

86 Instead, we consider two predefined metastable states A and B , which correspond to local minima
 87 on the potential energy landscape. We run only local molecular dynamics simulations around two
 88 metastable state to initialize our subsequent methods. In particular, we let $\Phi_{s:t}^{\rightarrow}(y)$ denote simulating
 89 the dynamics in Eq. (7) initialized at point y and time s for time $t - s$ and $t > s$. For $t < s$, we
 90 integrate backward in time and write $\Phi_{s:t}^{\leftarrow}(y)$.

91 For a short time interval s chosen as a design decision, our subsequent methods are initialized from
 92 the induced distributions $\mu_0(x) = (\Phi_{-s:0}^{\rightarrow})_{\#} \delta_A$ and $\mu_T(x) = (\Phi_{T+s:T}^{\leftarrow})_{\#} \delta_B$ where we abbreviate
 93 $\delta_A(y) = \delta(A - y)$, $\delta_B(y) = \delta(B - y)$.

94 We illustrate examples of μ_0 and μ_T resulting from different choices of s in Fig. 1, where A and B
 95 are indicated using green points and μ_0 and μ_T are indicated using red and blue points, respectively.
 96 Our hope is that short-run MD simulations provide a useful initialization for our Generalized Flow
 97 Matching transport problem with learned potentials.

98 3.2 Generalized Flow Matching

99 Inspired by the generalized Schrödinger bridge problem, we formulate the problem of finding feasible
 100 transition paths as a distribution matching problem such that we learn a vector field to transport
 101 between two given marginal distributions:

$$\begin{aligned} \mathcal{L}_{\text{GFM}} &= \min_{v_t^{\theta}, p_t} \int_0^T \int \left(\frac{1}{2} \|v_t^{\theta}(x_t)\|^2 + V_t(x_t) \right) p_t(x_t) dx dt \\ \text{s.t. } \partial_t p_t(x_t) &= -\nabla \cdot (p_t(x_t) v_t^{\theta}(x_t)), p_0 = \mu_0, p_T = \mu_T \end{aligned} \quad (8)$$

102 where μ_0 and μ_T are the density of the states explored by a local molecular dynamics simulation
 103 around the metastable states A and B , described above. In general, the kinetic energy corresponds to
 104 the speed of transport while the potential energy defines a state cost corresponding to how probable
 105 the states are (e.g. likelihood). Together these cost functions define an optimal interpolation of
 106 marginals p_t^* between μ_0 and μ_T , which reduces to the dynamic formulation of optimal transport
 107 for $V_t(\cdot) = 0$ [22]. It is worth noting that the potential energy function $V(\cdot)$ does not have to be the
 108 same as $U(\cdot)$ used in molecular dynamics simulation, and we now discuss how to learn a surrogate
 109 potential energy functions from μ_0, μ_T to guide the sampling of transition paths.

110 3.3 Inferring Kinetic and Potential Energy from Data

111 As one of the main goal of this study is to reduce the number of potential energy evaluation $U(\cdot)$, we
 112 show how we can learn a surrogate energy function $V(\cdot)$ from the local dynamics data in two ways.

113 *Latent interpolation:* We propose to learn an autoencoder that maps high-dimensional data into
 114 low-dimensional representations such that it preserves structural information [15]. The hypothesis is
 115 that the latent space compresses semantic information from data thus better measure distance than
 116 the ambient Euclidean space. Specifically, we map data (x_0, x_T) to the latent space as (z_0, z_T) . We

117 define the potential energy as the deviation of the state x_t from certain interpolation (e.g. spherical
118 interpolation) of x_0 and x_T in the latent space $I(z_0, z_T, t)$:

$$V(x_t) = \|x_t - \text{Decoder}(I(z_0, z_T, t))\|_2^2 \quad (9)$$

119 *Metric learning*: Another natural way to learn the potential energy is through metric learning such
120 that the metric informs how dense the data is around a particular location [4, 3]. Once the metric
121 $G : \mathbb{R}^{N \times 3} \rightarrow \mathbb{R}^{N \times N}$ is learned, the kinetic energy term in Eq. (8) will become:

$$\|v_t\|_G = v_t^T G(x_t) v_t = v_t^T v_t + v_t^T (G(x_t) - I) v_t = \|v_t\|_2 + V(x_t, v_t) \quad (10)$$

122 where the potential energy is implied as:

$$V(x_t, v_t) = v_t^T (G(x_t) - I) v_t \quad (11)$$

123 It is worth noting this can be equivalently considered as defining the kinetic energy under a learned
124 metric tensor G because potential energy often only depends on x_t , where in the other case, the
125 kinetic energy is defined under the trivial diagonal metric in Euclidean space. We leave the details
126 about metric learning in App. A.

127 3.4 Conditional Generalized Flow Matching Objective

128 To facilitate optimization of Eq. (8), we derive the following conditional or ‘bridge’ objective.

129 **Definition 1.** Assume the path of marginals p_t decomposes such that $p_t(x_t) =$
130 $\int p_{0,T}(x_0, x_T) p_{t|0,T}(x_t) dx_0 dx_T$ and there exists $v_{t|0,T}$ such that $\partial_t p_{t|0,T} = -\nabla \cdot (p_{t|0,T} v_{t|0,T})$.
131 We define the following *conditional* GFM objective

$$\mathcal{L}_{\text{cGFM}}(x_0, x_T) := \min_{p_{t|0,T}, v_{t|0,T}} \int_0^T \int \left(\frac{1}{2} \|v_{t|0,T}(x_t)\|^2 + V_t(x_t) \right) p_{t|0,T}(x_t | x_0, x_T) dx_t dt \quad (12)$$

s.t. $\partial_t p_{t|0,T}(x_t) = -\nabla \cdot (p_{t|0,T}(x_t) v_{t|0,T}(x_t))$, $p_{0|0,T} = \delta_{x_0}$, $p_{T|0,T} = \delta_{x_T}$.

132 where we abbreviate the boundary condition $p_{0|0,T}(x) = \delta(x - x_0)$ and optimize a conditional
133 objective for each sample from the joint distribution $(x_0, x_T) \sim p_{0,T}$.

134 The objective in Def. 1 is a deterministic analogue the conditional stochastic control objective in Prop.
135 2 of [15]. We next show that the optimizing the conditional objective yields an *upper bound* on the
136 marginal objective, which was not emphasized by [15]. See App. C.2 for proof.

137 **Proposition 1.** Taking the expectation of the conditional objective over $(x_0, x_T) \sim p_{0,T}$ and
138 enforcing that $p_{0,T} \in \Pi(\mu_0, \mu_T)$ satisfies the boundary conditions yields an upper bound on \mathcal{L}_{GFM} ,

$$\mathcal{L}_{\text{GFM}} \leq \mathbb{E}_{p_{0,T}(x_0, x_T)} [\mathcal{L}_{\text{cGFM}}(x_0, x_T)] \quad \text{s.t.} \quad p_0 = \mu_0, p_T = \mu_T. \quad (13)$$

139 For a learned potential energy (Sec. 3.3), our computational approach seeks to find an approximate
140 solution to the marginal GFM problem (Eq. (8)) by representing a coupling $p_{0,T}$ and parameterizing
141 $p_{t|0,T}$ and $v_{t|0,T}^\phi$. However, to generate unconditional transition paths from an initial x_0 or x_T and
142 facilitate a parameterization of $p_{0,T}$, we also consider learning a marginal vector field v_t^θ in Eq. (8).

143 **Neural Spline Parameterization of $v_{t|0,T}^\phi$** The conditional objective in Def. 1 requires $p_{t|0,T}$ and
144 $v_{t|0,T}^\phi$ which satisfy the continuity equation and respect the boundary conditions $\delta_{x_0}, \delta_{x_T}$. To satisfy
145 these desiderata, we use neural networks to parameterize a spline interpolation in the sample space
146 that satisfies the boundary conditions:

$$x_t^\phi = (1-t)x_0 + tx_T + t(1-t)\text{NN}_\phi(x_0, x_T, t) \quad (14)$$

147 This spline induces a path of marginal distributions $p_{t|0,T}$, where the velocity $v_{t|0,T}$ satisfying the
148 continuity equation in Eq. (12) is directly tractable as [1, 29]

$$v_{t|0,T}^\phi(x_t) = x_T - x_0 + t(1-t)\dot{\text{NN}}_\phi(x_0, x_T, t) + (1-2t)\text{NN}_\phi(x_0, x_T, t) \quad (15)$$

149 This choice of $v_{t|0,T}^\phi$ thus satisfies both the endpoint constraints in Eq. (14) and the conditional
150 continuity equation constraint for the marginals $p_{t|0,T}^\phi$ induced by x_t^ϕ sampling.

Algorithm 1 Generalized Flow Matching Algorithm

Require: data p_0, p_T , learned potential energy $V(\cdot)$, neural spline network NN_ϕ , velocity network v_θ (initialized as the solution to initial coupling, e.g. $\mu_0 \otimes \mu_T$), replay buffer \mathcal{B} , maximum replay buffer size $|\mathcal{B}|$

1: **while** not converged **do**

2: Sample $x_0 \sim p_0, x_T \sim \text{ODE}([0, T], x_0, v_\theta), t \sim \mathcal{U}(0, T)$

3: $x_t^\phi = (1-t)x_0 + tx_T + t(1-t)\text{NN}_\phi(x_0, x_T, t)$ eq. (14)

4: $v_{t|0,T}^\phi(x_t) = x_T - x_0 + t(1-t)\text{NN}_\phi(x_0, x_T, t) + (1-2t)\text{NN}_\phi(x_0, x_T, t)$ eq. (15)

5: $\mathcal{L}_{\text{spline}}^{p_0, T}(\phi) = \frac{1}{2}\|v_{t|0,T}^\phi\|^2 + V(x_t)$ eq. (17)

6: $\mathcal{L}_{\text{flow}}(\theta) = \|v_t^\theta(x_t) - \text{sg}(v_{t|0,T}^\phi(x_0, x_T, t))\|^2$ eq. (19)

7: Update ϕ, θ using $\nabla_\phi \mathcal{L}_{\text{spline}}(\phi), \nabla_\theta \mathcal{L}_{\text{flow}}(\theta)$

8: **if** Replay Buffer **then**

9: $\{x_0^i\}_{i=0}^N \sim p_0$

10: $\{\gamma^i\}_{i=0}^N = \text{ODE}([0, T], \{x_0^i\}_{i=0}^N, v_\theta), \gamma = x_{0:T}$

11: Compute weight \tilde{w}^i for γ^i , add (γ^i, \tilde{w}^i) to \mathcal{B} eq. (20)

12: **while** Replay Buffer not converged **do**

13: Sample $\gamma \sim \mathcal{B}, t \sim \mathcal{U}(0, T), x_0 \sim p_0, x_T \sim \text{ODE}([0, T], x_0, v_\theta)$

14: $\mathcal{L}_{\text{replay}}(\phi) = \|x_\phi(x_0, x_T, t) - \gamma_t\|^2 + \|v_\phi(x_0, x_T, t)\|^2$ eq. (21)

15: $\mathcal{L}_{\text{flow}}(\theta) = \|v_t^\theta(x_t) - \text{sg}(v_{t|0,T}^\phi(x_0, x_T, t))\|^2$

16: Update ϕ, θ using $\nabla_\phi \mathcal{L}_{\text{replay}}(\phi), \nabla_\theta \mathcal{L}_{\text{flow}}(\theta)$

17: **end while**

18: **end if**

19: **end while**

20: **return** NN_ϕ, v_θ

151 **Definition 2.** Define $\mathcal{L}_{\text{cGFM}}^\phi(x_0, x_T)$ as the value of the objective in Eq. (12) for possibly suboptimal
 152 $x_t^\phi \sim p_{t|0,T}$ and $v_{t|0,T}^\phi$ which satisfy the constraints, as is guaranteed by our parameterization above.
 153 Using Eq. (12) and Eq. (13), we have

$$\mathcal{L}_{\text{cGFM}}(x_0, x_T) \leq \mathcal{L}_{\text{cGFM}}^\phi(x_0, x_T), \quad \mathcal{L}_{\text{GFM}} \leq \mathbb{E}_{p_{0,T}} \left[\mathcal{L}_{\text{cGFM}}^\phi(x_0, x_T) \right] \quad \forall p_{0,T} \in \Pi(\mu_0, \mu_T). \quad (16)$$

154 Using this parameterization, we can tractably optimize the objective $\mathcal{L}_{\text{cGFM}}^\phi(x_0, x_T)$ as a function
 155 of ϕ . Expanding to write an expectation over samples from $p_{0,T}$, we minimize the following the
 156 objective over $v_{t|0,T}^\phi$,

$$\begin{aligned} \mathcal{L}_{\text{spline}}^{p_0, T}(\phi) &= \mathbb{E}_{p_{0,T}} \left[\mathcal{L}_{\text{cGFM}}^\phi(x_0, x_T) \right] \\ &= \mathbb{E}_{p_{0,T}(x_0, x_T)} \left[\mathbb{E}_{\mathcal{U}(t)} \mathbb{E}_{p_{t|0,T}(x_t)} \left[\frac{1}{2} \|v_\phi(x_0, x_T, t)\|_2^2 + V(x_\phi(x_0, x_T, t)) \right] \right] \end{aligned} \quad (17)$$

157 where we have replaced the integral over $\int_0^T (\cdot) dt = \mathbb{E}_{\mathcal{U}(t)}[(\cdot)]$ with an expectation over the uniform
 158 distribution. In practice, we sample time points $t \sim \mathcal{U}(t)$ during training and use an empirical
 159 expectation over $(x_0, x_T) \sim p_{0,T}$.

160 **Marginal Vector Field** v_t^θ To perform unconditional sampling, we need to learn a marginal vector
 161 field v_t^θ which simulates the marginals p_t induced by a particular $p_{0,T}, p_{t|0,T}^\phi$, and $v_{t|0,T}^\phi$. This also
 162 corresponds to translating a candidate solution for the objective $\mathbb{E}_{p_{0,T}}[\mathcal{L}_{\text{cGFM}}^\phi(x_0, x_T)]$ in Eq. (13)
 163 into a solution to the marginal problem in Eq. (8). Following similar arguments as [14, 1, 29], one
 164 can show that the marginal vector field

$$v_t^{(p_{0,T}, \phi)}(x_t) = \mathbb{E}_{p_{0,T}(x_0, x_T)} \left[\frac{p_{t|0,T}^\phi(x_t)}{p_t(x_t)} v_{t|0,T}^\phi(x_t) \right] \quad (18)$$

165 satisfies the continuity equation for the induced marginals $p_t(x_t) = \mathbb{E}_{p_{0,T}}[p_{t|0,T}^\phi(x_t)]$, with $\partial p_t =$
 166 $-\nabla \cdot (p_t v_t^{(p_{0,T}, \phi)})$. We use the notation $v_t^{(p_{0,T}, \phi)}$ to indicate that the appropriate vector field is
 167 induced from our current $p_{0,T}$, $p_{t|0,T}^\phi$, and $v_{t|0,T}^\phi$. See App. C Lemma 1,

168 We finally can use the flow matching objective to learn an approximate $v_t^\theta \approx v_t^{(p_{0,T}, \phi)}$ for our current
 169 $p_{0,T}$ and ϕ [14, 1, 29], where $\text{sg}(\cdot)$ indicates stop-gradient,

$$\mathcal{L}_{\text{flow}}(\theta) = \mathbb{E}_{\mathcal{U}(t)} \mathbb{E}_{p_{0,T}} \mathbb{E}_{p_{t|0,T}^\phi} \left\| v_t^\theta(x_t, t) - \text{sg} \left(v_{t|0,T}^\phi(x_0, x_T, t) \right) \right\|^2 \quad (19)$$

170 Note that we amortize this optimization over training steps by alternating updates for v_t^θ , $p_{0,T}$, and ϕ ,
 171 so that v_t^θ may not exactly match $v_t^{(p_{0,T}, \phi)}$.

172 **Coupling Parameterization** Finally, the above conditional objectives require samples from a
 173 coupling distribution $p_{0,T}$, which should satisfy the boundary constraints $p_0 = \mu_0$, $p_T = \mu_T$ in
 174 Eq. (13). Note that Prop. 1 requires optimization over $p_{0,T}$ and our initial data from MD simulation is
 175 unpaired ($p_{0,T} = \mu_0 \otimes \mu_T$). Thus, optimizing the neural spline on fixed, independent coupling will
 176 not lead an optimal solution. We introduce two ideas for maintaining and updating couplings $p_{0,T}$.

177 *Minibatch Optimal Transport.* Following [23, 29], we can consider solving for (entropic) optimal
 178 transport couplings over an empirical batch of samples $(x_0, x_T) \sim \mu_0 \otimes \mu_T$. However, the choice
 179 of cost is crucial to defining the solutions. While the objective in Eq. (8) corresponds to dynamical
 180 OT with the squared error or Euclidean cost for $V(\cdot) = 0$, calculating the appropriate dynamical
 181 cost with nonzero potential energy is an optimization problem in its own right [23]. To simplify the
 182 algorithm, we use OT couplings with the Euclidean cost and introduce resampling based on a replay
 183 buffer below.

184 *Rectified Flow.* Following [17, 25], an alternative technique to solve for optimal transport couplings is
 185 to iteratively simulate the marginal vector field v_t^θ . For example, we could sample $p_{0,T}^\theta$ by sampling
 186 an initial $x_0 \sim \mu_0$ and simulating to obtain x_T . For the Euclidean cost or a further family of convex
 187 costs, iterative simulation and matching v_t^θ approaches the solution to the optimal transport problem
 188 [16]. However, again, since our GFM problem involves a dynamical cost, we simplify by using the
 189 rectified coupling induced from v_t^θ and introduce corrections via resampling.

190 3.5 Resampling and Replay Buffer

191 Noting the fact that our coupling parameterizations above do not reflect the desired dynamical cost
 192 where $V(\cdot) := U(\cdot)$, we introduce a resampling procedure to reweight paths induced by our coupling
 193 $p_{0,T}$ and spline parameterization of $p_{t|0,T}^\phi$ or $x_t^\phi(x_0, x_T)$.

$$\tilde{w}(x|x_0, x_1) = \int_0^T \left(\frac{1}{2} \|v_\theta(x_t)\|^2 + U(x_t) \right) dt \quad (20)$$

194 We aim to assign greater weights to the lower cost paths and smaller weights to the higher cost paths.

195 We normalize the cost as weight $w_i = \exp(-\tilde{w}(x^i)) / \sum_{n=1}^N \exp(-\tilde{w}(x^i))$.

196 After sampling lower cost transition paths, we push them into a replay buffer \mathcal{B} by their importance
 197 weights. In addition, we refine the learned neural spline by drawing samples from the replay buffer.
 198 To do so, we sample paths $\gamma = \{x_t\}_{t=0}^T$ and optimize the following objective to align the neural
 199 spline with the paths:

$$\mathcal{L}_{\text{replay}}(\phi) = \mathbb{E}_{\mathcal{U}(t)} \mathbb{E}_{p_{0,T}} \mathbb{E}_\gamma \|x_\phi(x_0, x_T, t) - \gamma_t\|^2 + \|v_\phi(x_0, x_T, t)\|^2 \quad (21)$$

200 In addition, we apply a kinetic energy loss function to ensure the smoothness of the learned spline.

201 4 Experiment


202 4.1 Experiment Set-up

203 **Müller-Brown Potential.** We first employ the Müller-Brown potential which is a commonly used
 204 mathematical model to study transition paths between metastable states. The energy landscape is

205 characterized by three local minima and two saddle points connecting them and can be written down
 206 analytically in App. D.1. To simulate this system, we run the first-order Langevin dynamics around
 207 each local minima.

$$x_{t+1} = x_t - \nabla_x V(x_t) \cdot dt + \sqrt{dt} \cdot \text{diag}(\xi) \cdot \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1) \quad (22)$$

208 where we apply Euler discretization of the continuous dynamics. In our experiment, we set $dt = 10^{-4}$
 209 and $\xi = 5$.



(a) Shorter-run dataset (b) Longer-run dataset (a) Sampled paths (shorter) (b) Sampled paths (longer)

Figure 1: Both datasets in the figure contain 2000 pairs of data points, but randomly sampled from simulation of 4K and 12K steps, respectively. Figure 2: Sampled paths from models trained on both the shorter-run and longer-run datasets (Saddle points are stared).

210 **Alanine Dipeptide.** We validate our proposed method on a real-world molecular system, Alanine
 211 Dipeptide, which contains 22 atoms. The transition between two metastable states (C7eq and
 212 C7ax) is characterized by a two-dimensional free energy surface (ϕ, ψ dihedral angles). The
 213 molecular configurations are sampled by conventional molecular dynamics (cMD). MD simulations
 214 are performed in vacuum for 1.2 ns with 2 fs time step for each metastable state with the AMBER99SB-
 215 ILDN force field. Trajectories and CVs are recorded every 40 fs. Langevin integrator are used to
 216 maintain the system temperature to 300 K. HBonds were constrained during simulations.

217 The free energy surface (FES) is obtained by 800-ns well-tempered metadynamics (WT-MetaD)
 218 simulations. Two backbone dihedral angles are chosen as collective variables (CV). Each CV axis
 219 is evenly discretized with 25 grid points and the Gaussian bias potential are deposited every 2 ps
 220 along the CVs grids. The height and width of the bias potential are set to 0.2 kJ/mol and 0.05 radians
 221 respectively. The 2D FES is collected and summed from the Gaussian deposits. FES is converged to
 222 0.1 kcal/mol after 800 ns by examining the RMSE of FES between adjacent time stamps. To improve
 223 the convergence, we employ the well-tempered version of MetaD with a scaling factor of 8. The
 224 simulation is performed via the OpenMM¹ software.

225 To achieve translation and rotation equivariance, we use the internal coordinate system in addition to
 226 the Cartesian coordinate system, more details can be found in App. B.

227 **Hardware.** All experiments are conducted on two NVIDIA GeForce RTX 4090 GPU cards.

228 4.2 2D Toy Potential: Muller-Brown Potential

Method	Evaluations	MinMax Energy	Max Energy	Distance d_1^2	Distance d_2^3
MCMC	1.03B	-40.21	-17.80 ± 14.77	-	-
Doob's Lagrangian	1.28M	-40.56	-14.81 ± 13.73	-	-
Linear (Random)	N/A	-37.01	7.51 ± 13.0	0.62 ± 0.16	0.13 ± 0.09
Linear (OT)	N/A	-38.98	5.22 ± 17.82	0.59 ± 0.22	0.20 ± 0.14
Ours (shorter)	4K	-40.67	-19.97 ± 12.20	0.29 ± 0.14	0.09 ± 0.07
Ours (longer)	12K	-40.67	-27.98 ± 21.76	0.15 ± 0.11	0.18 ± 0.15
Ours-10K	20K	-40.67	-32.71 ± 16.95	0.12 ± 0.084	0.13 ± 0.12

Table 1: Müller-Brown potential quantitative evaluation. 1,000 paths sampled from models trained on both datasets in Figure 1 and one additional 20K simulation steps dataset. For each sampled 1,000 paths, we report the distribution of maximum energy state along the path and minimum energy of the maximum energy states, and the shortest distance points to the two saddle points.

¹<https://openmm.org/>

229 The Müller-Brown potential has three local minima and two saddle points following the closed-form
 230 potential energy surface in Equation (44). Starting from the initial and final local minima located at
 231 $(-0.56, 1.44)$ and $(-0.05, 0.47)$ in Figure 1, we generate the training data by simulating the first-order
 232 Langevin Dynamics Equation (22) for 4,000, 12,000, and 20,000 steps with $dt = 10^{-4}$. As shown
 233 in Table 1, a minimum of 4,000 steps simulation of molecular dynamics around local minima can
 234 give us decent paths that are close to the saddle points. As the simulation increases to 12,000 and
 235 20,000 steps, it further improves the performance. In addition, our method requires much fewer data
 236 compared to MCMC’s 1.03B and Doob’s Lagrangian’s 1.28M simulation steps. It is worth noting
 237 that our methods cannot sample from the true transition path distribution as MCMC [7] and Doob’s
 238 Lagrangian [8], but we show we can hit around the saddle points quite efficiently.

239 **4.3 Molecular System: Alanine Dipeptide**

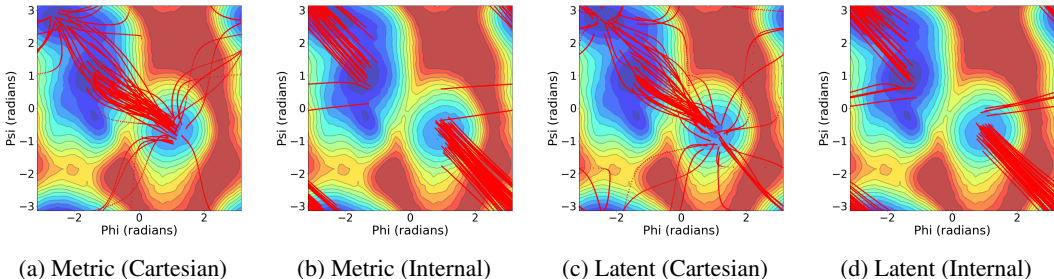


Figure 3: Alanine Dipeptide qualitative evaluation. 50 randomly sampled transition paths are shown for both parameterization in Cartesian and internal coordinate systems with two learned potential energies. Each models are trained over 30,000 data sampled uniformly from a 1.2ns simulation on each metastable states.




Figure 4: Alanine Dipeptide low-energy path visualization. A total of 500 timesteps from one metastable state to another going through an energy barrier.

240 In Figure 3, we visualize 50 randomly sampled transition paths from our method, we can observe
 241 that most of the paths find the correct collective variable (ϕ, ψ) dihedral angles in a much higher
 242 (66) dimensional space only by learning potential energy from a short-run molecular dynamics
 243 simulation around the local minima. In addition, the two ways of learning potential energy result
 244 in similar sampled path distributions. Nevertheless, we find the sampled transition paths with the
 245 flipped dihedral angles. In Figure 4, we demonstrate a qualitative showcase from a low-energy path,
 246 transitioning between the two metastable states.

247 As shown in Table 2, the sampled paths by our method is way better than linear interpolating the two
 248 metastable states, in both Cartesian and internal coordinate systems. In addition, we observe that
 249 the distributions of the maximum energy state over a path sampled by both approaches of learning
 250 potential energies are also close with the latent interpolation method being way faster. Next, we show
 251 how different sample sizes (5K, 10K, 20K and 30K) affect the quality of the sampled paths. In general,
 252 we find the performance improves with more data samples. We also compare different coupling

² d_1 is the shortest distance to the first saddle point $(-0.77, 0.64)$

³ d_2 is the shortest distance to the second saddle point $(0.22, 0.3)$

parameterizations and we observe that the OT coupling is way better than the product measure as initial coupling. However, reflow effectively improves over the product coupling and approach the performance of the OT coupling. In the end, we validate the effectiveness of the resampling procedures. We find in general resampling improves the performance but saturates quickly after one iteration. It is also worth noting that resampling can be expensive as it requires calculating importance weights over an entire trajectory. Additional experimental results and visualizations can be found in App. D.

Potential	Coord	Resampling	Coupling	# sample pairs	# evals	MinMax Energy	Max Energy	Run Time (min)
Linear	Cartesian	N/A	Product	N/A	N/A	2981.64	$2.12e+24 \pm 2.12e+25$	N/A
Linear	Cartesian	N/A	OT	N/A	N/A	1044.95	$4.98e+09 \pm 3.39e+10$	N/A
Linear	Internal	N/A	Product	N/A	N/A	558.05	$1.75e+18 \pm 1.75e+19$	N/A
Linear	Internal	N/A	OT	N/A	N/A	580.10	$3.83e+15 \pm 2.22e+16$	N/A
MCMC*	Cartesian	N/A	N/A	N/A	25.82M	28.67	$1,212.81 \pm 19,444.46$	N/A
MCMC	Cartesian	N/A	N/A	N/A	1.29B*	60.52	288.46 ± 128.31	N/A
Metric	Cartesian	N/A	OT	30K	1.2M	678.13	987.02 ± 210.77	100
Metric	Internal	N/A	OT	30K	1.2M	526.26	698.83 ± 152.48	112
Latent	Cartesian	N/A	OT	30K	1.2M	600.09	949.69 ± 235.92	18
Latent	Internal	N/A	OT	30K	1.2M	525.68	971.97 ± 1673.66	29
Metric	Cartesian	N/A	OT	5K	1.2M	862.38	1581.29 ± 487.10	N/A
Metric	Cartesian	N/A	OT	10K	1.2M	705.65	1058.79 ± 275.15	N/A
Metric	Cartesian	N/A	OT	20K	1.2M	588.96	979.26 ± 267.62	N/A
Metric	Cartesian	N/A	OT	30K	1.2M	664.36	977.28 ± 215.98	N/A
Metric	Cartesian	N/A	Product	30K	1.2M	1312.88	2903.58 ± 1221.16	N/A
Metric	Cartesian	N/A	Reflow-1	30K	1.2M	1023.45	2316.64 ± 829.71	N/A
Metric	Cartesian	N/A	Reflow-2	30K	1.2M	898.30	1960.67 ± 684.71	N/A
Metric	Cartesian	N/A	Reflow-3	30K	1.2M	944.93	1922.53 ± 629.63	N/A
Metric	Cartesian	Resample-1	OT	30K	16.2M	538.86	841.33 ± 176.78	N/A
Metric	Cartesian	Resample-1	OT	30K	31.2M	539.78	848.56 ± 167.92	N/A
Metric	Cartesian	Resample-1	OT	30K	46.2M	603.21	883.41 ± 189.68	N/A

Table 2: Alanine Dipeptide quantitative results. Top 100-weighted paths are evaluated for each setup with different ways of learned potential energies, coordinate systems, coupling parameterizations and resampling steps. In addition, we report the training time for learning potential energy in two different ways. (MCMC results are taken from [8], * indicates variable-length MCMC.)

260 5 Conclusion, Limitation and Future Work

261 In this paper, we propose to simulate transition dynamics in molecular systems by inferring dynamics
 262 from the local dynamics around one metastable state to another. We propose a generalized flow
 263 matching algorithm that additionally optimizes a learned potential energy of the system. In our
 264 scenario, the potential energy is obtained through metric learning or latent space interpolation. We
 265 further employ an importance sampling technique and replay buffer to improve the convergence of
 266 the method. Experimental results demonstrate that the proposed method is capable of finding good
 267 transition paths, and good approximations of transition states with a significantly small number of
 268 energy evaluations.

269 One main limitation of the current framework is that the sampled path distribution could be far
 270 from the transition path distribution. One future direction is to use the learned path distribution as a
 271 proposal distribution for sampling-based approaches, e.g. [12]; another direction is to use the current
 272 framework as an initialization for transition state search methods.

273 References

274 [1] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A
 275 unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

276 [2] Akashnathan Aranganathan, Xinyu Gu, Dedi Wang, Bodhi Vani, and Pratyush Tiwary. Modeling
 277 boltzmann weighted structural ensembles of proteins using ai based methods. 2024.

278 [3] Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the
 279 curvature of deep generative models. In *International Conference on Learning Representations*,
 280 2018.

- 281 [4] Georgios Arvanitidis, Soren Hauberg, and Bernhard Schölkopf. Geometrically enriched latent
282 spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 631–639.
283 PMLR, 2021.
- 284 [5] Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. Transition path
285 sampling: Throwing ropes over rough mountain passes, in the dark. *Annual review of physical*
286 *chemistry*, 53(1):291–318, 2002.
- 287 [6] Avishek Das, Dominic C Rose, Juan P Garrahan, and David T Limmer. Reinforcement learning
288 of rare diffusive dynamics. *The Journal of Chemical Physics*, 155(13), 2021.
- 289 [7] Christoph Dellago, Peter G Bolhuis, and Phillip L Geissler. Transition path sampling. *Advances*
290 *in chemical physics*, 123:1–78, 2002.
- 291 [8] Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes,
292 Alan Aspuru-Guzik, and Kirill Neklyudov. Doob’s lagrangian: A sample-efficient variational
293 approach to transition path sampling. In *ICML 2024 AI for Science Workshop*.
- 294 [9] Chenru Duan, Yuanqi Du, Haojun Jia, and Heather J Kulik. Accurate transition state generation
295 with an object-aware equivariant elementary reaction diffusion model. *Nature Computational*
296 *Science*, 3(12):1045–1055, 2023.
- 297 [10] Chenru Duan, Guan-Horng Liu, Yuanqi Du, Tianrong Chen, Qiyuan Zhao, Haojun Jia, Carla P
298 Gomes, Evangelos A Theodorou, and Heather J Kulik. React-ot: Optimal transport for generat-
299 ing transition state in chemical reactions. *arXiv preprint arXiv:2404.13430*, 2024.
- 300 [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances*
301 *in neural information processing systems*, 33:6840–6851, 2020.
- 302 [12] Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Berend Ensing, and Max Welling. Stochas-
303 tic optimal control for collective variable free sampling of molecular transition paths. *Advances*
304 *in Neural Information Processing Systems*, 36, 2024.
- 305 [13] Kacper Kapusniak, Peter Potapchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael
306 Bronstein, Avishek Joey Bose, and Francesco Di Giovanni. Metric flow matching for smooth
307 interpolations on the data manifold. *arXiv preprint arXiv:2405.14780*, 2024.
- 308 [14] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
309 matching for generative modeling. In *The Eleventh International Conference on Learning*
310 *Representations*.
- 311 [15] Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou,
312 and Ricky TQ Chen. Generalized schrödinger bridge matching. In *The Twelfth International*
313 *Conference on Learning Representations*.
- 314 [16] Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint*
315 *arXiv:2209.14577*, 2022.
- 316 [17] Xingchao Liu, Chengyue Gong, et al. Flow straight and fast: Learning to generate and transfer
317 data with rectified flow. In *The Eleventh International Conference on Learning Representations*.
- 318 [18] Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching:
319 Learning stochastic dynamics from samples. In *International conference on machine learning*,
320 pages 25858–25889. PMLR, 2023.
- 321 [19] Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, and Alireza Liu,
322 Qiang Makhzani. A computational framework for solving wasserstein lagrangian flows. In
323 *Forty-first International Conference on Machine Learning*, 2024.
- 324 [20] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling
325 equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147,
326 2019.

- 327 [21] Lagnajit Pattanaik, John B Ingraham, Colin A Grambow, and William H Green. Generating
328 transition states of isomerization reactions with deep learning. *Physical Chemistry Chemical*
329 *Physics*, 22(41):23618–23626, 2020.
- 330 [22] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data
331 science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- 332 [23] Aram-Alexandre Pooladian, Carles Domingo-Enrich, Ricky TQ Chen, and Brandon Amos.
333 Neural optimal transport with lagrangian costs. In *The 40th Conference on Uncertainty in*
334 *Artificial Intelligence*, 2024.
- 335 [24] Dominic C Rose, Jamie F Mair, and Juan P Garrahan. A reinforcement learning approach to
336 rare trajectory sampling. *New Journal of Physics*, 23(1):013013, 2021.
- 337 [25] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger
338 bridge matching. *Advances in Neural Information Processing Systems*, 36, 2024.
- 339 [26] Søren Smidstrup, Andreas Pedersen, Kurt Stokbro, and Hannes Jónsson. Improved initial guess
340 for minimum energy path calculations. *The Journal of chemical physics*, 140(21), 2014.
- 341 [27] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and
342 Ben Poole. Score-based generative modeling through stochastic differential equations. In
343 *International Conference on Learning Representations*.
- 344 [28] Lixin Sun, Jonathan Vandermause, Simon Batzner, Yu Xie, David Clark, Wei Chen, and Boris
345 Kozinsky. Multitask machine learning of collective variables for enhanced sampling of rare
346 events. *Journal of Chemical Theory and Computation*, 18(4):2341–2353, 2022.
- 347 [29] Alexander Tong, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks,
348 Kilian Fatras, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative
349 models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- 350 [30] Eric Vanden-Eijnden et al. Transition-path theory and path-finding algorithms for the study of
351 rare events. *Annual review of physical chemistry*, 61:391–420, 2010.
- 352 [31] Jiawei Yan, Hugo Touchette, and Grant M Rotskoff. Learning nonequilibrium control forces to
353 characterize dynamical phase transitions. *Physical Review E*, 105(2):024115, 2022.

354 A Additional Details on Metric Learning

355 Following [13], given a dataset of data samples $\{x_i\}_{i=1}^N$, we learn a metric from the Radial Basis
 356 Function (RBF) such that $G(x) = (\text{diag}(h(x)) + \epsilon I)^{-1}$, where $\epsilon > 0$ and the function $h(x)$ is
 357 defined as:

$$h(x) = \sum_{k=1}^K \omega_k \exp\left(-\frac{\lambda_k}{2} \|x - \bar{x}_k\|^2\right) \quad (23)$$

358 where \bar{x}_k is the centroid of the K clusters found by k -means clustering algorithm, w_k is learned
 359 weights and λ_k is a bandwidth associated with each cluster. The bandwidth around each cluster C_k is
 360 defined as follows:

$$\lambda_k = \frac{1}{2} \left(\frac{\kappa}{|C_k|} \sum_{x \in C_k} \|x - \bar{x}_k\|^2 \right)^{-2} \quad (24)$$

361 where C_k is the k -th cluster, \bar{x}_k is the centroid of the cluster, and κ is a hyperparameter that controls
 362 the decay rate of the weight for cluster of different shapes. We then use the following objective
 363 function to learn the weights w_k for each cluster in the dataset \mathcal{D} :

$$\mathcal{L}(\{\omega_k\}_{k=0}^K) = \sum_{x \in \mathcal{D}} (1 - h(x))^2 = \sum_{x \in \mathcal{D}} \left(1 - \sum_{k=1}^K \omega_k \exp\left(-\frac{\lambda_k}{2} \|x - \bar{x}_k\|^2\right) \right)^2 \quad (25)$$

364 B Translational and Rotational Invariance

365 A function f is G -equivariant if $\forall x \in \mathbb{R}^d, g \in G$, we have $f \circ g(x) = g \circ f(x)$. A special case of
 366 equivariant function is the invariant function such that a function f is G -invariant if $\forall x \in \mathbb{R}^d, g \in G$,
 367 we have $f \circ g(x) = f(x)$. The kinetic and potential energy of each state of a molecular system is
 368 invariant to the Euclidean group while the Cartesian coordinate and vector field are equivariant to the
 369 Euclidean group. We represent the molecular system in internal coordinates following [20] instead
 370 of Cartesian coordinates which are constructed by invariant quantities, i.e. distances, angles and
 371 dihedral angles. In this scenario, we remove the unnecessary degrees of freedom, i.e. translations,
 372 rotations and reflections.

373 C Proofs

374 C.1 Conditional Probability Paths

375 Toward deriving our conditional objective, we begin by showing the following lemma [29].

376 **Lemma 1.** *For a given joint distribution $p_{0,T}$ and conditional distribution $p_{t|0,T}$ such that $p_t(x_t) =$
 377 $\mathbb{E}_{p_{0,T}}[p_{t|0,T}(x_t)]$, and a conditional vector field $v_{t|0,T}$ which satisfies the continuity equation,*

$$\frac{\partial}{\partial t} p_{t|0,T}(x_t) = -\nabla \cdot (p_{t|0,T}(x_t) v_{t|0,T}(x_t)), \quad (26)$$

378 *then, under mild conditions, the vector field*

$$v_t(x_t) = \mathbb{E}_{p_{0,T}(x_0, x_T)} \left[\frac{p_{t|0,T}(x_t)}{p_t(x_t)} v_{t|0,T}(x_t) \right] \quad (27)$$

379 *satisfies the continuity equation*

$$\partial_t p_t(x_t) = \mathbb{E}_{p_{0,T}} [\partial_t p_{t|0,T}(x_t) v_{t|0,T}(x_t)] = -\nabla \cdot (p_t(x_t) v_t(x_t)) \quad (28)$$

380 *Proof.* Assuming the Leibniz rule holds, we decompose

$$\frac{\partial}{\partial t} p_t(x_t) = \frac{\partial}{\partial t} \mathbb{E}_{p_{0,T}(x_0, x_T)} [p_{t|0,T}(x_t | x_0, x_T)] = \mathbb{E}_{p_{0,T}(x_0, x_T)} \left[\frac{\partial}{\partial t} p_{t|0,T}(x_t | x_0, x_T) \right] \quad (29)$$

381 We will introduce a vector field $v_{t|0,T}$ which is constrained to satisfy the continuity equation for
 382 $p_{t|0,T}$. Omitting explicit conditioning in the arguments, we write

$$\frac{\partial}{\partial t} p_{t|0,T}(x_t) = -\nabla \cdot (p_{t|0,T}(x_t) v_{t|0,T}(x_t)). \quad (30)$$

383 Finally, we would like to relate the conditional vector field $v_{t|0,T}$ to the marginal vector field v_t in
 384 Eq. (35). Following [29] Thm 3.1, we confirm that

$$v_t(x_t) = \mathbb{E}_{p_{0,T}|t(x_0, x_T|x_t)} [v_{t|0,T}(x_t)] = \frac{1}{p_t(x_t)} \mathbb{E}_{p_{0,T}(x_0, x_T)} [p_{t|0,T}(x_t) v_{t|0,T}(x_t)] \quad (31)$$

385 satisfies the continuity equation relationships in Eq. (29)-Eq. (30), namely

$$\frac{\partial}{\partial t} p_t(x_t) = \mathbb{E}_{p_{0,T}(x_0, x_T)} \left[\frac{\partial}{\partial t} p_{t|0,T}(x_t) \right] = \mathbb{E}_{p_{0,T}(x_0, x_T)} [-\nabla \cdot (p_{t|0,T}(x_t) v_{t|0,T}(x_t))] \quad (32)$$

$$= -\nabla \cdot \mathbb{E}_{p_{0,T}(x_0, x_T)} [p_{t|0,T}(x_t) v_{t|0,T}(x_t)] \quad (33)$$

$$= -\nabla \cdot (p_t(x_t) v_t(x_t)) = \frac{\partial}{\partial t} p_t(x_t) \quad (34)$$

386 where, in the second line, we use the linearity of divergence operator and expectation to swap their
 387 order and, in the third line, we substitute the identity in Eq. (31) to recover $\frac{\partial}{\partial t} p_t(x_t)$. \square

388 C.2 Proof of Conditional Objective

389 *Proof.* We begin from the Fokker-Planck equation formulation of the GFM objective

$$\begin{aligned} \mathcal{L}_{\text{GFM}} &= \min_{v_t, p_t} \int_0^T \int \left(\frac{1}{2} \|v_t(x_t)\|^2 + V_t(x_t) \right) p_t(x_t) dx dt \\ \text{s.t. } \quad \partial_t p_t(x_t) &= -\nabla \cdot (p_t(x_t) v_t(x_t)), \quad p_0 = \mu_0, \quad p_T = \mu_T \end{aligned} \quad (35)$$

390 Assume that a path of marginals $p_t(x_t)$ can be decomposed as $p_t(x_t) = \mathbb{E}_{p_{0,T}} [p_{t|0,T}(x_t)]$, where we
 391 assume the joint distribution $p_{0,T} \in \Pi(\mu_0, \mu_T)$ satisfies the endpoint constraints and the conditional
 392 distribution $p_{t|0,T}$ is suitably smooth (e.g. absolute continuity) such that there exists $v_{t|0,T}$ satisfying
 393 $\partial_t p_{t|0,T} = -\nabla \cdot (p_{t|0,T} v_{t|0,T})$. Under these assumptions, we show that the objective in Eq. (35), as
 394 a function of p_t, v_t , can be upper bounded in terms of the conditional $p_{t|0,T}, v_{t|0,T}$.⁴

395 For any p_t, v_t satisfying the above, we begin by rewriting the continuity equation constraint

$$\partial_t p_t(x_t) = \mathbb{E}_{p_{0,T}(x_0, x_T)} [\partial_t p_{t|0,T}(x_t)] = \mathbb{E}_{p_{0,T}(x_0, x_T)} [-\nabla \cdot (p_{t|0,T} v_{t|0,T})] \quad (36)$$

396 where assumed $p_{0,T}$ satisfies the endpoint marginal constraints.

397 Turning to the objective, we write the expectation using $\mathbb{E}_{p_t} [V_t(x_t)] = \mathbb{E}_{p_{0,T}} \mathbb{E}_{p_{t|0,T}} [V_t(x_t)]$,

$$\int_0^T \int \left(\frac{1}{2} \|v_t(x_t)\|^2 + V_t(x_t) \right) p_t(x_t) dx dt \quad (37)$$

$$= \int_0^T \int \frac{1}{2} \|v_t(x_t)\|^2 p_t(x_t) dx dt + \mathbb{E}_{p_{0,T}(x_0, x_T)} \left[\int_0^T \int V_t(x_t) p_{t|0,T}(x_t) dx dt \right] \quad (38)$$

398 Finally, we would like to express the v_t term in terms of $v_{t|0,T}$ to match the constraint in (36). Using
 the identity Eq. (31) and then Jensen's inequality, we write

$$= \int_0^T \int \frac{1}{2} \|\mathbb{E}_{p_{0,T}|t} [v_{t|0,T}(x_t)]\|^2 p_t(x_t) dx dt + \mathbb{E}_{p_{0,T}} \left[\int_0^T \int V_t(x_t) p_{t|0,T}(x_t) dx dt \right] \quad (39)$$

$$\leq \int_0^T \int \frac{1}{2} \mathbb{E}_{p_{0,T}|t} \|v_{t|0,T}(x_t)\|^2 p_t(x_t) dx dt + \mathbb{E}_{p_{0,T}} \left[\int_0^T \int V_t(x_t) p_{t|0,T}(x_t) dx dt \right] \quad (40)$$

$$= \mathbb{E}_{p_{0,T}} \left[\int_0^T \int \left(\frac{1}{2} \|v_{t|0,T}(x_t)\|^2 + V_t(x_t) \right) p_{t|0,T}(x_t) dx dt \right] \quad (41)$$

⁴We assume the decomposition into suitable $p_t(x_t) = \mathbb{E}_{p_{0,T}} [p_{t|0,T}(x_t)]$ does not change the value of the optimization. Otherwise, a further bound would be induced.

399 where, in the last line, we use the fact that $p_{0,T|t}p_t = p_{0,t}, t, T = p_{0,T}p_{t|0,T}$ for any x_0, x_t, x_T .

400 We have thus shown that any feasible, factorizable p_t, v_t corresponds to an upper bound on the
 401 objective in terms of a coupling $p_{0,T}$ and conditional $p_{t|0,T}, v_{t|0,T}$. To formalize our conclusions, we
 402 define the conditional objective for particular $x_0, x_T \sim p_{0,T}$

$$\mathcal{L}_{\text{cGFM}}(x_0, x_T) := \min_{p_{t|0,T}, v_{t|0,T}} \int_0^T \int \left(\frac{1}{2} \|v_{t|0,T}^\phi(x_t)\|^2 + V_t(x_t) \right) p_{t|0,T}(x_t|x_0, x_T) dx_t dt \quad (42)$$

s.t. $\partial_t p_{t|0,T}(x_t) = -\nabla \cdot (p_{t|0,T}(x_t)v_{t|0,T}^\phi(x_t)), \quad p_{0|0,T} = \delta_{x_0}, \quad p_{T|0,T} = \delta_{x_T}$

403 where we abbreviate the boundary condition $p_{0|0,T}(x) = \delta(x - x_0)$.

404 Finally, we reintroduce the outer expectation over $p_{0,T}$ and consider optimizing over $p_{0,T} \in$
 405 $\Pi(\mu_0, \mu_T)$ (such that $p_0 = \mu_0, p_T = \mu_T$). We can thus conclude that optimizing the conditional
 406 objective upper bounds \mathcal{L}_{GFM} , due to our use of decomposable p_t and Jensen’s inequality in (40)

$$\mathcal{L}_{\text{GFM}} \leq \min_{p_{0,T}} \mathbb{E}_{p_{0,T}} \left[\mathcal{L}_{\text{cGFM}}(x_0, x_T) \right] \quad \text{s.t.} \quad p_0 = \mu_0, \quad p_T = \mu_T \quad (43)$$

407

□

408 D Additional Experimental Details

409 D.1 Müller-Brown Potential


410 The Müller-Brown potential can be written down analytically:

$$\begin{aligned} V(x, y) = & -200 \exp(-(x-1)^2 - 10y^2) \\ & -100 \exp(-x^2 - 10(y-0.5)^2) \\ & -170 \exp(-6.5(x+0.5)^2 + 11(x+0.5)(y-1.5) - 6.5(y-1.5)^2) \\ & +15 \exp(0.7(x+1)^2 + 0.6(x+1)(y-1) + 0.7(y-1)^2) \end{aligned} \quad (44)$$

411 The potential energy function, neural spline network, and velocity network are trained for 100
 412 epochs, respectively with a batch size of 256. We use the Adam optimizer for all training with 10^{-2} ,
 413 10^{-5} and 10^{-3} , respectively. For each neural network, we use a three-layer MLP, with 128 hidden
 414 units per layer, and the SELU activation function. The clustering bandwidth for metric learning in
 415 Equation (24) is set to $\kappa = 1.5$. The number of clusters K is set to 100 in Equation (23).

416 D.2 Alanine Dipeptide


417 The potential energy function, neural spline network, and velocity network are trained for 400, 100
 418 and 100 epochs, respectively with a batch size of 512. We use the Adam optimizer for all training
 419 with 10^{-2} , 10^{-5} and 10^{-3} , respectively. For each neural network, we use a three-layer MLP, with
 420 128 hidden units per layer, and the SELU activation function. The clustering bandwidth for metric
 421 learning in Equation (24) is set to $\kappa = 1.5$. The number of clusters K is set to 150 in Equation (23).
 422 And in metric formula, we set ϵ to 0.001. To obtain the latent space for interpolation, we train a
 423 32-dimension latent space VAE with a three-layer MLP encoder with 64 hidden units and another
 424 three-layer MLP decoder with 64 hidden units. Note the learning rate is 10^{-3} instead of 10^{-2} . For
 425 all experiments, we use the scale of 0.1 Angstrom as unit length.



(a) Linear with OT

(b) Linear without OT

Figure 5: Linear interpolation paths for the Longer-run dataset with and without OT. Saddle points are labeled. Only 100 paths are selected.




(a) Cartesian without OT

(b) Cartesian with OT

(c) Internal without OT

(d) Internal with OT

Figure 6: Linear interpolation paths for Alanine Dipeptide with 30K data, of which 50 pathways are randomly selected.




(a) Without resampling

(b) One iterations

(c) Two iterations

(d) Three iterations

Figure 7: Qualitative evaluation of resampling procedures on Alanine Dipeptide. Results for resampling in the Cartesian coordinate system with multiple iterations are reported.



(a) Without reflow

(b) One iterations

(c) Two iterations

(d) Three iterations

Figure 8: Qualitative evaluation of reflow procedures on Alanine Dipeptide. Results for reflow iterations in the Cartesian coordinate system with multiple iterations are reported. 50 paths are randomly selected.




Figure 9: Qualitative evaluation of initial coupling on Alanine Dipeptide. We show results for both product measure and OT couplings.

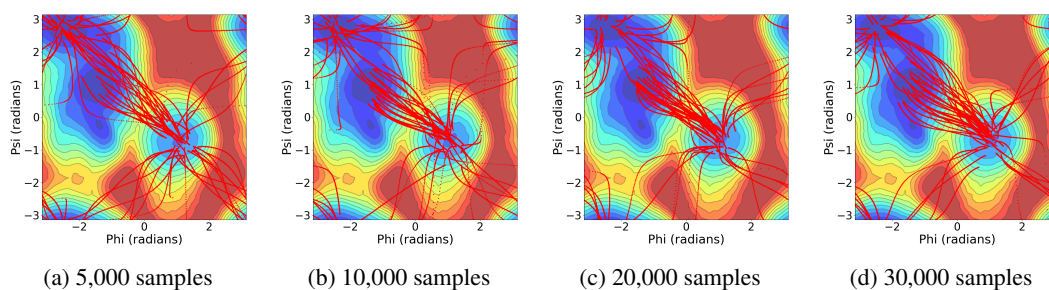


Figure 10: Qualitative evaluation of data sample sizes on Alanine Dipeptide. 50 randomly selected transition paths are shown for models trained with 5,000, 10,000, 20,000, and 30,000 data samples.