DISCRETE ACTION MATCHING: LEARNING STOCHASTIC DYNAMICS FROM SAMPLES VIA STATE GRAPHS

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning population dynamics from unpaired temporal marginals is an ill-posed inverse problem that requires structural assumptions on the underlying dynamics. We introduce *Discrete Action Matching* (DAM), a finite-state counterpart of Action Matching based on discrete Wasserstein geometry on graphs. For a prescribed marginal path and transport geometry, we derive an action-minimization objective for its canonical minimum-kinetic-energy current. Our key observation is that the density dependence of the discrete action reduces to neighboring density ratios. Along an empirical interpolation of the snapshots, DAM first estimates these ratios and then learns an action potential. The learned fields also define a graph-supported Markov sampler. Experiments on controlled synthetic dynamics and real mouse gastrulation data evaluate marginal reconstruction and interpolation. Additional experiments approximate numerical surface-transport paths from paired samples.

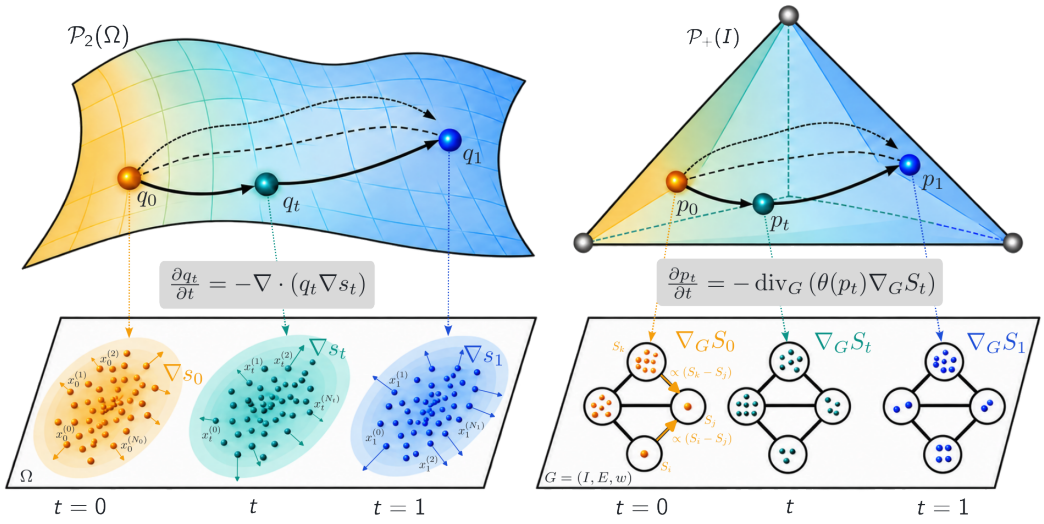


Figure 1: **Action Matching (Neklyudov et al., 2023, left) and Discrete Action Matching (ours, right).** Continuity equations relate prescribed marginal paths (top) to potential fields and samples in the corresponding state spaces (bottom). Dashed curves illustrate alternative paths. Notation is introduced at the end of Section 1.

1 INTRODUCTION

Learning population dynamics from temporal snapshots is a ubiquitous inverse problem arising when populations can be observed repeatedly but individual trajectories cannot be tracked (Lavenant et al., 2024; Schiebinger et al., 2019; Tong et al., 2020). In temporal single-cell data, for example, successive measurements describe changing cell populations without revealing the transitions of individual cells (Pijuan-Sala et al., 2019). Each snapshot provides samples from a population distribution, while the motion connecting these distributions remains unobserved. Learning dynamics from such data therefore requires a principle for selecting among compatible motions.

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.

Optimal-transport geometry (Lott, 2007) provides a natural selection principle: among all motions producing the same instantaneous change in a distribution, choose the one with minimum kinetic energy (Benamou & Brenier, 2000; Ambrosio et al., 2008). Action Matching (AM) (Neklyudov et al., 2023) turns this principle into a learning objective for a prescribed marginal path on a continuous state space. It recovers the canonical velocity field from marginal samples without observing velocities or individual trajectories. Subsequent work extends this variational perspective to optimizing marginal paths and learning population mechanics (Neklyudov et al., 2024; Guan et al., 2026a).

Many population models describe individuals through discrete states, such as cell types or epidemiological compartments (Pijuan-Sala et al., 2019; Kermack & McKendrick, 1927). A graph specifies which states can exchange probability. Discrete Wasserstein geometry (Li, 2025) equips such graphs with continuity equations and kinetic energies (Maas, 2011; Chow et al., 2012; Mielke, 2011). In this geometry, the representation of minimum-energy population motion by graph potentials is an established result (Chow et al., 2019; Li, 2022). Related potential constructions also appear in kinetic-optimal discrete generative flows (Shaul et al., 2025). These results provide the geometric ingredients for a finite-state counterpart of Action Matching.

The remaining difficulty concerns learning from samples. In continuous AM, the kinetic term is an expectation over the marginal distribution. On a graph, its counterpart involves an edge mobility that depends on the probability masses at both neighboring states. For nonlinear mobilities, these unknown masses prevent direct evaluation of the kinetic term from a single marginal sample. A sample-based method must therefore account for the marginal-dependent mobility as well as the potential driving the motion. We introduce *Discrete Action Matching* (DAM), a finite-state counterpart of Action Matching that addresses this difficulty through local density ratios. Such ratios are established quantities in normalization-free estimation on discrete spaces (Hyvärinen, 2007; Dawid et al., 2012; Kanamori & Takenouchi, 2017) and in discrete score-based models (Meng et al., 2022; Lou et al., 2024). Our key observation is that positive 1-homogeneity expresses an edge mobility as the probability of one endpoint multiplied by a function of the neighboring density ratio. The endpoint probability enters through the sampling distribution, leaving the ratio as the quantity to estimate. This makes the discrete action locally computable from samples and separates density-ratio estimation from learning the action potential.

The graph and mobility function are prescribed in our formulation. For a given marginal path, DAM selects its canonical minimum-kinetic-energy current and provides a Markov construction for trajectory generation. With finite observations, we specify the empirical interpolation used for training and assess agreement with held-out population measurements separately. This distinction separates the choice of the marginal path from the canonical motion along it; historical individual trajectories remain unidentifiable from unpaired marginals (Lavenant et al., 2024; Guan et al., 2024).

Contributions. Our main contributions are as follows:

- **Discrete Action Matching.** In §3.1, we derive a finite-state action-gap identity and show that minimizing the population objective recovers the canonical minimum-kinetic-energy current of a prescribed positive marginal path under a chosen graph transport geometry.
- **Learning from unpaired marginal samples.** We estimate neighboring density ratios with a local objective (§3.2), then freeze the ratio model and learn the action potential (§3.3). Both stages use samples from the same positive empirical interpolant.
- **Markov realization.** In §3.3, we construct graph-supported nonnegative transition rates and establish the conditions under which the resulting Markov chain reproduces the prescribed marginal path. Appendix B.3 gives the derivation and the effect of ratio-estimation errors.

In §5 and Appendix D, we evaluate marginal reconstruction and interpolation on controlled synthetic dynamics and mouse gastrulation data. Comparisons with the adapted graph-gradient-flow learner of Rancati et al. (2026) show lower mean marginal errors, with histogram-based DAM remaining competitive with its learned-ratio counterpart. Supplementary sphere and hand experiments approximate numerical surface-transport paths (Lavenant et al., 2018) using paired samples (Appendix D.6).

The remainder of the paper is organized as follows. Section 2 introduces the observation model, Action Matching, and discrete Wasserstein geometry. Section 3 develops DAM and its training and sampling procedures. Section 4 discusses related methods, Section 5 presents the main experiments,

and Section 6 discusses limitations and future directions. Detailed derivations and experimental protocols are provided in the Appendix.

Notation. We use Ω for a continuous state space and $I = \{1, \dots, N\}$ for a finite state space. For $t \in [0, 1]$, the corresponding marginals are denoted by $q(t)$ and $p(t)$, where $q(t)$ is a density and $p(t)$ is a probability vector with entries $p_i(t)$. We write $\mathcal{P}_2(\Omega)$ for probability measures with finite second moment and $\mathcal{P}_+(I) = \{p \in \mathbb{R}^N : p_i > 0,\ \sum_i p_i = 1\}$ for the positive probability simplex. The graph $G = (I, E, \omega)$ is connected and undirected, with edge set E and symmetric positive edge weights ω_{ij}. The scalar fields $s(t, \cdot)$ and $S(t, \cdot)$ are continuous and discrete action potentials, respectively. Their gradients ∇s and $\nabla_G S$ describe the associated continuous velocity and graph edge field; div_G is the graph divergence (§2.3). The notation $\theta(p)$ denotes edgewise mobility, with entries $\theta(p_i, p_j)$. Figure 1 uses the shorthand $q_t = q(t)$, $p_t = p(t)$, $s_t = s(t, \cdot)$, and $S_t = S(t, \cdot)$, and writes $w_{ij} = \omega_{ij}$. Its upper panels show marginal paths, while its lower panels show potential fields and samples $x_t^{(m)}$, with N_t denoting the snapshot size.

2 Background

We formulate learning population dynamics as an inverse problem from unpaired temporal marginals and review the geometric principle that selects canonical dynamics for a prescribed marginal path. We first introduce the observation model and distinguish the evolving population distribution from the unobserved trajectories of individual samples (§2.1). We then review continuous Action Matching (§2.2) and the graph differential operators and discrete Wasserstein geometry needed to transfer this construction to finite state spaces (§2.3). Detailed derivations are provided in Appendix A.

2.1 Learning dynamics from unpaired marginal snapshots

Let $\mathcal{X}$ denote the state space of an unknown stochastic process $(X_t)_{t\in[0,1]}$, and let $\rho(t) := \mathcal{L}(X_t)$ denote its one-time marginal law (E et al., 2019). We use $\mathcal{X} = \Omega$ and write $q(t)$ when the state space is continuous, and use $\mathcal{X} = I$ and write $p(t)$ when it is finite. The process induces a probability law on trajectories, whereas the collection $(\rho(t))_{t\in[0,1]}$ contains only its one-time marginals. At finitely many observation times $0 = t_0 < t_1 < \cdots < t_K = 1$, we observe independent samples

$$\mathcal{D}_k := \{x_k^{(m)}\}_{m=1}^{N_k}, \qquad x_k^{(m)} \overset{\text{i.i.d.}}{\sim} \rho(t_k), \tag{1}$$

with no correspondence between observations collected at different times. This observation model arises naturally when measurements are cross-sectional or destructive, as in temporal single-cell data (Schiebinger et al., 2019; Pijuan-Sala et al., 2019), and more generally whenever population snapshots are available but individual trajectories cannot be tracked (Tong et al., 2020). Marginal observations do not determine how probability mass moves between states: different microscopic dynamics can generate the same marginal evolution, and finite snapshots additionally leave the intermediate marginal path unknown (Lavenant et al., 2024; Guan et al., 2024). Population-dynamics inference therefore requires structural assumptions, such as path-space regularization (Chizat et al., 2022), Wasserstein gradient flows (Bunne et al., 2022; Terpin et al., 2024; Persiianov et al., 2026; Guan et al., 2026b), or action-based principles (Neklyudov et al., 2023; 2024; Guan et al., 2026a).

We seek the *canonical dynamics* selected by a prescribed transport geometry and a sufficiently regular marginal path. With finite observations, this requires a choice of the path between snapshots; our training procedure uses the positive empirical interpolant defined in equation 17. The chosen path is part of the modeling assumptions.

A natural way to select canonical dynamics is to choose, among all motions producing the same infinitesimal change of the marginal, the one with minimum kinetic energy. This principle underlies optimal-transport geometry (Benamou & Brenier, 2000; Otto, 2001; Ambrosio et al., 2008) and, for a prescribed marginal path, Action Matching (Neklyudov et al., 2023). We review this continuous construction next before transferring it to a finite state space.

2.2 Action Matching on continuous sample spaces

Consider a smooth, uniformly positive path $q(t)$ of probability densities on a compact connected Riemannian domain (Ω, g^Ω). The gradient, divergence, and volume element dx are induced by the

metric g^Ω (Lee, 2018). We use periodic or zero normal probability-flux boundary conditions and the regularity assumptions detailed in Appendix A.1. A velocity field $v(t,\cdot)$ generates the marginal evolution through the *continuity equation*

$$\partial_t q(t) + \operatorname{div}(q(t)v(t)) = 0, \tag{2}$$

which expresses conservation of probability mass (Benamou & Brenier, 2000). The marginal tangent $\partial_t q(t)$ does not determine $v(t)$ uniquely: adding a field with zero $q(t)$-weighted divergence leaves the marginal evolution unchanged. Wasserstein geometry selects the compatible velocity of smallest $L^2(q(t))$ norm (Ambrosio et al., 2008; Lott, 2007). In the smooth positive setting above, this velocity is a gradient field $\nabla s^\star(t)$, with $s^\star(t,\cdot)$ unique up to an additive constant:

$$\nabla s^\star(t) \in \operatorname*{arg\,min}_{\partial_t q(t) = -\operatorname{div}(q(t)v)} \frac{1}{2}\int_\Omega g_x^\Omega(v(x),v(x))\, q(t,x)\, dx. \tag{3}$$

The general theory formulates Wasserstein tangent spaces using the weighted L^2 closure of gradients (Ambrosio et al., 2008, Theorem 8.3.1 and Section 8.4). Here the smooth setting allows us to work directly with potentials. The relation between this minimum-energy characterization, the weighted Hodge decomposition, and dynamical optimal transport is reviewed in Appendix A.1.

Action Matching (Neklyudov et al., 2023) learns $s^\star$ without observing either $\partial_t q(t)$ or trajectories. For a sufficiently regular candidate potential s, define the *action gap*

$$\mathrm{AG}(s, s^\star) := \frac{1}{2}\int_0^1 \mathbb{E}_{x\sim q(t)}\left[\|\nabla s(t,x) - \nabla s^\star(t,x)\|^2_{g_x^\Omega}\right] dt, \tag{4}$$

where $\|u\|^2_{g_x^\Omega} := g_x^\Omega(u,u)$. This discrepancy cannot be evaluated directly because $s^\star$ is unknown. Integration by parts in space and time removes the unknown target from the optimization. The Euclidean identity of Neklyudov et al. (2023, Theorem 2.2) extends under the assumptions above to $\mathrm{AG}(s,s^\star) = \mathcal{L}_{\mathrm{AM}}(s) - \mathcal{L}_{\mathrm{AM}}(s^\star)$, where

$$\begin{aligned}\mathcal{L}_{\mathrm{AM}}(s) := {}& \mathbb{E}_{x\sim q(0)}[s(0,x)] - \mathbb{E}_{x\sim q(1)}[s(1,x)] \\ & + \int_0^1 \mathbb{E}_{x\sim q(t)}\left[\partial_t s(t,x) + \frac{1}{2}\|\nabla s(t,x)\|^2_{g_x^\Omega}\right] dt.\end{aligned} \tag{5}$$

Thus minimizing $\mathcal{L}_{\mathrm{AM}}$ over admissible potentials recovers the canonical gradient field. Every occurrence of $q(t)$ in this objective appears through an expectation, so the loss can be estimated from marginal samples without evaluating the density, its time derivative, or velocities (Neklyudov et al., 2023). This requires sample access at the integration times; finite snapshots alone do not provide samples from the unknown intermediate marginals.

For our development, the key feature of equation 5 is its geometric structure. The construction requires a *gradient* that maps a scalar potential to a local velocity and a density-dependent *inner product* that measures kinetic energy. Discrete Wasserstein geometry provides graph counterparts of these objects (Maas, 2011; Chow et al., 2012; Mielke, 2011). We introduce them next.

2.3 Graph Wasserstein geometry and canonical discrete dynamics

We now specialize to a finite state space $I = \{1,\dots,N\}$ equipped with a connected undirected graph $G = (I, E, \omega)$ (Li, 2016). Its symmetric weights satisfy $\omega_{ij} = \omega_{ji} > 0$ for $\{i,j\} \in E$, and $\mathcal{N}(i)$ denotes the neighbors of i. The graph specifies which pairs of states may exchange probability locally. We work on the positive probability simplex $\mathcal{P}_+(I) := \left\{p \in \mathbb{R}^N : p_i > 0,\ \sum_{i=1}^N p_i = 1\right\}$. Throughout, E contains unordered edges, and each edge is counted once in sums over E.

For a scalar potential $S \in \mathbb{R}^N$ and a skew-symmetric edge field v, meaning $v_{ij} = -v_{ji}$, define

$$(\nabla_G S)_{ij} := \sqrt{\omega_{ij}}(S_i - S_j), \tag{6}$$

$$(\operatorname{div}_G v)_i := \sum_{j\in\mathcal{N}(i)} \sqrt{\omega_{ij}}\, v_{ji}. \tag{7}$$

With this sign convention, $-\operatorname{div}_G$ is the adjoint of ∇_G under the unweighted vertex and edge pairings (Chow et al., 2019; Li, 2022); see Appendix A.2.

The remaining ingredient is the counterpart of multiplication by the continuous density. Probability masses live on vertices, whereas velocities live on edges, so a graph transport geometry assigns an effective *mobility* to each edge from the masses at its endpoints (Maas, 2011; Mielke, 2011). We use a symmetric mobility $\theta : [0,\infty)^2 \to [0,\infty)$ that is smooth and strictly positive on $(0,\infty)^2$ and positively 1-homogeneous: $\theta(\lambda a, \lambda b) = \lambda\theta(a,b), \lambda > 0$. The full adopted assumptions and the arithmetic, geometric, harmonic, and logarithmic examples are given in Definition A.1. Different choices of θ generally induce different transport geometries. Write $\theta_{ij}(p) := \theta(p_i, p_j)$ and interpret $\theta(p)v$ edgewise. The mobility-weighted edge inner product and its norm are

$$\langle v, w\rangle_{p,\theta} := \sum_{\{i,j\}\in E} \theta_{ij}(p) v_{ij} w_{ij}, \qquad \|v\|^2_{p,\theta} := \langle v, v\rangle_{p,\theta}.$$

The *graph continuity equation* for a potential field is

$$\dot{p}(t) = -\operatorname{div}_G(\theta(p(t))\nabla_G S(t)) =: L_\theta(p(t))S(t), \tag{8}$$

or, componentwise,

$$\dot{p}_i(t) = \sum_{j\in\mathcal{N}(i)} \omega_{ij}\theta(p_i(t), p_j(t))\big(S_i(t) - S_j(t)\big). \tag{9}$$

The corresponding kinetic energy is

$$\mathcal{K}_{G,\theta}(p, S) := \frac{1}{2}\|\nabla_G S\|^2_{p,\theta} = \frac{1}{2}\sum_{\{i,j\}\in E} \omega_{ij}\theta(p_i, p_j)(S_i - S_j)^2. \tag{10}$$

These operators underlie finite-state transport constructions (Maas, 2011; Chow et al., 2012; Mielke, 2011). Here p_i denotes probability mass; the conversion to the stationary-relative densities used in reversible-chain formulations is explained in Appendix A.3.

The representation of marginal tangents by potential fields is standard in discrete Wasserstein geometry (Maas, 2011, Section 3, Riemannian structure). For positive p on a connected graph, the mobility-weighted Laplacian $L_\theta(p)$ is positive semidefinite, has kernel $\operatorname{span}\{\mathbf{1}\}$, and has range $T_p\mathcal{P}_+(I) = \{\sigma \in \mathbb{R}^N : \mathbf{1}^\top\sigma = 0\}$. Consequently, every marginal tangent $\sigma \in T_p\mathcal{P}_+(I)$ admits a potential $S^\star$, unique up to an additive constant, such that

$$\sigma = L_\theta(p)S^\star = -\operatorname{div}_G(\theta(p)\nabla_G S^\star)\,. \tag{11}$$

The weighted graph Hodge decomposition identifies $\nabla_G S^\star$ as the unique minimum-energy edge field with this tangent (Maas, 2011, Lemma 3.6); see also Chow et al. (2019, Lemma 4):

$$\nabla_G S^\star \in \operatorname*{arg\,min}_{\substack{v_{ij}=-v_{ji}\\ \sigma=-\operatorname{div}_G(\theta(p)v)}} \frac{1}{2}\|v\|^2_{p,\theta}. \tag{12}$$

Appendix A.3 reviews this construction in our notation, including the graph-Laplacian representation and the continuous–discrete operator dictionary in Table 1. In particular, along a smooth positive probability path there is a canonical potential $S^\star(t)$, unique up to a time-dependent constant

$$\dot{p}(t) = L_\theta(p(t))S^\star(t). \tag{13}$$

This construction supplies the geometric ingredients needed for a finite-state counterpart of Action Matching. A new difficulty arises when learning from samples: the discrete kinetic energy depends on the unknown endpoint masses through $\theta(p_i(t), p_j(t))$. For nonlinear mobilities, this dependence prevents direct evaluation of an edge contribution from an individual marginal sample. In §3, we use positive 1-homogeneity to reduce this dependence to neighboring density ratios and obtain a sample-based action objective.

3 METHOD

We first derive the discrete action objective (§3.1) and describe density-ratio estimation along an empirical interpolant (§3.2). We then learn the action potential with the ratio model held fixed and construct a Markov realization of the dynamics (§3.3).

3.1 Discrete Action Matching

Let $p \in C^2([0,1]; \mathcal{P}_+(I))$ be a prescribed probability path on the fixed connected graph G, with a symmetric mobility that is smooth and positive on the interior. By the canonical representation in equation 11, there exists a potential $S^\star(t) \in \mathbb{R}^N$, unique up to an additive time-dependent constant, satisfying equation 13. Its graph gradient is the minimum-kinetic-energy representative of the marginal tangent $\dot{p}(t)$. Our goal is to learn this canonical potential from marginal samples, without observing either $\dot{p}(t)$ or trajectories. Following continuous Action Matching (Neklyudov et al., 2023), for a candidate $S \in C^1([0,1]; \mathbb{R}^N)$ we define the *discrete action gap*

$$\mathrm{AG}_{G,\theta}(S, S^\star) := \frac{1}{2} \int_0^1 \|\nabla_G S(t) - \nabla_G S^\star(t)\|^2_{p(t),\theta} \, dt, \tag{14}$$

using the mobility-weighted edge norm introduced in the §2. Direct minimization of equation 14 is unavailable because the target $S^\star$ is unknown. Graph summation by parts, followed by integration by parts in time, eliminates this target from the optimization.

Theorem 3.1 (Discrete Action Matching)**.** *Fix a connected undirected graph with symmetric positive edge weights and a symmetric mobility θ that is smooth and strictly positive on $(0,\infty)^2$. Let $p \in C^2([0,1]; \mathcal{P}_+(I))$ and choose $S^\star(t) = L_\theta(p(t))^\dagger \dot{p}(t)$, where $\dagger$ denotes the Moore–Penrose pseudoinverse (Moore, 1920; Penrose, 1955). This is the zero-mean canonical potential; the constant rank of $L_\theta(p(t))$ gives $S^\star \in C^1([0,1]; \mathbb{R}^N)$. For $S \in C^1([0,1]; \mathbb{R}^N)$, define*

$$\begin{aligned} \mathcal{L}_{\mathrm{DAM}}(S) := \mathbb{E}_{i\sim p(0)}[S(0,i)] - \mathbb{E}_{i\sim p(1)}[S(1,i)] \\ + \int_0^1 \left\{ \mathbb{E}_{i\sim p(t)}[\partial_t S(t,i)] + \mathcal{K}_{G,\theta}(p(t), S(t)) \right\} dt, \end{aligned} \tag{15}$$

where kinetic energy $\mathcal{K}_{G,\theta}(p,S)$ defined at equation 10. Then $\mathcal{L}_{\mathrm{DAM}}(S) - \mathcal{L}_{\mathrm{DAM}}(S^\star) = \mathrm{AG}_{G,\theta}(S, S^\star)$. Consequently, the minimizers over $C^1([0,1]; \mathbb{R}^N)$ are precisely $S(t) = S^\star(t) + c(t)\mathbf{1}$, $c \in C^1([0,1])$, with equality at every $t \in [0,1]$.

The proof in Appendix A.3 follows the graph counterpart of the continuous Action Matching identity of Neklyudov et al. (2023, Theorem 2.2). The theorem identifies the unrestricted population target for a prescribed path and geometry; a restricted network class or finite-data optimization need not attain this target. When the probabilities are available, equation 15 can be evaluated directly. For learning from samples, its boundary and time-derivative terms already have the required expectation form. For nonlinear mobilities, the remaining difficulty is the kinetic energy, whose edge contributions depend on the unknown quantities $\theta(p_i(t), p_j(t))$.

For sample-based evaluation, assume additionally that θ is positively 1-homogeneous. With $\kappa_\theta(u) := \theta(1, e^u)$, homogeneity gives $\theta(p_i, p_j) = p_i \kappa_\theta(\log(p_j/p_i))$, hence

$$\mathcal{K}_{G,\theta}(p,S) = \frac{1}{4} \mathbb{E}_{i\sim p} \left[\sum_{j \in \mathcal{N}(i)} \omega_{ij} \kappa_\theta\left(\log \frac{p_j}{p_i} \right) (S_i - S_j)^2 \right]. \tag{16}$$

Thus local ratio estimates suffice to evaluate the action without reconstructing the full probability vector.[1] Appendix B.2 derives this representation, including the factor $1/4$.

3.2 Learning local density ratios

We estimate neighboring ratios through a scalar log-density field, following normalization-free estimation on discrete spaces (Hyvärinen, 2007; Kanamori & Takenouchi, 2017). The field $U_\phi(t,i)$ targets $U^\star(t,i) = \log p_i(t) + c(t)$. Its edge differences represent log-ratios and enforce reciprocity and cycle consistency (Appendix B); S_η separately parameterizes the dynamics.

To obtain intermediate-time samples from finite snapshots, let $\widehat{p}_k$ be the empirical distribution of $\mathcal{D}_k$, set $r_i = 1/N$, and fix $0 < \delta < 1$. For $u = (t - t_k)/(t_{k+1} - t_k)$ and $h(u) = 3u^2 - 2u^3$, declare

$$\widetilde{p}(t) = (1-\delta)\big[(1 - h(u))\widehat{p}_k + h(u)\widehat{p}_{k+1}\big] + \delta r, \qquad t \in [t_k, t_{k+1}]. \tag{17}$$

[1] For arithmetic mobility, the kinetic action admits a ratio-free sample estimator. Our Markov realization still uses density ratios; see Appendix B.2.

Sampling this mixture requires no pairing across snapshots and ensures $\widetilde{p}_i(t) \geq \delta/N$. Both training stages use this same path and its reference-mixed endpoints. Interpolation and reference mixing define the statistical target and can bias it relative to the unknown population evolution. The path is C^1 and piecewise smooth; the action identity extends intervalwise as detailed in Appendix B.3.

Using fresh draws $T_b \sim \mathrm{Unif}[0,1]$ and $I_b \sim \widetilde{p}(T_b)$, we minimize

$$\widehat{\mathcal{R}}(\phi) := \frac{1}{B}\sum_{b=1}^{B} \sum_{j \in \mathcal{N}(I_b)} \omega_{I_b j}\, \mathrm{softplus}(U_\phi(T_b, j) - U_\phi(T_b, I_b))\,, \tag{18}$$

where $\mathrm{softplus}(z) := \log(1+e^z)$. For each positive marginal, the population minimizers are exactly $U = \log p + c\mathbf{1}$ (Appendix B.2, equation 57). Under simultaneous realizability and exact population minimization, the time-integrated objective therefore recovers the ratios of $\widetilde{p}$ almost everywhere; continuity of the fitted ratios extends equality to every time. This does not establish finite-sample recovery of the unknown path.

3.3 Training and Markov realization

After fitting U_ϕ, we freeze ϕ and optimize S_η in the estimated geometry. Define

$$\widehat{\mathcal{K}}_{\eta,\phi}(t,i) := \frac{1}{4}\sum_{j \in \mathcal{N}(i)} \omega_{ij}\kappa_\theta(U_\phi(t,j) - U_\phi(t,i))\,\big(S_\eta(t,j) - S_\eta(t,i)\big)^2. \tag{19}$$

With fresh (T_b, I_b) and independent endpoint samples $I_b^0 \sim \widetilde{p}(0)$ and $I_b^1 \sim \widetilde{p}(1)$, minimize

$$\widehat{\mathcal{L}}_{\mathrm{DAM}}(\eta;\phi) := \frac{1}{B}\sum_{b=1}^{B}\Big[S_\eta(0, I_b^0) - S_\eta(1, I_b^1) + \partial_t S_\eta(T_b, I_b) + \widehat{\mathcal{K}}_{\eta,\phi}(T_b, I_b)\Big]. \tag{20}$$

Time derivatives use automatic differentiation in PyTorch (Paszke et al., 2019). Conditional on the data and frozen U_ϕ, this estimator is unbiased at fixed S_η for the action with the effective positive symmetric conductances in equation 58. Exact ratios recover the intended-mobility objective. Appendix B.3 gives regularity and estimation details.

For sampling, let $\psi_\varepsilon(z) := \varepsilon \log(1 + e^{z/\varepsilon})$, $\varepsilon > 0$, so $\psi_\varepsilon(z) - \psi_\varepsilon(-z) = z$. For neighboring states $i \neq j$, set

$$Q_t^{\eta,\phi}(i,j) := \omega_{ij}\kappa_\theta(U_\phi(t,j) - U_\phi(t,i))\,\psi_\varepsilon(S_\eta(t,j) - S_\eta(t,i))\,. \tag{21}$$

Other off-diagonal entries vanish and $Q_t^{\eta,\phi}(i,i) := -\sum_{j\neq i} Q_t^{\eta,\phi}(i,j)$. Finite continuous fields give a bounded graph-supported Markov generator (Anderson, 1991). Its column-vector marginal law satisfies $\dot{\mu} = Q_t^\top \mu$ (Pauline et al., 2025). At exact ratios for a positive target path p, $Q_t^\top p = L_\theta(p) S_\eta$. If additionally S_η is canonical and $\mu(0) = p(0)$, uniqueness gives $\mu(t) = p(t)$ throughout. The split selects one of multiple compatible Markov realizations.

Imperfect ratios preserve generator validity but can separate its current from the effective training current, even at zero potential difference (Appendix B.3). For finite snapshots, ideal reproduction concerns the declared $\widetilde{p}$; approximate fits generally yield a different law μ.

4 Related Work

Action Matching. Action Matching learns continuous population dynamics from independent temporal marginals through a sample-based variational objective (Neklyudov et al., 2023). Subsequent work extends this principle to Wasserstein Lagrangian Flows (Neklyudov et al., 2024), parameter-dependent dynamics (Berman et al., 2024), and more general Wasserstein mechanics (Guan et al., 2026a). DAM transfers the fixed-path Action Matching principle to finite state spaces using mobility-dependent discrete Wasserstein geometry.

Discrete Wasserstein geometry and gradient-flow learning. Discrete transport geometries on graphs provide finite-state analogues of Wasserstein continuity equations, kinetic energies, and gradient flows (Maas, 2011; Chow et al., 2012; Erbar, 2014). These constructions have subsequently

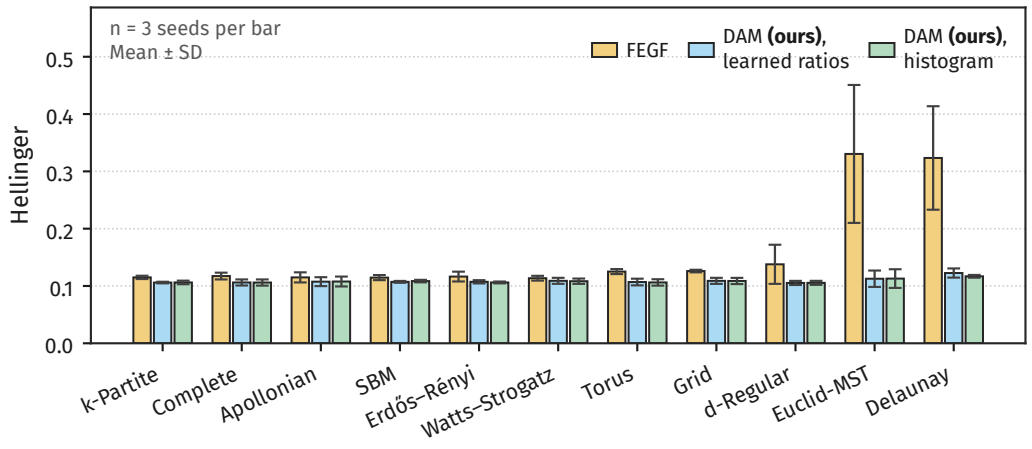


Figure 2: Synthetic reconstruction across eleven graph families with $N = 60$, $M = 1{,}000$ samples per split and time, and $\beta = 0.1$. DAM uses logarithmic mobility; lower is better.

been developed as Riemannian geometries on the probability simplex and used to study discrete gradient-flow dynamics (Chow et al., 2018; Li & Montúfar, 2018; Li, 2022). Most closely related to our work, Rancati et al. (2026) learn finite-state population dynamics by assuming an autonomous free-energy gradient flow and recovering its energy through discrete-JKO optimality conditions, analogous to the continuous inverse-JKO approach of Terpin et al. (2024). Their method requires numerical computation of discrete-Wasserstein geodesics between observed marginals. DAM addresses a complementary inverse problem: it learns the canonical, potentially time-dependent velocity of a prescribed marginal evolution and uses local density ratios to make the corresponding action directly estimable from samples without discrete-geodesic preprocessing.

Further connections to discrete transport, inverse gradient-flow learning, and discrete generative models are discussed in Appendix C.1.

5 Experiments

We adapt the synthetic and mouse-gastrulation benchmarks of Rancati et al. (2026) and compare DAM with their free-energy gradient-flow method (FEGF), starting from the official implementation.[2] Methods receive matched sampled inputs, with validation restricted to retained observation times. Learned-ratio and histogram DAM share the positive empirical interpolant and action-training settings. Figures report mean Hellinger errors and sample standard deviations across three seeds. Protocol adaptations, including FEGF's time-indexed gastrulation potential and the methods' different transport geometries, are detailed in Appendix D.2.

5.1 Synthetic graph dynamics

We retain five unpaired marginals from synthetic free-energy dynamics on eleven graph families. Models start from the training empirical initial distribution and propagate without resets; errors against independent test histograms are averaged over the four remaining observed times. This evaluates observed-time reconstruction. With logarithmic mobility, both DAM variants achieve lower mean errors than FEGF in every graph family at the base setting (Figure 2). This advantage largely persists across entropy, graph-size and sample-budget sweeps; histogram and learned-ratio DAM have comparable pooled errors (Appendix D.3).

[2] https://github.com/dariorancati/discrete-gradient-flow

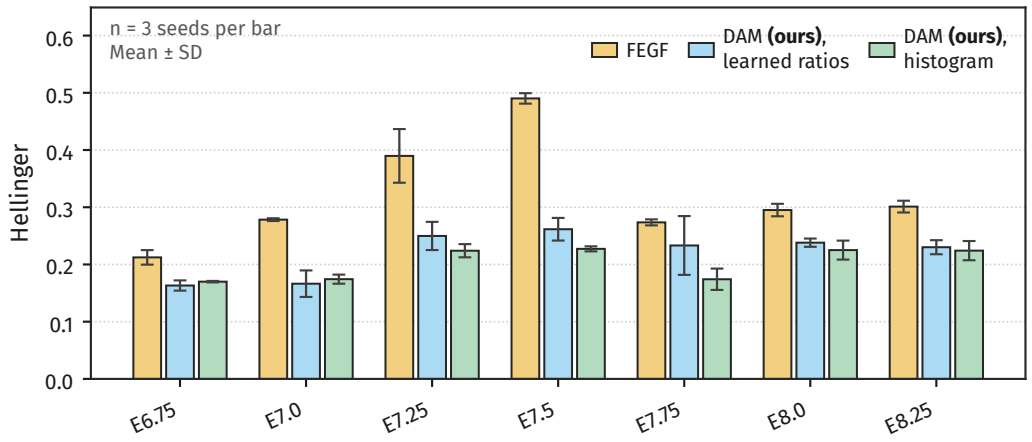


Figure 3: Gastrulation interpolation against omitted smoothed atlas marginals. DAM uses harmonic mobility; lower is better.

5.2 Mouse gastrulation

The gastrulation benchmark uses 34 cell types at nine stages from the atlas of Pijuan-Sala et al. (2019), with a lineage-weighted graph and weak background edges. We omit each interior stage, refit on the other eight (including later observations) and propagate from the preceding retained empirical marginal. The omitted pseudocount-smoothed atlas marginal is scoring-only, so this evaluates interpolation. Both DAM variants achieve lower mean error than FEGF (Figure 3). Harmonic mobility gives the lowest DAM means among four tested choices, without independent mobility selection; all results appear in Table 3. Seed variability reflects resampling and optimization on the same atlas. Comparable accuracy to histogram DAM motivates learned log-ratios as a route to larger state spaces. Appendix D gives full protocols and supplementary mesh experiments.

6 Discussion & Future Work

Limitations. DAM recovers the geometry-selected minimum-action current rather than the historical microscopic dynamics, which are not identifiable from temporal marginals alone (Lavenant et al., 2024; Guan et al., 2024). Likewise, the induced Markov generator is not unique: our smooth positive split provides one trajectory-level realization of the canonical signed current, while other decompositions may produce different microscopic paths with the same marginal evolution.

Toward discrete Wasserstein Lagrangian Flows. A natural extension is a discrete analogue of Wasserstein Lagrangian Flows (Neklyudov et al., 2024). While Action Matching selects the canonical velocity of a prescribed path (Neklyudov et al., 2023), Wasserstein Lagrangian Flows additionally optimize the interpolation between finitely observed marginals. Combining this principle with DAM would require a parameterized marginal path and an outer objective enforcing observation constraints, replacing our predefined empirical interpolant with a variationally selected path. Conditional and unbalanced (Chizat, 2017) extensions for cortical biomarkers (Huang et al., 2024) with changing total burden provide another future direction (see extended discussion in Appendix C.2).

Relation to free-energy models. DAM complements methods such as FEGF (Rancati et al., 2026), which assume that the evolution is generated by a single autonomous free-energy functional. DAM instead allows a time-dependent potential and therefore applies to more general prescribed marginal evolutions. We hope that proposed local density-ratio estimation provides a route for transferring further Wasserstein variational methods from continuous to discrete population dynamics.

AI USE STATEMENT

Generative AI tools, including Codex and Claude, assisted with literature searches, code development and debugging, mathematical derivations, and manuscript drafting and editing (including teaser image generation). The authors reviewed and verified the AI-assisted content and take full responsibility for the theoretical claims, implementation, experimental results, and final manuscript.

ETHICS STATEMENT

This work develops methods for learning population dynamics from temporal snapshots on finite state spaces. The reported experiments use synthetic data, existing mouse gastrulation atlas data, and samples generated from numerical transport on surface meshes. No human participants were recruited and no new biological data were collected for this study. The inferred dynamics are model-dependent descriptions of population change and do not establish biological causation or clinical validity. Cortical prediction is discussed as a future application.

REPRODUCIBILITY STATEMENT

The geometric assumptions and theoretical proofs are provided in Appendix A. Appendices B.3 and D document the training objectives, empirical interpolation, support treatment, dataset preprocessing, baseline adaptations, configurations, evaluation metrics, and numerical checks.

REFERENCES

Shun-ichi Amari. *Information Geometry and Its Applications*. Springer Japan, 2016. ISBN 9784431559788. doi: 10.1007/978-4-431-55978-8. URL http://dx.doi.org/10.1007/978-4-431-55978-8.

Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2000.

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.

William J. Anderson. *Continuous-Time Markov Chains*. Springer New York, 1991. ISBN 9781461230380. doi: 10.1007/978-1-4612-3038-0. URL http://dx.doi.org/10.1007/978-1-4612-3038-0.

Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993, 2021.

Nihat Ay, Jürgen Jost, Hông Vân Lê, and Lorenz Schwachhöfer. Information geometry and sufficient statistics. *Probability Theory and Related Fields*, 162(1-2):327–364, July 2014. ISSN 1432-2064. doi: 10.1007/s00440-014-0574-8. URL http://dx.doi.org/10.1007/s00440-014-0574-8.

Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

Jules Berman, Tobias Blickhan, and Benjamin Peherstorfer. Parametric model reduction of mean-field and stochastic systems via higher-order action matching. *Advances in Neural Information Processing Systems*, 37:56588–56618, 2024.

Tobias Blickhan, Jules Berman, Andrew Stuart, and Benjamin Peherstorfer. Dice: Discrete inverse continuity equation for learning population dynamics. *arXiv preprint arXiv:2507.05107*, 2025.

Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.

Charlotte Bunne, Laetitia Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal optimal transport modeling of population dynamics. In *International Conference on Artificial Intelligence and Statistics*, pp. 6511–6528. PMLR, 2022.

Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.

Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. In *Forty-first International Conference on Machine Learning*, 2024.

Lenaic Chizat. *Unbalanced optimal transport: Models, numerical methods, applications*. PhD thesis, Université Paris sciences et lettres, 2017.

Lénaïc Chizat, Stephen Zhang, Matthieu Heitz, and Geoffrey Schiebinger. Trajectory inference via mean-field langevin in path space. *Advances in Neural Information Processing Systems*, 35: 16731–16742, 2022.

Shui-Nee Chow, Wen Huang, Yao Li, and Haomin Zhou. Fokker–planck equations for a free energy functional or markov process on a graph. *Archive for Rational Mechanics and Analysis*, 203(3): 969–1008, 2012.

Shui-Nee Chow, Wuchen Li, and Haomin Zhou. Entropy dissipation of fokker-planck equations on graphs. *Discrete & Continuous Dynamical Systems-A*, 38(10):4929–4950, 2018.

Shui-Nee Chow, Wuchen Li, and Haomin Zhou. A discrete schrödinger equation via optimal transport on graphs. *Journal of Functional Analysis*, 276(8):2440–2469, 2019.

Shui-Nee Chow, Wuchen Li, and Haomin Zhou. Wasserstein hamiltonian flows. *Journal of Differential Equations*, 268(3):1205–1219, 2020.

Tingting Dan, Jiaqi Ding, Ziquan Wei, Shahar Kovalsky, Minjeong Kim, Won Hwa Kim, and Guorong Wu. Re-think and re-design graph neural networks in spaces of continuous graph diffusion functionals. *Advances in Neural Information Processing Systems*, 36:59375–59387, 2023.

Oscar Davis, Samuel Kessler, Mircea Petrache, İsmail İ Ceylan, Michael Bronstein, and Avishek J Bose. Fisher flow matching for generative modeling over discrete data. *Advances in Neural Information Processing Systems*, 37:139054–139084, 2024.

A Philip Dawid, Steffen Lauritzen, and Matthew Parry. Proper local scoring rules on discrete sample spaces. *The Annals of Statistics*, pp. 593–608, 2012.

Gilberto Diaz-Garcia, Cesar Uribe, and Nicanor Quijano. Population dynamics for discrete wasserstein gradient flows over networks. In *LatinX in AI at International Conference on Machine Learning 2021*, LXAI at ICML 2021. Journal of LatinX in AI Research, July 2021. doi: 10.52591/202107241. URL http://dx.doi.org/10.52591/202107241.

Weinan E, Tiejun Li, and Eric Vanden-Eijnden. *Applied Stochastic Analysis*. American Mathematical Society, May 2019. ISBN 9781470476571. doi: 10.1090/gsm/199. URL http://dx.doi.org/10.1090/gsm/199.

Matthias Erbar. Gradient flows of the entropy for jump processes. In *Annales de l'IHP Probabilités et statistiques*, volume 50, pp. 920–945, 2014.

Matthias Erbar and Jan Maas. Ricci curvature of finite markov chains via convexity of the entropy. *Archive for Rational Mechanics and Analysis*, 206(3):997–1038, 2012.

Matthias Erbar, Jan Maas, and Melchior Wirth. On the geometry of geodesics in discrete optimal transport: M. erbar et al. *Calculus of variations and partial differential equations*, 58(1):19, 2019.

Matthias Erbar, Martin Rumpf, Bernhard Schmitzer, and Stefan Simon. Computation of optimal transport on discrete metric measure spaces: M. erbar et al. *Numerische Mathematik*, 144(1): 157–200, 2020.

Montacer Essid and Justin Solomon. Quadratically regularized optimal transport on graphs. *SIAM Journal on Scientific Computing*, 40(4):A1961–A1986, 2018.

Alessio Figalli and Federico Glaudo. *An invitation to optimal transport, Wasserstein distances, and gradient flows*, volume 1. EMS press Berlin, 2021.

Peyré Gabriel and Cuturi Marco. Computational optimal transport with applications to data sciences. *Foundations and trends® in machine learning*, 11(5-6):355–607, 2019.

Wilfrid Gangbo, Wuchen Li, and Chenchen Mou. Geodesics of minimal length in the set of probability measures on graphs. *ESAIM: Control, Optimisation and Calculus of Variations*, 25:78, 2019.

Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching. *Advances in Neural Information Processing Systems*, 37: 133345–133385, 2024.

David Gentile and James M Murphy. Static and dynamic approaches to computing barycenters of probability measures on graphs. *arXiv preprint arXiv:2603.26940*, 2026.

Vincent Guan, Joseph Janssen, Hossein Rahmani, Andrew Warren, Stephen Zhang, Elina Robeva, and Geoffrey Schiebinger. Identifying drift, diffusion, and causal structure from temporal snapshots. *arXiv preprint arXiv:2410.22729*, 2024.

Vincent Guan, Lazar Atanackovic, and Kirill Neklyudov. A call to lagrangian action: Learning population mechanics from temporal snapshots. In *Forty-third International Conference on Machine Learning*, 2026a. URL https://openreview.net/forum?id=MjPv4kRu3H.

Vincent Guan, Joseph Janssen, Nicolas Lanzetti, Antonio Terpin, Geoffrey Schiebinger, and Elina Robeva. Gradient-flow SDEs have unique transient population dynamics. In *The 29th International Conference on Artificial Intelligence and Statistics*, 2026b. URL https://openreview.net/forum?id=gfF0rwtRMp.

Søren Hauberg. Differential geometry for generative modeling. Tutorial at the Asian Conference on Machine Learning (ACML), 2021. URL http://acml-conf.org.

Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in neural information processing systems*, 34:12454–12465, 2021.

Yanquan Huang, Tingting Dan, Won Hwa Kim, and Guorong Wu. Uncovering cortical pathways of prion-like pathology spreading in alzheimer's disease by neural optimal mass transport. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 498–508. Springer, 2024.

Aapo Hyvärinen. Some extensions of score matching. *Computational statistics & data analysis*, 51 (5):2499–2512, 2007.

Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4):695, 2005.

Takafumi Kanamori and Takashi Takenouchi. Graph-based composite local bregman divergences on discrete sample spaces. *Neural Networks*, 95:44–56, 2017.

Leonid V Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pp. 199–201, 1942.

F.P. Kelly. *Reversibility and Stochastic Networks*. Cambridge Mathematical Library. Cambridge University Press, 2011. ISBN 9781107401150. URL https://books.google.ru/books?id=28kG73yXwjkC.

William Ogilvy Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.

Nikita Kornilov and Alexander Korotin. On the equivalence of optimal transport problem and action matching with optimal vector fields. *Russian Mathematical Surveys*, 80(5):925–927, 2025.

Nikita Maksimovich Kornilov, Petr Mokrov, Alexander Gasnikov, and Alexander Korotin. Optimal flow matching: Learning straight trajectories in just one step. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

John D Lafferty. The density manifold and configuration space quantization. *Transactions of the American Mathematical Society*, 305(2):699–741, 1988.

Hugo Lavenant, Sebastian Claici, Edward Chien, and Justin Solomon. Dynamical optimal transport on discrete surfaces. *ACM Transactions on Graphics (TOG)*, 37(6):1–16, 2018.

Hugo Lavenant, Stephen Zhang, Young-Heon Kim, and Geoffrey Schiebinger. Toward a mathematical theory of trajectory inference. *The Annals of Applied Probability*, 34(1A), February 2024. ISSN 1050-5164. doi: 10.1214/23-aap1969. URL http://dx.doi.org/10.1214/23-AAP1969.

John M. Lee. *Introduction to Riemannian Manifolds*. Springer International Publishing, 2018. ISBN 9783319917559. doi: 10.1007/978-3-319-91755-9. URL http://dx.doi.org/10.1007/978-3-319-91755-9.

Wha Jin Lee, Jesse A Brown, Hye Ryun Kim, Renaud La Joie, Hanna Cho, Chul Hyoung Lyoo, Gil D Rabinovici, Joon-Kyung Seong, and William W Seeley. Regional aβ-tau interactions promote onset and acceleration of alzheimer's disease tau spreading. *Neuron*, 110(12):1932–1943, 2022.

Wuchen Li. *A study of stochastic differential equations and Fokker-Planck equations with applications*. PhD thesis, Georgia Institute of Technology, 2016.

Wuchen Li. Transport information geometry: Riemannian calculus on probability simplex. *Information Geometry*, 5(1):161–207, 2022.

Wuchen Li. Geometric calculations on probability manifolds from reciprocal relations in master equations. *arXiv preprint arXiv:2504.19368*, 2025.

Wuchen Li and Guido Montúfar. Natural gradient via optimal transport. *Information Geometry*, 1 (2):181–214, 2018.

John Lott. Some geometric calculations on wasserstein space. *Commun. Math. Phys.*, 277(2):423–437, November 2007.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 32819–32848, 2024.

Jan Maas. Gradient flows of the entropy for finite markov chains. *Journal of Functional Analysis*, 261(8):2250–2292, 2011.

Michele Mascherpa, Axel Ringh, Amirhossein Taghvaei, and Johan Karlsson. A convex approach for markov chain estimation from aggregate data via inverse optimal transport. *European Journal of Control*, pp. 101592, 2026.

Chenlin Meng, Kristy Choi, Jiaming Song, and Stefano Ermon. Concrete score matching: Generalized score matching for discrete data. *Advances in Neural Information Processing Systems*, 35: 34532–34545, 2022.

Panayotis Mertikopoulos and William H Sandholm. Riemannian game dynamics. *Journal of Economic Theory*, 177:315–364, 2018.

Peter W Michor. *Topics in differential geometry*, volume 93. American Mathematical Soc., 2008.

Alexander Mielke. A gradient structure for reaction–diffusion systems and for energy-drift-diffusion systems. *Nonlinearity*, 24(4):1329–1346, 2011.

Eliakim H Moore. On the reciprocal of the general algebraic matrix. *Bulletin of the american mathematical society*, 26:294–295, 1920.

Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action matching: Learning stochastic dynamics from samples. In *International conference on machine learning*, pp. 25858–25889. PMLR, 2023.

Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A computational framework for solving Wasserstein lagrangian flows. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 37461–37485. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/neklyudov24a.html.

F Otto. The geometry of dissipative evolution equations: the porous medium equation. *Comm. Partial Differential Equations*, 26:101–174, 2001.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Vincent Pauline, Tobias Höppe, Kirill Neklyudov, Alexander Tong, Stefan Bauer, and Andrea Dittadi. Foundations of diffusion models in general state spaces: A self-contained introduction. *arXiv preprint arXiv:2512.05092*, 2025.

Roger Penrose. A generalized inverse for matrices. In *Mathematical proceedings of the Cambridge philosophical society*, volume 51, pp. 406–413. Cambridge University Press, 1955.

Mikhail Persiianov, Jiawei Chen, Petr Mokrov, Alexander Tyurin, Evgeny Burnaev, and Aleksandr Korotin. Learning of population dynamics: Inverse optimization meets jko scheme. In *International Conference on Learning Representations*, volume 2026, pp. 113434–113473, 2026.

P. Petersen. *Riemannian Geometry*. Graduate Texts in Mathematics. Springer New York, 2006. ISBN 9780387294032. URL https://books.google.ru/books?id=9cekXdo52hEC.

Gabriel Peyré. Entropic approximation of wasserstein gradient flows. *SIAM Journal on Imaging Sciences*, 8(4):2323–2351, 2015.

Gabriel Peyré. Optimal transport for machine learners. *arXiv preprint arXiv:2505.06589*, 2025.

Blanca Pijuan-Sala, Jonathan A Griffiths, Carolina Guibentif, Tom W Hiscock, Wajid Jawaid, Fernando J Calero-Nieto, Carla Mulas, Ximena Ibarra-Soria, Richard CV Tyser, Debbie Lee Lian Ho, et al. A single-cell molecular map of mouse gastrulation and early organogenesis. *Nature*, 566(7745):490–495, 2019.

Giovanni Pistone. Information geometry of the probability simplex: A short course. *Nonlinear Phenomena in Complex Systems*, 23(2):221–242, July 2020. ISSN 1561-4085. doi: 10.33581/1561-4085-2020-23-2-221-242. URL http://dx.doi.org/10.33581/1561-4085-2020-23-2-221-242.

Dario Rancati, Jan Maas, and Francesco Locatello. Learning discrete diffusion on graphs via free-energy gradient flows. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=VaQZwmnPGv.

Filippo Santambrogio. *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*. Springer International Publishing, 2015. ISBN 9783319208282. doi: 10.1007/978-3-319-20828-2. URL http://dx.doi.org/10.1007/978-3-319-20828-2.

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.

Gunter Schwarz. *Hodge decomposition - A method for solving boundary value problems*. Lecture Notes in Mathematics. Springer, Berlin, Germany, 1995 edition, July 1995.

Neta Shaul, Itai Gat, Marton Havasi, Daniel Severo, Anuroop Sriram, Peter Holderrieth, Brian Karrer, Yaron Lipman, and Ricky TQ Chen. Flow matching with general discrete paths: A kinetic-optimal perspective. In *International Conference on Learning Representations*, volume 2025, pp. 16397–16429, 2025.

Hannes Stark, Bowen Jing, Chenyu Wang, Gabriele Corso, Bonnie Berger, Regina Barzilay, and Tommi Jaakkola. Dirichlet flow matching with applications to dna sequence design. In *International Conference on Machine Learning*, pp. 46495–46513. PMLR, 2024.

Haoran Sun, Hanjun Dai, Bo Dai, Haomin Zhou, and Dale Schuurmans. Discrete langevin samplers via wasserstein gradient flow. In *International Conference on Artificial Intelligence and Statistics*, pp. 6290–6313. PMLR, 2023a.

Haoran Sun, Lijun Yu, Bo Dai, Dale Schuurmans, and Hanjun Dai. Score-based continuous-time discrete diffusion models. In *The Eleventh International Conference on Learning Representations*, 2023b. URL https://openreview.net/forum?id=BYWWwSY2G5s.

Antonio Terpin, Nicolas Lanzetti, Martín Gadea, and Florian Dörfler. Learning diffusion at lightspeed. *Advances in Neural Information Processing Systems*, 37:6797–6832, 2024.

Panagiotis Theodoropoulos, Juno Nam, Evangelos Theodorou, and Jaemoo Choi. Generalized schrödinger bridge on graphs. In *Forty-third International Conference on Machine Learning*, 2026.

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*, pp. 9526–9536. PMLR, 2020.

Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

Jacob W Vogel, Yasser Iturria-Medina, Olof T Strandberg, Ruben Smith, Elizabeth Levitis, Alan C Evans, and Oskar Hansson. Spread of pathological tau proteins through communicating neurons in human alzheimer's disease. *Nature communications*, 11(1):2612, 2020.

Bohan Zhou, Shu Liu, Xinzhe Zuo, and Wuchen Li. Accelerated markov chain monte carlo algorithms on discrete states. *arXiv preprint arXiv:2505.12599*, 2025.

CONTENTS

A Continuous and discrete Wasserstein geometry **16**

A.1 Optimal transport and Action Matching on continuous sample spaces 17

A.2 Dynamical optimal transport on discrete sample spaces 19

A.3 Wasserstein geometry on the probability simplex 22

B Discrete scores and local density ratios **25**

B.1 Local density ratios and discrete scores 26

B.2 Density ratios in Discrete Action Matching 28

B.3 Empirical training and Markov realization 30

C Extended discussions **32**

C.1 Extended related work 32

C.2 Cortical prediction from unpaired observations 35

D Additional experimental details **35**

D.1 Finite-snapshot synthetic protocol 35

D.2 Training and baseline specification 36

D.3 Synthetic results and ablations 37

D.4 Mouse gastrulation 37

D.5 Sphere and hand surface transport 39

D.6 Surface teacher, sampling and learning protocol 39

D.7 Mesh estimator and integration diagnostics 40

D.8 Accuracy versus computational cost 40

A CONTINUOUS AND DISCRETE WASSERSTEIN GEOMETRY

This section provides the geometric foundations used in Sections 2 and 3. We follow the dynamical viewpoint on optimal transport, in which transport is formulated as a minimum-action problem over probability paths (Benamou & Brenier, 2000; Otto, 2001; Ambrosio et al., 2008). In continuous spaces, this viewpoint yields the formal Riemannian structure of the 2-Wasserstein space associated with Otto calculus (Otto, 2001; Lott, 2007). On finite state spaces, graph gradients, mobility-weighted continuity equations, and kinetic energies give an analogous geometry on the relative interior of the probability simplex (Maas, 2011; Chow et al., 2012; Erbar, 2014). Related finite-dimensional differential-geometric formulations are developed by Li & Montúfar (2018); Li (2022; 2025).

Our presentation is organized around the distinction between a *global transport problem*, which optimizes the probability path itself, and a *fixed-path tangent problem*, which selects a canonical velocity for a prescribed marginal evolution. The latter is the geometric principle underlying Action Matching (Neklyudov et al., 2023). We first review this distinction in continuous spaces. We then introduce its graph counterpart for a general admissible mobility θ, and finally make the induced Riemannian structure on the probability simplex explicit through a mobility-weighted graph Laplacian. The latter representation also gives the existence, uniqueness, and minimum-action characterization of the canonical discrete tangent used in the main text.

For comprehensive treatments of optimal transport, we refer to Villani et al. (2009); Santambrogio (2015); Ambrosio et al. (2008); computational aspects are reviewed by Gabriel & Marco (2019).

Accessible introductions include Figalli & Glaudo (2021); Peyré (2025). Our continuous calculations use standard notions from Riemannian geometry; see, e.g., Petersen (2006); Michor (2008); Lee (2018), or Hauberg (2021) for a machine-learning-oriented introduction.

A.1 Optimal transport and Action Matching on continuous sample spaces

Let (Ω, g^Ω) be a finite-dimensional smooth Riemannian manifold with volume form dx. For every $x \in \Omega$, the Riemannian metric $g^\Omega_x : T_x\Omega \times T_x\Omega \longrightarrow \mathbb{R}$ is an inner product on the tangent space $T_x\Omega$; we write

$$\|v\|^2_{g^\Omega_x} := g^\Omega_x(v, v).$$

We denote by $\mathcal{P}_2(\Omega)$ the Borel probability measures on Ω with finite second moment. Whenever a measure is absolutely continuous with respect to dx, we identify it with its density. For the displayed differential calculations, take a compact connected smooth Riemannian domain, $q \in C^1([0,1]; C^2(\Omega))$ uniformly positive, smooth spatial vector fields, and candidates $s \in C^1([0,1]; C^2(\Omega))$. Use periodic conditions or zero normal probability flux; assume the canonical potential has the same regularity whenever its derivatives are invoked. These are sufficient formal assumptions, not a general theorem for every absolutely continuous Wasserstein curve.

Benamou–Brenier and potential velocities. Given two endpoint distributions $q^0, q^1 \in \mathcal{P}_2(\Omega)$, the dynamical formulation of optimal transport writes the quadratic Wasserstein distance as

$$\frac{1}{2}W_2^2(q^0, q^1) = \inf_{(q_t, v_t)} \frac{1}{2}\int_0^1 \int_\Omega \|v_t(x)\|^2_{g^\Omega_x} q_t(x)\, dx\, dt,$$

where the infimum is over probability paths $t \mapsto q_t$ and velocity fields v_t satisfying the continuity equation

$$\partial_t q_t + \operatorname{div}(q_t v_t) = 0, \qquad q_0 = q^0, \qquad q_1 = q^1. \tag{22}$$

This formulation originates with Benamou & Brenier (2000); its relation to absolutely continuous curves in Wasserstein space is developed systematically in Ambrosio et al. (2008). Equation equation 22 expresses conservation of probability mass under the velocity v_t; see also Figalli & Glaudo (2021, Lemma 4.1.1).

The potential form of the optimal velocity can be seen directly from the constrained problem, as reviewed in Neklyudov et al. (2023, Appendix B). Introduce a Lagrange multiplier $s_t : \Omega \to \mathbb{R}$ for equation 22. Ignoring endpoint terms that do not depend on v, the velocity-dependent part of the augmented functional is

$$\int_0^1 \int_\Omega \left[\frac{1}{2}\|v_t\|^2_{g^\Omega} - g^\Omega_x(\nabla s_t, v_t)\right] q_t\, dx\, dt,$$

where we used the Riemannian integration-by-parts identity

$$\int_\Omega s\, \operatorname{div}(qv)\, dx = -\int_\Omega q\, g^\Omega_x(\nabla s, v)\, dx$$

under the stated boundary assumptions (Lee, 2018). Completing the square,

$$\frac{1}{2}\|v\|^2_{g^\Omega} - g^\Omega(\nabla s, v) = \frac{1}{2}\|v - \nabla s\|^2_{g^\Omega} - \frac{1}{2}\|\nabla s\|^2_{g^\Omega},$$

shows that the minimizing velocity is $v_t = \nabla s_t$. The corresponding continuity equation is therefore

$$\partial_t q_t = -\operatorname{div}\left(q_t \nabla s_t\right).$$

Only differences of the potential matter, so s_t is defined up to an additive time-dependent constant.

The fixed-path tangent problem. The Benamou–Brenier problem simultaneously chooses a probability path and a velocity field. Action Matching instead starts from a prescribed marginal path and asks which velocity should represent its instantaneous motion (Neklyudov et al., 2023). This distinction is most transparent through the local Wasserstein tangent problem.

Formally, consider the space of smooth strictly positive densities $\mathcal{P}_+(\Omega) := \left\{q \in C^\infty(\Omega) : q > 0, \quad \int_\Omega q\, dx = 1\right\}$. Its tangent space consists of zero-mass perturbations,

$T_q\mathcal{P}_+(\Omega) = \left\{\sigma : \int_\Omega \sigma(x)\,dx = 0\right\}$. This density-manifold viewpoint is classical in information geometry and optimal transport (Lafferty, 1988; Otto, 2001; Lott, 2007).

For a prescribed tangent $\sigma \in T_q\mathcal{P}_+(\Omega)$, the equation $\sigma = -\operatorname{div}(qv)$ does not determine an arbitrary velocity v uniquely. If w satisfies $\operatorname{div}(qw) = 0$, then $v + w$ generates the same tangent. In the weighted $L^2(q)$ geometry, this ambiguity is resolved by the weighted Helmholtz–Hodge decomposition: the gradient fields form the orthogonal complement of the q-divergence-free fields (Ambrosio et al., 2008, Lemma 8.4.2); see also Schwarz (1995); Chow et al. (2018). Formally,

$$v = \nabla s + w, \qquad \operatorname{div}(qw) = 0,$$

with $\int_\Omega q\, g_x^\Omega(\nabla s, w)\,dx = 0$. Consequently,

$$\frac{1}{2}\int_\Omega q\,\|v\|_{g^\Omega}^2\,dx = \frac{1}{2}\int_\Omega q\,\|\nabla s\|_{g^\Omega}^2\,dx + \frac{1}{2}\int_\Omega q\,\|w\|_{g^\Omega}^2\,dx. \tag{23}$$

The gradient component is therefore the unique minimum-norm representative of the tangent. This motivates the Wasserstein kinetic energy of σ at q,

$$\mathcal{K}_W(q,\sigma) := \inf_{\substack{v \\ \sigma = -\operatorname{div}(qv)}} \frac{1}{2}\int_\Omega q(x)\|v(x)\|_{g_x^\Omega}^2\,dx.$$

For absolutely continuous Wasserstein curves, the existence and minimum-norm characterization of the canonical velocity are made rigorous in Ambrosio et al. (2008, Theorem 8.3.1 and Section 8.4). The minimizer is the gradient representative

$$\sigma = -\operatorname{div}\left(q\nabla s\right),$$

and

$$\mathcal{K}_W(q,\sigma) = \frac{1}{2}\int_\Omega q(x)\|\nabla s(x)\|_{g_x^\Omega}^2\,dx.$$

Along an absolutely continuous Wasserstein curve, the same canonical velocity also arises infinitesimally from optimal transport between nearby marginals (Ambrosio et al., 2008, Proposition 8.4.6); see also Neklyudov et al. (2023, Appendix B.1).

Wasserstein metric and path action. The map

$$C^\infty(\Omega)/\mathbb{R} \longrightarrow T_q\mathcal{P}_+(\Omega), \qquad [s] \longmapsto -\operatorname{div}(q\nabla s),$$

identifies potentials modulo constants with tangent directions under suitable regularity assumptions (Otto, 2001; Lott, 2007). If

$$\sigma_s = -\operatorname{div}(q\nabla s), \qquad \sigma_r = -\operatorname{div}(q\nabla r),$$

the formal Wasserstein inner product is

$$g_q^W(\sigma_s,\sigma_r) := \int_\Omega q(x)\, g_x^\Omega\left(\nabla s(x), \nabla r(x)\right)dx.$$

Thus $\mathcal{K}_W(q,\sigma) = \frac{1}{2}g_q^W(\sigma,\sigma)$. The family $g^W = \{g_q^W\}_q$ is the formal Riemannian structure underlying Otto calculus (Otto, 2001; Lott, 2007); rigorous formulations use Wasserstein tangent spaces, metric derivatives, and absolutely continuous curves (Ambrosio et al., 2008).

For a smooth path q_t with canonical potential s_t,

$$\dot{q}_t = -\operatorname{div}(q_t\nabla s_t),$$

the corresponding kinetic action is

$$\mathcal{A}_W[q] := \int_0^1 \mathcal{K}_W(q_t,\dot{q}_t)\,dt = \frac{1}{2}\int_0^1\int_\Omega q_t(x)\|\nabla s_t(x)\|_{g_x^\Omega}^2 dx\,dt.$$

The Benamou–Brenier problem can therefore be written as

$$\frac{1}{2}W_2^2(q^0,q^1) = \inf_{\substack{q_0=q^0 \\ q_1=q^1}} \mathcal{A}_W[q]. \tag{24}$$

The local tangent problem selects the canonical velocity of a *given* path, whereas equation 24 additionally minimizes over the path connecting the endpoints.

Fixed-path Action Matching. Action Matching (Neklyudov et al., 2023) keeps the complete marginal path q_t fixed and learns its canonical potential $s_t^\star$. By Neklyudov et al. (2023, Theorem 2.1), for an absolutely continuous Wasserstein path there exists a unique canonical gradient field $\nabla s_t^\star$ satisfying

$$\dot{q}_t = -\operatorname{div}(q_t \nabla s_t^\star), \tag{25}$$

with $s_t^\star$ unique up to an additive function of time.

For a candidate potential s, define the action gap

$$\mathrm{AG}_W(s, s^\star) := \frac{1}{2}\int_0^1 \int_\Omega q_t(x)\, \|\nabla s_t(x) - \nabla s_t^\star(x)\|^2_{g_x^\Omega}\, dx\, dt.$$

Expanding the square gives the cross term $\int_\Omega q_t g^\Omega(\nabla s_t, \nabla s_t^\star)\, dx$. Using spatial integration by parts and equation 25,

$$\int_\Omega q_t g^\Omega(\nabla s_t, \nabla s_t^\star)\, dx = \int_\Omega s_t\, \dot{q}_t\, dx.$$

Integration by parts in time then yields

$$\int_0^1 \int_\Omega s_t\, \dot{q}_t\, dx\, dt = \int_\Omega s_1 q_1\, dx - \int_\Omega s_0 q_0\, dx - \int_0^1 \int_\Omega q_t\, \partial_t s_t\, dx\, dt. \tag{26}$$

Substituting equation 26 into the expanded action gap gives $\mathrm{AG}_W(s, s^\star) = \mathcal{L}_{\mathrm{AM}}(s) - \mathcal{L}_{\mathrm{AM}}(s^\star)$, where

$$\begin{aligned}\mathcal{L}_{\mathrm{AM}}(s) :=& \mathbb{E}_{X\sim q_0}[s_0(X)] - \mathbb{E}_{X\sim q_1}[s_1(X)] \\ &+ \int_0^1 \mathbb{E}_{X\sim q_t}\left[\partial_t s_t(X) + \frac{1}{2}\|\nabla s_t(X)\|^2_{g_X^\Omega}\right] dt.\end{aligned} \tag{27}$$

This is the Action Matching identity of Neklyudov et al. (2023, Theorem 2.2). Its significance for learning from snapshots is that the prescribed path enters equation 27 only through expectations under its marginals. The discrete construction below follows the same fixed-path principle, but its kinetic energy contains an additional density-dependent edge mobility.

A.2 Dynamical optimal transport on discrete sample spaces

We now replace the continuous state space by a finite set $I = \{1, \dots, N\}$ equipped with a connected undirected weighted graph $G = (I, E, \omega)$, where E is the set of unoriented edges and

$$\omega_{ij} = \omega_{ji} > 0, \qquad \{i, j\} \in E.$$

The graph specifies which pairs of states may exchange probability locally. Such graph-based differential structures underlie several formulations of discrete dynamical optimal transport (Maas, 2011; Chow et al., 2012; Mielke, 2011; Li & Montúfar, 2018). We write $\mathcal{N}(i) := \{j \in I : \{i, j\} \in E\}$ for the neighborhood of i. The probability simplex and its relative interior are

$$\mathcal{P}(I) := \left\{p \in \mathbb{R}^N : p_i \geq 0,\ \mathbf{1}^\top p = 1\right\}, \tag{28}$$

$$\mathcal{P}_+(I) := \{p \in \mathcal{P}(I) : p_i > 0,\ i \in I\}. \tag{29}$$

Graph gradient and divergence. A graph vector field is a skew-symmetric edge field

$$v_{ij} = -v_{ji}, \qquad \{i, j\} \in E.$$

For a scalar potential $S \in \mathbb{R}^N$, define the graph gradient

$$(\nabla_G S)_{ij} := \sqrt{\omega_{ij}}(S_i - S_j).$$

This is the weighted finite-difference convention used in graph formulations of discrete Wasserstein geometry (Chow et al., 2012; Li & Montúfar, 2018; Li, 2022). The graph divergence is

$$(\operatorname{div}_G v)_i := \sum_{j\in\mathcal{N}(i)} \sqrt{\omega_{ij}}\, v_{ji}.$$

With these signs, $-\operatorname{div}_G$ is the adjoint of ∇_G. Indeed, for any potential S and skew-symmetric edge field v,

$$\begin{aligned} S^\top[-\operatorname{div}_G v] &= -\sum_{i\in I} S_i \sum_{j\in\mathcal{N}(i)} \sqrt{\omega_{ij}}\, v_{ji} \\ &= \frac{1}{2}\sum_{i\in I}\sum_{j\in\mathcal{N}(i)} \sqrt{\omega_{ij}}(S_i - S_j)v_{ij} \\ &= \sum_{\{i,j\}\in E} (\nabla_G S)_{ij} v_{ij}. \end{aligned} \tag{30}$$

The factor $1/2$ in the directed double sum compensates for counting every unoriented edge twice. This identity is the finite-graph analogue of integration by parts and appears throughout discrete transport calculus (Chow et al., 2012; 2019; Rancati et al., 2026).

Mobility-weighted transport. In the continuous continuity equation, the pointwise density multiplies a velocity defined at the same spatial location. On a graph, probability masses live on vertices while velocities live on edges. A discrete transport geometry must therefore specify how the two endpoint masses determine the amount of probability available for transport across an edge. This role is played by a mobility or mean (Maas, 2011; Mielke, 2011; Erbar et al., 2019).

Definition A.1 (Admissible mobility)**.** We adopt the following sufficient common assumptions, rather than identify them with every discrete transport framework: $\theta : [0,\infty)^2 \to [0,\infty)$ is continuous, smooth and strictly positive on $(0,\infty)^2$, symmetric, nondecreasing in each argument, jointly concave, normalized by $\theta(1,1) = 1$, and positively homogeneous:

$$\theta(\lambda a, \lambda b) = \lambda\theta(a,b), \qquad \lambda > 0.$$

We do not require $\theta(0,b) = 0$, unlike the zero-boundary class in Maas (2011, Assumption 3.1); arithmetic mobility is therefore allowed. Connectedness and positive conductances suffice for the tangent representation; temporal regularity supports the action identity; homogeneity enables ratio factorization. The softplus minimizer itself does not depend on the mobility.

The normalization and homogeneity imply $\theta(a,a) = a$. Standard examples include the arithmetic mean

$$\theta^{\mathrm{ari}}(a,b) = \frac{a+b}{2},$$

the logarithmic mean

$$\theta^{\log}(a,b) = \begin{cases} \dfrac{a-b}{\log a - \log b}, & a \neq b, \\ a, & a = b, \end{cases}$$

and the geometric mean

$$\theta^{\mathrm{geo}}(a,b) = \sqrt{ab}.$$

We also use the harmonic mean $\theta^{\mathrm{har}}(a,b) = 2ab/(a+b)$, extended by zero when either argument vanishes. Different choices induce different discrete transport geometries. In the stationary-relative reversible-chain construction, the logarithmic mean makes the heat flow a gradient flow of relative entropy (Maas, 2011; Erbar & Maas, 2012).

For $p \in \mathcal{P}_+(I)$, write $\theta_{ij}(p) := \theta(p_i, p_j)$. The mobility-weighted edge pairing is

$$\langle v, \widetilde{v}\rangle_{p,\theta} := \sum_{\{i,j\}\in E} \theta_{ij}(p) v_{ij}\widetilde{v}_{ij}.$$

In particular,

$$|\nabla_G S|^2_{p,\theta} := \sum_{\{i,j\}\in E} \omega_{ij}\theta(p_i,p_j)(S_i - S_j)^2.$$

Taking $v = \theta(p)\nabla_G R$ in equation 30 yields the mobility-weighted integration-by-parts identity

$$\langle \nabla_G S, \nabla_G R\rangle_{p,\theta} = S^\top\left[-\operatorname{div}_G\left(\theta(p)\nabla_G R\right)\right]. \tag{31}$$

This is the discrete counterpart of

$$\int q\, g^{\Omega}(\nabla s, \nabla r) = \int s[-\operatorname{div}(q\nabla r)].$$

A differentiable probability path $t \mapsto p(t)$ generated by a potential $S(t)$ then satisfies the graph continuity equation

$$\dot{p}(t) + \operatorname{div}_G\left(\theta(p(t))\nabla_G S(t)\right) = 0. \tag{32}$$

In coordinates,

$$\dot{p}_i(t) = \sum_{j\in\mathcal{N}(i)} \omega_{ij}\theta\left(p_i(t), p_j(t)\right)\left(S_i(t) - S_j(t)\right).$$

Graph continuity equations of this form are central to the dynamical transport constructions of Maas (2011); Chow et al. (2012); Mielke (2011). Compared with the continuous equation $\dot{q} = -\operatorname{div}(q\nabla s)$, the pointwise density is replaced by an edge mobility because the state and velocity variables live on vertices and edges, respectively.

Kinetic action and discrete dynamical transport. The instantaneous kinetic energy associated with p and S is

$$\mathcal{K}_{G,\theta}(p, S) := \frac{1}{2}\left|\nabla_G S\right|^2_{p,\theta} = \frac{1}{2}\sum_{\{i,j\}\in E} \omega_{ij}\theta(p_i, p_j)(S_i - S_j)^2. \tag{33}$$

For a path satisfying equation 32, its kinetic action is

$$\mathcal{A}_{G,\theta}[p, S] := \int_0^1 \mathcal{K}_{G,\theta}\left(p(t), S(t)\right) dt.$$

For positive endpoints p^0, p^1, define the interior-path distance using absolutely continuous $p : [0,1] \to \mathcal{P}_+(I)$ and measurable potentials of finite kinetic action satisfying the continuity equation almost everywhere. The graph analogue of the Benamou–Brenier problem is then

$$\frac{1}{2}W^2_{G,\theta}(p^0, p^1) := \inf_{\substack{p,S \\ \dot{p}+\operatorname{div}_G(\theta(p)\nabla_G S)=0 \\ p(0)=p^0,\ p(1)=p^1}} \mathcal{A}_{G,\theta}[p, S]. \tag{34}$$

Discrete dynamical transport distances of this type were introduced and studied in various closely related forms by Maas (2011); Chow et al. (2012); Mielke (2011); the precise global behavior near the boundary of the simplex depends on the mobility. Our method uses the induced geometry on $\mathcal{P}_+(I)$, where all edge mobilities are strictly positive.

Remark A.2 (Relation to static Kantorovich transport)**.** The distance in equation 34 should not be confused with applying the classical Kantorovich problem directly to the finite metric space (I, d^G). The latter gives

$$W^2_{\mathrm{K}}(p^0, p^1) := \min_{\pi\in\Pi(p^0,p^1)} \sum_{i,j\in I} d^G(i,j)^2\pi_{ij},$$

where $\Pi(p^0, p^1)$ is the set of couplings with the prescribed marginals (Kantorovich, 1942; Villani et al., 2009). In general,

$$W_{G,\theta} \neq W_{\mathrm{K}}.$$

The static Kantorovich problem is determined by the ground cost and a coupling of the endpoints, whereas the dynamical graph construction additionally specifies how probability moves infinitesimally through the graph and depends on the choice of mobility θ (Maas, 2011; Erbar, 2014). It is this latter infinitesimal geometry that is required for the fixed-path Action Matching construction.

The continuity equation equation 32 and kinetic energy equation 33 already contain the two ingredients needed for Action Matching. We next make the resulting tangent-space geometry explicit and show that every tangent direction has a unique minimum-energy gradient representative.

A.3 WASSERSTEIN GEOMETRY ON THE PROBABILITY SIMPLEX

The relative interior $\mathcal{P}_+(I)$ is an $(N-1)$-dimensional smooth manifold with tangent space (Amari & Nagaoka, 2000; Amari, 2016; Ay et al., 2014; Pistone, 2020).

$$T_p\mathcal{P}_+(I) = \left\{\sigma \in \mathbb{R}^N : \mathbf{1}^\top \sigma = 0\right\}.$$

Discrete dynamical transport equips this manifold with a density-dependent, mobility-dependent Riemannian metric (Maas, 2011; Chow et al., 2012; Mielke, 2011; Li, 2022; 2025), rather than the Fisher–Rao geometry commonly considered in information geometry. We derive this transport metric directly for the general mobility θ used throughout the paper.

Mobility-weighted graph Laplacian. Fix an arbitrary orientation of the unoriented edges and let $D \in \mathbb{R}^{|E|\times N}$ be the weighted incidence matrix

$$D_{e,k} = \begin{cases} \sqrt{\omega_{ij}}, & k = i, \\ -\sqrt{\omega_{ij}}, & k = j, \\ 0, & \text{otherwise}, \end{cases} \qquad e = (i,j).$$

Thus DS is the oriented representation of $\nabla_G S$. For $p \in \mathcal{P}_+(I)$ define the diagonal mobility matrix

$$\Theta_\theta(p) := \operatorname{diag}\left(\theta(p_i,p_j)\right)_{\{i,j\}\in E}$$

and the mobility-weighted graph Laplacian

$$L_\theta(p) := D^\top \Theta_\theta(p) D. \tag{35}$$

The operator does not depend on the arbitrary edge orientation. With the sign conventions of Section A.2,

$$L_\theta(p)S = -\operatorname{div}_G\left(\theta(p)\nabla_G S\right).$$

Hence the graph continuity equation can be written compactly as $\dot{p}(t) = L_\theta(p(t))S(t)$.

For any $S, R \in \mathbb{R}^N$,

$$S^\top L_\theta(p) R = (DS)^\top \Theta_\theta(p)(DR) = \langle \nabla_G S, \nabla_G R\rangle_{p,\theta}. \tag{36}$$

Since $p \in \mathcal{P}_+(I)$ and θ is positive on $(0,\infty)^2$, all diagonal entries of $\Theta_\theta(p)$ are strictly positive. Therefore

$$S^\top L_\theta(p) S = \sum_{\{i,j\}\in E} \omega_{ij}\theta(p_i,p_j)(S_i - S_j)^2 \geq 0.$$

Moreover, $S^\top L_\theta(p)S = 0$ if and only if $S_i = S_j$ on every edge. Since G is connected,

$$\ker L_\theta(p) = \operatorname{span}\{\mathbf{1}\}. \tag{37}$$

Thus $\operatorname{rank} L_\theta(p) = N-1$ and, because $L_\theta(p)\mathbf{1} = 0$,

$$\operatorname{Ran} L_\theta(p) = T_p\mathcal{P}_+(I). \tag{38}$$

This is the finite-dimensional mechanism behind the identification of tangents with discrete gradients in the Maas–Erbar–Chow geometry (Maas, 2011; Chow et al., 2019; Rancati et al., 2026).

Potentials therefore naturally live in the quotient $\mathbb{R}^N/\operatorname{span}\{\mathbf{1}\}$, and $[S] \longmapsto L_\theta(p)S$ is an isomorphism from this quotient to $T_p\mathcal{P}_+(I)$.

Canonical tangent representation and graph Hodge decomposition. We review the canonical tangent representation in equation 11 and its minimum-energy characterization in equation 12. The argument uses the mobility-weighted graph Hodge decomposition (Maas, 2011, Lemma 3.6); see also Chow et al. (2019, Lemma 4).

Derivation of the canonical tangent representation. Let $\sigma \in T_p\mathcal{P}_+(I)$. By equation 38, there exists $S^\star \in \mathbb{R}^N$ such that $L_\theta(p)S^\star = \sigma$, or equivalently,

$$\sigma = -\operatorname{div}_G\left(\theta(p)\nabla_G S^\star\right). \tag{39}$$

If R is another solution, then $L_\theta(p)(R-S^\star)=0$. Equation equation 37 implies $R-S^\star=c\mathbf{1}$, so the potential class $[S^\star]$ and the gradient $\nabla_G S^\star$ are unique. It remains to establish the minimum-energy property. Let v be any skew-symmetric edge field satisfying

$$\sigma = -\operatorname{div}_G(\theta(p)v). \tag{40}$$

Define $u := v - \nabla_G S^\star$. Subtracting equation 39 from equation 40 gives $\operatorname{div}_G(\theta(p)u)=0$. By the mobility-weighted integration-by-parts identity equation 31,

$$\langle \nabla_G S^\star, u\rangle_{p,\theta} = S^{\star\top}[-\operatorname{div}_G(\theta(p)u)] = 0.$$

Therefore

$$\begin{aligned}\frac{1}{2}\|v\|_{p,\theta}^2 &= \frac{1}{2}\|\nabla_G S^\star + u\|_{p,\theta}^2 \\ &= \frac{1}{2}\|\nabla_G S^\star\|_{p,\theta}^2 + \frac{1}{2}\|u\|_{p,\theta}^2 \\ &\geq \frac{1}{2}\|\nabla_G S^\star\|_{p,\theta}^2. \end{aligned} \tag{41}$$

Strict positivity of the mobility on $\mathcal{P}_+(I)$ implies equality only when $u=0$. Hence $\nabla_G S^\star$ is the unique minimum-energy edge velocity inducing σ. □

Equation equation 41 is the discrete analogue of the weighted Hodge decomposition equation 23. In particular, an arbitrary graph velocity can be decomposed as

$$v = \nabla_G S^\star + u, \qquad \operatorname{div}_G(\theta(p)u) = 0,$$

where the second component is invisible to the marginal tangent but contributes additional kinetic energy. This is precisely why the scalar potential in Discrete Action Matching is a geometric selection principle rather than merely a parameterization choice.

Primal representation of the metric. Let $L_\theta(p)^\dagger$ denote the Moore–Penrose pseudoinverse. Since $\ker L_\theta(p) = \operatorname{span}\{\mathbf{1}\}$, the pseudoinverse acts as the inverse of $L_\theta(p)$ on the tangent space. For

$$\sigma = L_\theta(p)S, \qquad \tau = L_\theta(p)R,$$

we have

$$\begin{aligned}\sigma^\top L_\theta(p)^\dagger \tau &= S^\top L_\theta(p) L_\theta(p)^\dagger L_\theta(p) R \\ &= S^\top L_\theta(p) R \\ &= \langle \nabla_G S, \nabla_G R\rangle_{p,\theta}. \end{aligned} \tag{42}$$

Thus the same metric can be written either in potential coordinates,

$$g^G_{p,\theta}(L_\theta(p)S, L_\theta(p)R) := \langle \nabla_G S, \nabla_G R\rangle_{p,\theta},$$

or directly in tangent coordinates,

$$g^G_{p,\theta}(\sigma,\tau) := \sigma^\top L_\theta(p)^\dagger \tau.$$

Because θ is smooth and positive on the interior and $L_\theta(p)$ has constant rank $N-1$ on $\mathcal{P}_+(I)$, this defines a smooth positive-definite inner product on each tangent space. Hence $\mathcal{P}_+(I)$ carries the finite-dimensional Riemannian structure associated with the graph transport geometry; related constructions and their global metric properties are developed in Maas (2011); Mielke (2011); Li (2022; 2025).

The kinetic energy of a tangent vector is therefore

$$\mathcal{T}_{G,\theta}(p,\sigma) = \frac{1}{2} g^G_{p,\theta}(\sigma,\sigma) = \frac{1}{2}\sigma^\top L_\theta(p)^\dagger \sigma,$$

where $\sigma = L_\theta(p)S$. Consequently, the dynamical transport problem equation 34 admits the following primal form over the same interior path class

$$\frac{1}{2}W^2_{G,\theta}(p^0,p^1) = \inf_{\substack{p_t\in\mathcal{P}_+(I)\\ p_0=p^0,\ p_1=p^1}} \frac{1}{2}\int_0^1 \dot p_t^\top L_\theta(p_t)^\dagger \dot p_t\, dt.$$

Thus the graph kinetic action is the Riemannian energy of the probability curve under the metric tensor represented by $L_\theta(p)^\dagger$ on $T_p\mathcal{P}_+(I)$.

Discrete Action Matching identity. The same structure gives the discrete counterpart of the continuous Action Matching derivation. Use the assumptions and zero-mean gauge of Theorem 3.1; write $p_t = p(t)$ and $S_t = S(t)$, so subscripts in this proof denote time, not state. Let $S_t^\star = L_\theta(p_t)^\dagger \dot{p}_t$. Then

$$\dot{p}_t = L_\theta(p_t) S_t^\star = -\operatorname{div}_G\left(\theta(p_t)\nabla_G S_t^\star\right). \tag{43}$$

For a candidate S_t, expand the discrete action gap

$$\mathrm{AG}_{G,\theta}(S, S^\star) = \frac{1}{2}\int_0^1 \Big(\|\nabla_G S_t\|^2_{p_t,\theta} + \|\nabla_G S_t^\star\|^2_{p_t,\theta} - 2\langle \nabla_G S_t, \nabla_G S_t^\star\rangle_{p_t,\theta}\Big) dt. \tag{44}$$

Using equation 31 and equation 43,

$$\langle \nabla_G S_t, \nabla_G S_t^\star\rangle_{p_t,\theta} = S_t^\top \dot{p}_t.$$

Temporal integration by parts gives

$$\int_0^1 S_t^\top \dot{p}_t \, dt = S_1^\top p_1 - S_0^\top p_0 - \int_0^1 p_t^\top \partial_t S_t \, dt.$$

Substituting into equation 44 yields

$$\begin{aligned}\mathrm{AG}_{G,\theta}(S, S^\star) = p_0^\top S_0 - p_1^\top S_1 + \int_0^1 \left[p_t^\top \partial_t S_t + \frac{1}{2}\|\nabla_G S_t\|^2_{p_t,\theta}\right] dt \\ + \frac{1}{2}\int_0^1 \|\nabla_G S_t^\star\|^2_{p_t,\theta} dt.\end{aligned} \tag{45}$$

The last term is independent of S. Evaluating the same expression at $S = S^\star$ shows that $\mathrm{AG}_{G,\theta}(S, S^\star) = \mathcal{L}_{\mathrm{DAM}}(S) - \mathcal{L}_{\mathrm{DAM}}(S^\star)$, which proves Theorem 3.1. This is the exact graph analogue of the fixed-path Action Matching identity of Neklyudov et al. (2023, Theorem 2.2). For nonlinear mobilities, the remaining density dependence is that of $\|\nabla_G S_t\|^2_{p_t,\theta}$ on $\theta(p_i(t), p_j(t))$; Appendix B shows how positive 1-homogeneity reduces this dependence to neighboring density ratios.

Relation to the Maas–Rancati parameterization. Our notation follows the probability-coordinate graph formulation used in the main text. The original construction of Maas (2011), and the recent presentation of Rancati et al. (2026), use a different but closely related parameterization based on a reversible Markov kernel K with invariant probability measure π. A probability measure is represented there by its density ρ relative to π,

$$p_i = \pi_i \rho_i, \qquad \rho_i = \frac{p_i}{\pi_i}.$$

The corresponding symmetric edge conductance has the form

$$c_{ij}^{\mathrm{MR}}(\rho) = \pi_i K_{ij} m(\rho_i, \rho_j),$$

where reversibility, $\pi_i K_{ij} = \pi_j K_{ji}$, makes the conductance symmetric (Maas, 2011; Rancati et al., 2026). Expressed directly in our probability coordinates,

$$c_{ij}^{\mathrm{MR}}(p) = \pi_i K_{ij} m\left(\frac{p_i}{\pi_i}, \frac{p_j}{\pi_j}\right). \tag{46}$$

Our construction instead writes the conductance as

$$c_{ij}(p) = \omega_{ij}\theta(p_i, p_j).$$

The two parameterizations should therefore be compared at the level of the edge conductances $c_{ij}(p) \longleftrightarrow c_{ij}^{\mathrm{MR}}(p)$, rather than by identifying ρ_i with p_i. For a homogeneous symmetric m, reversibility gives $c_{ij}^{\mathrm{MR}}(p) = m(K_{ij}p_i, K_{ji}p_j)$. For the logarithmic mean and nonuniform π, its ratio to $m(p_i, p_j)$ generally depends on p; it cannot generally be absorbed into fixed ω_{ij}. Uniform π permits a constant rescaling, and the geometric mean has the special factorization $c_{ij}^{\mathrm{MR}} = K_{ij}\sqrt{\pi_i/\pi_j}\sqrt{p_i p_j}$. These are sufficient special cases, not an exhaustive classification.

Table 1: Continuous–discrete dictionary for the dynamical transport structures underlying Action Matching.

Role	**Continuous sample space**	**Finite graph**
Probability state	$q \in \mathcal{P}_+(\Omega)$	$p \in \mathcal{P}_+(I)$
Tangent vector	$\sigma \in T_q\mathcal{P}_+(\Omega), \quad \int_\Omega \sigma\, dx = 0$	$\sigma \in T_p\mathcal{P}_+(I), \quad \mathbf{1}^\top \sigma = 0$
Potential	$[s] \in C^\infty(\Omega)/\mathbb{R}$	$[S] \in \mathbb{R}^N / \operatorname{span}\{\mathbf{1}\}$
Gradient	∇s	$(\nabla_G S)_{ij} = \sqrt{\omega_{ij}}(S_i - S_j)$
Potential-to-tangent map	$-\operatorname{div}(q\nabla s)$	$-\operatorname{div}_G(\theta(p)\nabla_G S) = L_\theta(p)S$
Continuity equation	$\dot{q}_t = -\operatorname{div}(q_t \nabla s_t)$	$\dot{p}_t = L_\theta(p_t)S_t$
Kinetic pairing	$\int_\Omega q\, g_x^\Omega(\nabla s, \nabla r)\, dx$	$\sum_{\{i,j\}\in E} \omega_{ij}\theta(p_i,p_j)(S_i - S_j)(R_i - R_j)$
Tangent metric	$g_q^W(\sigma,\tau)$	$g_{p,\theta}^G(\sigma,\tau) = \sigma^\top L_\theta(p)^\dagger \tau$
Kinetic energy	$\frac{1}{2} g_q^W(\sigma,\sigma)$	$\frac{1}{2} g_{p,\theta}^G(\sigma,\sigma)$
Path action	$\int_0^1 \mathcal{K}_W(q_t,\dot{q}_t)\, dt$	$\int_0^1 \mathcal{T}_{G,\theta}(p_t,\dot{p}_t)\, dt$
Invisible component	$\operatorname{div}(qw) = 0$	$\operatorname{div}_G(\theta(p)u) = 0$
Gauge invariance	$s \mapsto s + c$	$S \mapsto S + c\mathbf{1}$

Under the conductance-level interpretation, the unique-gradient statement of Maas (2011, Proposition 3.26), restated as Rancati et al. (2026, Proposition B.6), corresponds to the tangent representation proved above.

Table 1 summarizes the continuous and discrete constructions used throughout the paper.

The correspondence makes the common structure explicit. In both settings, a marginal tangent alone admits many velocity representations. The transport geometry projects this ambiguity onto a unique gradient representative, and its weighted squared norm defines the instantaneous kinetic energy. Continuous Action Matching exploits this structure because multiplication by q_t turns the kinetic term directly into an expectation under the observed marginal. On a graph, the corresponding weight is the edge mobility $\theta(p_i, p_j)$, which depends on two unknown probability masses. The next appendix section develops the local density-ratio representation that removes this remaining obstacle.

B Discrete scores and local density ratios

The Discrete Action Matching objective depends on the current marginal through the edge mobility $\theta(p_i(t), p_j(t))$. Positive 1-homogeneity reduces this dependence to ratios of neighboring probability masses, motivating the local density-ratio estimator introduced in Section 3.2. Such ratios have a longer history in normalization-free estimation on discrete spaces: ratio matching was introduced by Hyvärinen (2007), related constructions were placed within proper homogeneous local scoring rules by Dawid et al. (2012), and graph-based local Bregman formulations were developed by Kanamori & Takenouchi (2017). More recently, neighboring ratios were interpreted as discrete score-like quantities by Meng et al. (2022) and used directly in reverse-time discrete diffusion by Lou et al. (2024).

In DAM, these ratios are auxiliary marginal information rather than the learned population velocity: they determine the mobility-weighted kinetic energy, whereas the action potential S represents the canonical dynamics. This section provides the details omitted from the main text. We first review local density ratios and their relation to discrete scores, then derive their appearance in the DAM kinetic energy and prove consistency of the local softplus estimator. Finally, we discuss its relation to implicit score entropy and give the corresponding empirical training and Markov realization.

B.1 LOCAL DENSITY RATIOS AND DISCRETE SCORES

Let $G = (I, E, \omega)$ be a connected undirected graph and let $p \in \mathcal{P}_+(I)$. For every directed version of an edge $\{i, j\} \in E$, define the *local density ratio*

$$r_p(i, j) := \frac{p_j}{p_i}. \tag{47}$$

Although the underlying graph is undirected, the ratio is directed and satisfies the reciprocity identity $r_p(j, i) = r_p(i, j)^{-1}$. The collection $r_p(i) := (r_p(i, j))_{j \in \mathcal{N}(i)}$ therefore describes the masses of neighboring states relative to the mass at i. A basic property of equation 47 is invariance to global normalization. If $\widetilde{p}_i = cp_i$ for $c > 0$, then

$$\frac{\widetilde{p}_j}{\widetilde{p}_i} = \frac{p_j}{p_i}.$$

This observation is central to ratio-matching and homogeneous local-scoring methods for unnormalized discrete models (Hyvärinen, 2007; Dawid et al., 2012; Kanamori & Takenouchi, 2017).

Conversely, on a connected graph the complete collection of neighboring ratios determines the normalized probability distribution. For any path $i_0, i_1, \ldots, i_m$ in G,

$$\frac{p_{i_m}}{p_{i_0}} = \prod_{k=0}^{m-1} r_p(i_k, i_{k+1}). \tag{48}$$

Connectedness therefore determines all masses relative to an arbitrary reference state, after which the constraint $\sum_i p_i = 1$ fixes their common scale. Ratios arising from a probability distribution also satisfy cycle consistency: for every cycle $C = (i_0, \ldots, i_m = i_0)$,

$$\prod_{k=0}^{m-1} r_p(i_k, i_{k+1}) = 1.$$

Thus an arbitrary collection of positive directed edge values need not correspond to a globally defined probability distribution.

Ratio matching and local scoring rules. The use of neighboring probability ratios for statistical estimation predates their interpretation as discrete scores. Hyvärinen (2007) introduced ratio matching as a discrete counterpart of score matching (Hyvärinen & Dayan, 2005), initially for binary variables. The normalization constant cancels because the relevant conditional quantities can be expressed through ratios between probabilities of neighboring configurations.

A broader statistical interpretation was developed by Dawid et al. (2012). A scoring rule assigns a loss $\mathsf{S}(i, q)$ to a quoted distribution q after observing $i \in I$. It is *proper* if its expected value under a distribution p is minimized by quoting $q = p$, and *homogeneous* if

$$\mathsf{S}(i, cq) = \mathsf{S}(i, q), \qquad c > 0.$$

Homogeneous scores can therefore be evaluated from an unnormalized positive mass function. Dawid et al. (2012) characterize broad classes of proper local homogeneous scoring rules through clique-based homogeneous functions on an undirected neighborhood graph, with pseudo-likelihood and ratio matching among their examples.

Kanamori & Takenouchi (2017) develop the corresponding viewpoint through composite local Bregman divergences. For a graph-defined neighborhood system, their local potentials depend on vectors of ratios between neighboring masses. For example, their graph extension of the ratio-matching score contains terms of the form

$$\mathsf{S}_{\mathrm{RM}}(i, q) = \sum_{j \in \mathcal{N}(i)} \frac{1}{(1 + q_i/q_j)^2},$$

which is invariant to a global normalization of q. They further show that this graph ratio-matching construction is strictly proper when the neighborhood graph is connected (Kanamori & Takenouchi, 2017). This is consistent with the elementary reconstruction property in equation 48: a connected collection of exact local ratios determines the complete normalized distribution.

Concrete scores. Meng et al. (2022) reinterpret essentially the same local information as a discrete analogue of the continuous Stein score. For a neighborhood $\mathcal{N}(i)$, their *Concrete score* has components

$$c_p(i,j) := \frac{p_j - p_i}{p_i} = \frac{p_j}{p_i} - 1, \qquad j \in \mathcal{N}(i).$$

Hence

$$c_p(i,j) = r_p(i,j) - 1, \qquad r_p(i,j) = 1 + c_p(i,j),$$

so the Concrete score and the local density-ratio field contain exactly the same information.

This identity also explains the completeness result of Meng et al. (2022, Theorem 1): if the neighborhood-induced graph is connected, equality of Concrete-score fields implies equality of the underlying normalized distributions. The same work relates the construction to the continuous score. For

$$\mathcal{N}_\delta(x) = \{x + \delta e_k\}_{k=1}^d,$$

Meng et al. (2022, Proposition 1) show that

$$\frac{1}{\delta} c_p(x; \mathcal{N}_\delta) \longrightarrow \nabla_x \log p(x) \qquad \text{as } \delta \to 0.$$

Concrete Score Matching then provides a sample-based objective for estimating this local score field (Meng et al., 2022).

Score entropy and SEDD. The unshifted ratio itself appears directly in the time reversal of a continuous-time Markov chain. If $Q_t(i,j)$ denotes the forward transition rate from i to j and $p(t)$ its marginal distribution, the reverse-time off-diagonal rates contain the factor

$$\overline{Q}_t(i,j) = \frac{p_j(t)}{p_i(t)} Q_t(j,i), \qquad i \neq j,$$

under the usual row-generator convention (Kelly, 2011; Anderson, 1991; Sun et al., 2023b; Lou et al., 2024). Thus estimating the reverse dynamics of a prescribed forward diffusion requires the time-dependent ratios

$$r_t(i,j) = \frac{p_j(t)}{p_i(t)}.$$

Motivated by this observation, Lou et al. (2024) parameterize these positive ratios directly and introduce *score entropy*. For a positive prediction $s(i)_j$, target ratio $r = p_j/p_i$, and positive weight w_{ij}, the corresponding scalar contribution is

$$w_{ij} \left[s(i)_j - r \log s(i)_j + r(\log r - 1) \right].$$

The excess loss is $r \log \frac{r}{s} - r + s$, which is the Bregman divergence generated by the negative-entropy potential $F(a) = a \log a - a$. Consequently, its unique optimum is $s = r$. In particular, Lou et al. (2024, Proposition 3.2) establish consistency of score entropy for the density-ratio target under full support and positive weights.

The population score-entropy objective itself contains the unknown ratios. Lou et al. (2024, Proposition 3.3) eliminate them to obtain *implicit score entropy*,

$$\mathcal{L}_{\mathrm{ISE}}(s; p) = \mathbb{E}_{i \sim p} \left[\sum_{j \neq i} \left(w_{ij} s(i)_j - w_{ji} \log s(j)_i \right) \right], \tag{49}$$

up to a constant independent of the model. For the structured noising processes used in SEDD, Lou et al. (2024, Theorem 3.4) instead derive a denoising form based on the known forward transition kernel, avoiding the costly evaluation of the implicit objective over all alternative states.

These constructions establish that neighboring ratios are natural normalization-free statistics of a discrete marginal distribution. In SEDD, however, the ratios are learned because they directly determine the reverse-time generator. In Discrete Action Matching they play a different role: they determine the density-dependent *geometry* in which the canonical action is learned. The score-entropy/discrete-geometry connection is also discussed by Rancati et al. (2026, Appendix G); the derivation below explains its use for the DAM kinetic term without claiming priority for the general connection.

B.2 Density ratios in Discrete Action Matching

We now derive the appearance of local density ratios in Discrete Action Matching directly from the graph kinetic energy. Here no ratio-valued score is imposed as a dynamical parameterization, unlike in score-based discrete diffusion. The ratios arise solely because the mobility θ is positively 1-homogeneous.

Corollary B.1 (Sample-local kinetic energy)**.** *Under the assumptions of Theorem 3.1, assume additionally that θ is positively 1-homogeneous. Then the kinetic energy admits the sample-local representation equation 16.*

Recall from equation 33 that

$$\mathcal{K}_{G,\theta}(p,S) = \frac{1}{2}\sum_{\{i,j\}\in E} \omega_{ij}\theta(p_i,p_j)(S_i-S_j)^2. \tag{50}$$

Define the *log-ratio mobility profile* $\kappa_\theta(u) := \theta(1,e^u)$, $u \in \mathbb{R}$. For $p_i, p_j > 0$, positive 1-homogeneity gives

$$\theta(p_i,p_j) = p_i\theta\left(1,\frac{p_j}{p_i}\right) = p_i\kappa_\theta\left(\log\frac{p_j}{p_i}\right). \tag{51}$$

Thus the dependence of the edge mobility on two unknown probability masses separates into the source probability p_i and a function of one neighboring log-ratio.

Using equation 51, the undirected edge sum in equation 50 can be rewritten as a directed neighborhood sum:

$$\begin{aligned}\mathcal{K}_{G,\theta}(p,S) &= \frac{1}{4}\sum_{i\in I} p_i \sum_{j\in\mathcal{N}(i)} \omega_{ij}\kappa_\theta\left(\log\frac{p_j}{p_i}\right)(S_i-S_j)^2 \\ &= \frac{1}{4}\mathbb{E}_{i\sim p}\left[\sum_{j\in\mathcal{N}(i)} \omega_{ij}\kappa_\theta\left(\log\frac{p_j}{p_i}\right)(S_i-S_j)^2\right].\end{aligned} \tag{52}$$

To verify the factor $1/4$, note that the directed sum contains both orientations of every unoriented edge. For $\{i,j\}\in E$, the two contributions are

$$\frac{1}{4}\omega_{ij}p_i\kappa_\theta\left(\log\frac{p_j}{p_i}\right)(S_i-S_j)^2 + \frac{1}{4}\omega_{ij}p_j\kappa_\theta\left(\log\frac{p_i}{p_j}\right)(S_i-S_j)^2. \tag{53}$$

Each source-weighted mobility equals $\theta(p_i,p_j)$ by homogeneity and symmetry, so their sum is precisely $\frac{1}{2}\omega_{ij}\theta(p_i,p_j)(S_i-S_j)^2$. This proves Corollary B.1.

Equation equation 52 is the key statistical closure of Discrete Action Matching. Sampling $X\sim p$ accounts automatically for the source probability p_X. The remaining marginal information is only $\left\{\log\frac{p_j}{p_X}\right\}_{j\in\mathcal{N}(X)}$, rather than the complete probability vector p.

Arithmetic mobility. For $\theta_{\mathrm{ari}}(a,b) = (a+b)/2$, symmetry of the edge sum gives

$$\mathcal{K}_{G,\mathrm{ari}}(p,S) = \frac{1}{4}\sum_{\{i,j\}\in E}\omega_{ij}(p_i+p_j)(S_i-S_j)^2 = \frac{1}{4}\mathbb{E}_{i\sim p}\left[\sum_{j\in\mathcal{N}(i)}\omega_{ij}(S_i-S_j)^2\right]. \tag{54}$$

Each unordered edge contributes once for each sampled endpoint in the expectation, recovering its weight p_i+p_j. Thus the arithmetic action admits an estimator that requires marginal samples but no density ratios. This is an alternative to evaluating the general ratio-based expression.

The Markov realization in equation 21 still uses density ratios, since $\kappa_{\mathrm{ari}}(u) = (1+e^u)/2$. At exact ratios its off-diagonal rates are

$$Q_t(i,j) = \frac{\omega_{ij}}{2}\left(1+\frac{p_j(t)}{p_i(t)}\right)\psi_\varepsilon(S_j(t)-S_i(t)).$$

The simplification of the action therefore does not remove the density dependence of this rate construction.

A scalar log-density representation. The required log-ratios have a natural globally consistent representation. Define $U^\star(i) := \log p_i + c$, where $c \in \mathbb{R}$ is arbitrary. Then $U^\star(j) - U^\star(i) = \log \frac{p_j}{p_i}$. For a time-dependent marginal path,

$$U^\star(t, i) = \log p_i(t) + c(t). \tag{55}$$

Only differences of $U^\star$ enter the method, so the time-dependent gauge $c(t)$ is immaterial.

This scalar representation automatically enforces the structural constraints of a valid ratio field. Defining

$$r_U(i, j) := \exp\left(U(j) - U(i)\right),$$

we have identically

$$r_U(j, i) = r_U(i, j)^{-1}$$

and

$$\prod_{(i,j)\in C} r_U(i, j) = 1$$

for every graph cycle C. Thus, unlike a free directed-edge parameterization, a scalar field cannot produce mutually inconsistent ratios around a cycle.

Local softplus ratio estimation. We estimate the field equation 55 directly from samples. For a fixed full-support marginal p, define

$$\mathcal{R}_p(U) := \mathbb{E}_{i\sim p}\left[\sum_{j\in\mathcal{N}(i)} \omega_{ij}\,\mathrm{softplus}\left(U_j - U_i\right)\right], \tag{56}$$

where $\mathrm{softplus}(z) := \log(1 + e^z)$. Like the local scoring objectives reviewed above, equation 56 depends on p only through the outer expectation and therefore requires neither the individual masses p_i nor their normalization constant during empirical optimization.

The population target can be characterized edge by edge.

Proposition B.2 (Consistency of the local softplus ratio loss)**.** *Let G be connected and let $p \in \mathcal{P}_+(I)$. Then*

$$\operatorname*{argmin}_{U\in\mathbb{R}^N} \mathcal{R}_p(U) = \{\log p + c\mathbf{1} : c \in \mathbb{R}\}. \tag{57}$$

In particular, every minimizer satisfies

$$U_j - U_i = \log\frac{p_j}{p_i}, \qquad \{i, j\} \in E.$$

Proof. Consider an unoriented edge $\{i, j\} \in E$ and write $\delta_{ij} := U_j - U_i$. The two directed occurrences of this edge in equation 56 contribute

$$\ell_{ij}(\delta_{ij}) = \omega_{ij}\left[p_i\,\mathrm{softplus}(\delta_{ij}) + p_j\,\mathrm{softplus}(-\delta_{ij})\right].$$

Differentiation gives

$$\ell'_{ij}(\delta) = \omega_{ij}\frac{p_i e^\delta - p_j}{1 + e^\delta}.$$

Therefore

$$\ell'_{ij}(\delta) = 0 \quad\Longleftrightarrow\quad \delta = \log\frac{p_j}{p_i}.$$

Moreover,

$$\ell''_{ij}(\delta) = \omega_{ij}\frac{(p_i + p_j)e^\delta}{(1 + e^\delta)^2} > 0,$$

so this stationary point is the unique minimum of the edge contribution.

The collection $U_i^\star = \log p_i$ simultaneously realizes the unique minimizing difference on every edge. Hence any global minimizer must have $U_j - U_i = U_j^\star - U_i^\star$ for every edge. Since G is connected, $U - U^\star$ is constant on I, proving equation 57. □

Proposition B.2 explains why we parameterize the ratio estimator through a scalar field $U_\phi(t,i)$ rather than independent directed predictions. At the population optimum,

$$\exp\left(U_\phi(t,j) - U_\phi(t,i)\right) = \frac{p_j(t)}{p_i(t)},$$

while positivity, reciprocity, and cycle consistency hold automatically for arbitrary model parameters.

Relation to implicit score entropy. The local softplus loss is related to, but distinct from, implicit score entropy. To make the comparison explicit, restrict equation 49 to graph neighbors, take symmetric weights $w_{ij} = w_{ji} = \omega_{ij}$, and use the coherent positive-score parameterization $s_U(i)_j := \exp(U_j - U_i)$. Then $-\log s_U(j)_i = U_j - U_i$, and the implicit score-entropy objective becomes

$$\mathcal{L}_{\mathrm{ISE}}(U;p) = \mathbb{E}_{i\sim p}\left[\sum_{j\in\mathcal{N}(i)} \omega_{ij}\left(e^{U_j-U_i} + U_j - U_i\right)\right].$$

For one edge and $\delta = U_j - U_i$, its two orientations contribute

$$\omega_{ij}\left[p_i\left(e^{\delta}+\delta\right) + p_j\left(e^{-\delta}-\delta\right)\right].$$

The derivative vanishes uniquely at $\delta = \log\frac{p_j}{p_i}$, and the edge contribution is strictly convex. Consequently,

$$\operatorname*{argmin}_U \mathcal{L}_{\mathrm{ISE}}(U;p) = \operatorname*{argmin}_U \mathcal{R}_p(U) = \{\log p + c\mathbf{1}\}.$$

Thus the two objectives are *target-equivalent* under the scalar log-density parameterization, but they are not the same estimator. Their nonlinear terms are respectively $e^{U_j-U_i}$ and $\mathrm{softplus}(U_j-U_i)$, and hence their gradients and finite-sample behavior differ. In particular, implicit score entropy inherits an exponential dependence on the predicted log-ratio, whereas the derivative of the softplus term is logistic and bounded. Our method uses equation 56; the relation to Lou et al. (2024) is therefore a connection through the common population density-ratio target, rather than an identity of training losses.

With the local log-density field estimated, the kinetic energy becomes sample-computable through

$$\mathcal{K}_{G,\theta}(p,S) = \frac{1}{4}\mathbb{E}_{i\sim p}\left[\sum_{j\in\mathcal{N}(i)} \omega_{ij}\kappa_\theta\left(U_j^\star - U_i^\star\right)(S_i - S_j)^2\right].$$

We next use this representation to construct the empirical two-stage objective and the corresponding Markov dynamics.

B.3 Empirical training and Markov realization

Why declare an interpolant? The population action integrates expectations over all $t \in [0,1]$. Finite snapshots do not supply samples at arbitrary intermediate times. Equation equation 17 declares the distribution sampled at such times. It interpolates distributions, not matched individual trajectories. Conditional on the data, sample T uniformly, choose the reference with probability δ, otherwise choose the left or right snapshot with probabilities $1-h(u)$ and $h(u)$, and sample a state. At endpoints use the corresponding empirical/reference mixture. The resulting state law is exactly $\widetilde{p}(T)$, with $\widetilde{p}_i(t) \geq \delta/N$.

The original AM algorithm assumes intermediate-time sample access (Neklyudov et al., 2023, Algorithm 1). Merely replacing its time integral by derivative evaluations at finitely many observed knots changes the optimization problem. In particular, a candidate time-only gauge $c(t)\mathbf{1}$ changes that fixed-knot loss by $c(0)-c(1)+\sum_k \alpha_k c'(t_k)$. For $A>0$, choose $c_A(t) = A\int_0^t \prod_k (s-t_k)^2\,ds$. Every sampled derivative vanishes while $c_A(1) > 0$, so the unrestricted fixed-knot objective is unbounded below. All spatial differences, actions and rates are unchanged. Spatial zero-mean gauge fixing removes this example but not generally time-dependent spatial perturbations that vanish at the knots. This counterexample does not assert divergence of every restricted neural optimizer, nor that

interpolation is the only possible well-posed temporal formulation. DICE studies related finite-time AM difficulties (Blickhan et al., 2025, Section 6). We use the continuous-time estimator for the declared path, whose expected boundary and derivative terms cancel under a smooth time gauge; an individual Monte Carlo batch need not be gauge invariant.

Regularity of the practical path. A globally C^2 positive path satisfies Theorem 3.1. The cubic smoothstep construction is C^1 and piecewise smooth, generally not globally C^2. On each interval the theorem's calculation applies. For continuous, piecewise C^1 candidates, integration-by-parts boundary terms telescope because the path and candidate are continuous. The zero-mean canonical potential is continuous and piecewise C^1 for smoothstep, since the path tangent vanishes on both sides of each knot. Thus the identity extends in this piecewise class. A linear probability interpolant is continuous but can have tangent jumps; its canonical potential may jump. The identity still holds with the target kinetic energy as a finite additive constant, but a continuous network need not represent that discontinuous target exactly. Endpoint matching and regularity must be checked for the actual schedule, not inferred from a continuous-time network input.

Stage 1 and its guarantee. At a fixed positive marginal the local softplus loss in equation 56 has the minimizers in Proposition B.2. For a snapshot-only population loss $\sum_k w_k \mathcal{R}_{p(t_k)}(U(t_k))$, assume $w_k > 0$, simultaneous realizability at all knots, and global population minimization; the conclusion then holds only at those knots:

$$U_{\phi^\star}(t_k, j) - U_{\phi^\star}(t_k, i) = \log \frac{p_j(t_k)}{p_i(t_k)}.$$

The implemented objective instead estimates $\mathcal{R}_{\widetilde{p}}(\phi) = \int_0^1 \mathcal{R}_{\widetilde{p}(t)}(U_\phi(t))\, dt$. Its unrestricted optimum recovers $\log \widetilde{p}(t)$ modulo a gauge almost everywhere, with the same realizability and exact-minimization qualifications for a model class. This is not a finite-data recovery theorem for $p(t)$. Zero-count empirical states make unrestricted unsmoothed edge-ratio estimation problematic; reference mixing supplies positive support, at the cost of changing the target. Atlas pseudocounts, empirical resampling, and training reference mixing are separate operations (Appendix D.4).

Frozen geometry and conditional unbiasedness. For a fixed U and any positive path p, define

$$\widetilde{c}_{ij}^{U}(p) = \frac{\omega_{ij}}{2}\left[p_i \kappa_\theta(U_j - U_i) + p_j \kappa_\theta(U_i - U_j)\right], \tag{58}$$

$$\widetilde{\mathcal{K}}_U(p, S) = \frac{1}{2} \sum_{\{i,j\} \in E} \widetilde{c}_{ij}^{U}(p)(S_i - S_j)^2. \tag{59}$$

These positive symmetric conductances describe the training energy. They alone do not ensure the same sampling current. Indeed, set $a = \omega_{ij} p_i \kappa_\theta(U_j - U_i)$, $b = \omega_{ij} p_j \kappa_\theta(U_i - U_j)$ and $z = S_j - S_i$. The source-first training current is $(a + b)z/2$, whereas the rate construction yields

$$a\psi_\varepsilon(z) - b\psi_\varepsilon(-z) = \frac{a+b}{2} z + \frac{a-b}{2}\big[\psi_\varepsilon(z) + \psi_\varepsilon(-z)\big].$$

At $z = 0$ this equals $(a-b)\varepsilon \log 2$, which need not vanish. At exact ratios for p, $a = b = \omega_{ij}\theta(p_i, p_j)$ and $\widetilde{c}_{ij}^{U^\star}(p) = \omega_{ij}\theta(p_i, p_j)$. For fixed S and a frozen fitted U, fresh draws in equation 20 satisfy

$$\begin{aligned}\mathbb{E}[\widehat{\mathcal{L}}_{\mathrm{DAM}} \mid \mathcal{D}, U] = \widetilde{p}(0)^\top S(0) - \widetilde{p}(1)^\top S(1) \\ + \int_0^1 \left[\widetilde{p}(t)^\top \partial_t S(t) + \widetilde{\mathcal{K}}_U(\widetilde{p}(t), S(t))\right] dt.\end{aligned} \tag{60}$$

To verify this, average endpoint samples, then condition on $T = t$ and average the sampled state. Pairing the two orientations of each edge gives equation 59; finally average uniform time. At exact ratios for this same $\widetilde{p}$ the expression is its chosen-mobility action loss. This conditional estimator identity says nothing about unbiasedness of an optimized network or identification of the unknown intermediate population path.

A neighborhood sum $H_f(i) = \sum_j \omega_{ij} f(i, j)$ may be estimated with $J \sim q(\cdot \mid i)$, where $q(j \mid i) > 0$ on every neighbor:

$$\widehat{H}_f(i, J) = \frac{\omega_{iJ}}{q(J \mid i)} f(i, J), \qquad \mathbb{E}[\widehat{H}_f(i, J) \mid i] = H_f(i).$$

Independent repetitions reduce variance. The proposal is known and is not a density-ratio estimate. Freezing U prevents action optimization from changing the conductances rather than fitting the action in the estimated geometry.

Markov realization and its law. For finite continuous fields abbreviate $Q_t = Q_t^{S,U}$, with

$$Q_t(i,j) = \omega_{ij}\kappa_\theta(U_j - U_i)\psi_\varepsilon(S_j - S_i) \quad (j \in \mathcal{N}(i)), \qquad Q_t(i,i) = -\sum_{j \neq i} Q_t(i,j).$$

Nonneighbor rates vanish. Nonnegative finite off-diagonals and zero row sums give a valid generator; continuity makes it bounded on the compact time interval. Both S and U may be shifted by time-dependent spatial constants without changing rates. Let Y_t denote the constructed chain and $\mu(t)$ its column-vector law. Then $\dot{\mu}(t) = Q_t^\top \mu(t)$. If $U_j - U_i = \log(p_j/p_i)$, homogeneity and the split identity give

$$p_i Q_t(i,j) - p_j Q_t(j,i) = \omega_{ij}\theta(p_i,p_j)(S_j - S_i), \tag{61}$$

$$(Q_t^\top p)_i = \sum_j \omega_{ij}\theta(p_i,p_j)(S_i - S_j) = [L_\theta(p)S]_i. \tag{62}$$

If additionally $S = S^\star$, the prescribed path satisfies the same forward equation as μ. Matching initial laws and uniqueness yield $\mu = p$. This applies intervalwise for the piecewise-smooth construction. It does not equate a target path with a simulated law under approximate fits. The algebraic current identity also holds relative to $\widehat{p} = \mathrm{softmax}(U)$, but that vector is not automatically μ.

The smooth split obeys $\psi_\varepsilon(z) = \max(z,0) + \varepsilon\log(1 + e^{-|z|/\varepsilon})$. At exact ratios the added symmetric traffic satisfies detailed balance with p and leaves its instantaneous current unchanged, in the spirit of marginal-preserving generator additions (Campbell et al., 2024). This does not transfer an exact-current property to imperfect ratios.

Inference and trajectory simulation. For marginal evaluation, choose an initial distribution and time grid; freeze Q at each substep's evaluation time and propagate $\mu_{\ell+1} = \exp(\Delta t_\ell Q_{t_\ell}^\top)\mu_\ell$. The experiment-specific grids and refinement checks are in Appendix D. This deterministic calculation is different from drawing trajectories.

For an optional trajectory with a piecewise-constant generator, at state i and time t draw $\tau \sim \mathrm{Exp}(\lambda_i)$, where $\lambda_i = -Q_t(i,i)$. If $\lambda_i = 0$, remain at i until the next rate-change boundary. If $t + \tau$ reaches or crosses that boundary, advance to the boundary without a jump and restart using the new rates. Otherwise choose the first-jump destination with probability $Q_t(i,j)/\lambda_i$, advance to $t + \tau$, and repeat (Anderson, 1991). This is exact for the frozen generator and a time-discretization approximation for continuously varying rates.

C Extended discussions

C.1 Extended related work

This section provides additional context for the methods summarized in Section 4. We first discuss extensions of Action Matching and discrete Wasserstein gradient-flow methods, and then distinguish DAM from discrete diffusion, score-based, and flow-based generative models.

Extensions of Action Matching. Beyond the fixed-path formulation of Neklyudov et al. (2023), Wasserstein Lagrangian Flows additionally optimize the probability path subject to observed marginal constraints (Neklyudov et al., 2024), while Berman et al. (2024) consider parameter-dependent reduced dynamics. Wasserstein Hamiltonian formulations describe kinetic evolution directly on probability space (Chow et al., 2020), and Wasserstein Lagrangian Mechanics extends marginal-based learning to second-order population dynamics (Guan et al., 2026a). Connections between Action Matching and quadratic optimal transport under optimal vector-field parameterizations are studied by Kornilov & Korotin (2025). These developments concern continuous state spaces, whereas DAM focuses on the additional density-dependent geometry that appears on finite graphs.

Discrete Wasserstein geometry. The discrete transport geometry used by DAM originates from constructions that endow finite-state probability spaces with mobility-dependent dynamical transport metrics. Maas (2011) constructs a metric under which the heat equation of a reversible Markov chain is the gradient flow of relative entropy, while Chow et al. (2012) derive related discrete Fokker–Planck gradient flows on weighted graphs. Mielke (2011) studies generalized finite-state gradient structures. Subsequent developments include entropy dissipation in Chow et al. (2018), Riemannian formulations in Li & Montúfar (2018); Li (2022), and the study of geodesics and their support in Gangbo et al. (2019); Erbar et al. (2019). These results provide the graph gradient, continuity equation, and kinetic-energy metric underlying DAM.

Discrete Wasserstein geometry has also been used computationally for objectives different from population-dynamics inference. Diaz-Garcia et al. (2021) connect Wasserstein gradient flows with population-game dynamics, while Sun et al. (2023a) use discrete gradient flows to construct samplers for prescribed target distributions. In these settings the driving objective is given, whereas DAM infers a canonical velocity from observed temporal marginals.

Inverse Wasserstein gradient flows. A related inverse problem assumes that the observed population evolution is generated by the Wasserstein gradient flow of an unknown energy functional. In continuous spaces, Bunne et al. (2022) formulate this problem through repeated JKO steps and Brenier maps (Brenier, 1991). Terpin et al. (2024) derive a first-order alternative based on precomputed optimal-transport information, while Persiianov et al. (2026) remove this preprocessing through a saddle-point formulation over the energy and transport maps.

Rancati et al. (2026) extend the first-order JKO approach to finite graphs using discrete Wasserstein geometry. Their construction recovers an autonomous free-energy functional after computing graph transport geodesics between observed marginals. DAM instead treats the marginal evolution itself as prescribed and recovers its canonical minimum-action velocity, which may depend explicitly on time. The two methods therefore impose different structural assumptions: inverse JKO learning seeks one energy functional explaining the complete evolution, whereas DAM selects the minimum-action tangent compatible with the marginal path.

A deterministic map cannot generally push an atomic source to an arbitrary target: a point mass cannot map to a two-point mixture. Thus a map-based inverse-JKO formulation (Persiianov et al., 2026) needs adaptation when transport splits mass. Graph fluxes provide one alternative; Li & Montúfar (2018) discuss the distinct dynamical graph geometry. This does not exclude every finite-state analogue of a transport map.

Discrete diffusion and score-based models. Discrete generative modeling provides a separate use of finite-state stochastic dynamics. Multinomial diffusion and D3PM construct prescribed corruption processes whose reversal is learned from data (Hoogeboom et al., 2021; Austin et al., 2021), while continuous-time formulations represent forward and reverse dynamics through continuous-time Markov chains (Campbell et al., 2022). Concrete Score Matching (Meng et al., 2022) and SEDD (Lou et al., 2024) use neighboring probability ratios as discrete analogues of continuous scores or as sufficient quantities for reverse-time dynamics.

DAM uses the same type of local density-ratio information for a different purpose. The ratios do not parameterize the population velocity or a reverse diffusion process; they determine the marginal-dependent mobility needed to evaluate the discrete Wasserstein kinetic energy. A separate potential then represents the canonical population dynamics. Appendix B develops this connection in detail.

Discrete flow-based generative models. Discrete Flow Matching (Gat et al., 2024) and related multimodal generative-flow constructions (Campbell et al., 2024) learn dynamics along prescribed or designed probability paths between endpoint distributions. Other approaches construct flows on the probability simplex using Dirichlet (Stark et al., 2024) or Fisher–Rao geometry (Davis et al., 2024). Their primary goal is generative transport between endpoint distributions, with the intermediate probability path chosen as part of the modeling procedure. DAM instead addresses population-dynamics inference, where intermediate temporal marginals are observed data and the objective is to recover the canonical graph velocity associated with their evolution.

Graph transport computation and population geometry. Lavenant et al. (2018) develop dynamical transport on discrete surfaces, and Erbar et al. (2020) give computational methods for dynamical transport on graphs. Gentile & Murphy (2026) study static and dynamic graph barycenters. These works optimize paths or distances between supplied distributions. The quadratically regularized graph transport problem of Essid & Solomon (2018) uses linear ground distance, a different objective from DAM's kinetic action. Peyré (2015) uses entropic regularization of transport plans to compute gradient flows; this regularization is distinct from entropy in a learned free energy.

Diaz-Garcia et al. (2021) connect Wasserstein gradient flows to population games, while Mertikopoulos & Sandholm (2018) study game dynamics through Riemannian costs of motion. The simplex course of Pistone (2020) supplies information-geometric background; its Fisher geometry differs from the graph-Wasserstein metric used here. Erbar (2014) develops a nonlocal jump-kernel construction on Euclidean space, providing broader context rather than the direct finite-state source of each displayed operator.

Attribution and the scope of the geometry. We use probability coordinates and a prescribed symmetric conductance, as in graph formulations associated with Chow and collaborators; Maas's reversible-chain construction uses stationary-relative density, and Mielke develops finite-state gradient structures. These are connected constructions, not an exhaustive division into two schools. The conversion and its limitations are in equation 46. Geodesic work (Gangbo et al., 2019; Erbar et al., 2019) studies global and boundary questions; our action identity only uses the positive interior tangent problem and does not assume one entire admissibility framework contains another. More specifically, Erbar et al. (2019, Theorem 3.4) give a Hamilton–Jacobi duality proof building on the earlier duality result of Gangbo et al. (2019). The former use an irreducible reversible Markov triple and an admissible mean; the latter formulate symmetric graph weights and a smooth, concave, homogeneous interior-positive mobility with finite $\int_0^1 g(r, 1-r)^{-1/2}\, dr$, treating geodesics that can meet the boundary. Their global results require more than the interior assumptions used here.

Path optimization and mechanics. The WLF formulation of Neklyudov et al. (2024) expresses a primal action over probability paths through a dual problem: the action potential is fitted for a path, and an outer objective optimizes the path. Differentiation through a parameterized path is distinct from backpropagating through a simulated trajectory. We make no historical "first" claim for that dual principle. Guan et al. (2026a) instead introduce Wasserstein Lagrangian Mechanics to learn second-order population evolution from snapshots; this is a distinct objective from WLF and is not a directly interchangeable finite-state baseline. Kornilov et al. (2024) study the optimal-transport connection for Flow Matching (FM), whereas Kornilov & Korotin (2025) establish the corresponding AM equivalence.

Hamiltonian and accelerated sampling. Zhou et al. (2025) use mobility-weighted graph kinetic energy in a damped Hamiltonian system for a probability distribution and momentum potential, with a jump-process realization. DAM shares a potential-based continuity equation after matching conductances and signs, but fits a prescribed path rather than imposing their additional momentum equation to approach a target distribution. This relationship does not turn the learned action potential into a target log density for classical Hamiltonian Monte Carlo, nor imply accelerated mixing or greater sampling accuracy. Such guarantees require a separately specified dynamics and analysis.

Time-discrete and aggregate-observation population inference. DICE (Blickhan et al., 2025) learns population dynamics on continuous state spaces using an objective derived from a time-discrete weak continuity equation. Its loss is invariant under spatially constant, time-dependent shifts of the potential. This discretization concerns time, whereas DAM provides a finite-state counterpart of Action Matching under a chosen graph mobility. With finite observations, DAM explicitly specifies the interpolating marginal path used by its learning objective.

Mascherpa et al. (2026) estimate a common Markov transition matrix jointly with transport plans between aggregate observations through a convex entropic optimal-transport formulation. DAM instead learns the canonical, potentially time-dependent velocity of a prescribed marginal path under its chosen mobility. Its local density-ratio representation supplies the marginal information needed to evaluate the action from samples without graph-geodesic preprocessing. The FEGF adaptations

used in our experiments, including the time-indexed potential for gastrulation, are specified in Appendix D.2.

Potential constructions and graph bridges. Potential representations and the associated Riemannian formulation on the probability simplex are developed by Chow et al. (2019). In discrete generative modeling, Shaul et al. (2025) construct kinetic-optimal flows using graph-Laplacian potentials along prescribed or designed probability paths. Campbell et al. (2024) show how detailed-balance additions to transition rates preserve a prescribed marginal evolution. These works establish closely related potential constructions and the freedom to choose among compatible Markov realizations. DAM combines these established structures with a mobility-dependent action objective and local-ratio estimation for a training path specified from unpaired temporal snapshots.

Generalized Schrödinger bridges on graphs (Theodoropoulos et al., 2026) optimize path laws under endpoint constraints, a relative-entropy penalty against a reference process, and additional running costs. DAM fits a prescribed marginal path and graph mobility.

C.2 Cortical prediction from unpaired observations

Graph-based studies model tau propagation through brain connectivity and its interaction with amyloid burden (Vogel et al., 2020; Lee et al., 2022), while X-FlowNet and NeuralOMT learn transport-related dynamics from longitudinal regional or cortical measurements (Dan et al., 2023; Huang et al., 2024). An interesting direction is to combine these models with the unpaired marginal objective of DAM **(ours)**, potentially exploiting observations without matched temporal samples and reducing dependence on complete longitudinal records. Such training could complement paired supervision and improve data utilization; whether it improves next-visit accuracy remains to be tested. The present formulation does not by itself identify subject-specific transitions from unmatched cohorts. Unpaired state samples are not automatically unpaired subjects: personalized prediction additionally requires a justified conditional population model and registered, common state labels.

The sphere and hand experiments provide an intermediate step from controlled synthetic dynamics to cortical surface modeling. To extend this progression to cortical biomarker evolution with changing total burden, we propose developing unbalanced action matching with source and sink terms, alongside conditioning on baseline information and addressing cohort and disease-stage differences. Normalizing PET maps discards absolute burden, and balanced probability transport alone cannot model its increase; NeuralOMT already permits variable mass (Huang et al., 2024). Evaluation should use held-out longitudinal subjects and distinguish predictive accuracy from biological identification of inferred transport. At inference, an unavailable follow-up or future target burden must not enter the closure, initialization or rescaling. Fitting observed times does not establish extrapolation, and neither improved prediction nor clinical benefit follows from the present experiments.

D Additional experimental details

We describe the synthetic reconstruction and gastrulation interpolation benchmarks, compare mobility and density-ratio choices, and evaluate transport on two surface meshes.

D.1 Finite-snapshot synthetic protocol

Data and graphs. We adapt the graph free-energy benchmark of Rancati et al. (2026) to five sampled marginals at times $0, 1/4, 1/2, 3/4, 1$. Paths follow free-energy descent with energy $F(p) = \sum_i p_i V_i + \beta \sum_i p_i \log(p_i/\pi_i)$, where π is the stationary distribution and $V_i \sim \mathrm{Unif}[-1, 1]$ independently. Initial relative densities $\rho_i = p_i/\pi_i$ are sampled uniformly on $[0, 100]$ and normalized to satisfy $\sum_i \pi_i \rho_i(0) = 1$.

We use eleven graph families: complete, Erdős–Rényi, regular, Watts–Strogatz, stochastic block model, Delaunay, Euclidean minimum spanning tree, grid, k-partite, torus and Apollonian. Geometric constructions use jitter 0.05 and an edge-weight length scale equal to 0.25 times the mean edge length. Row-normalizing symmetric weights gives the transition matrix K; normalized weighted degrees give π. DAM uses edge weights $\omega_{ij} = \pi_i K_{ij}$.

Table 2: Synthetic and gastrulation training settings. All models use Adam, float64 arithmetic, weight decay 10^{-6}, gradient-clipping norm 1 and 200 updates per epoch.

Setting	DAM	FEGF
Architecture	16-dimensional node embedding, scalar time, two width-64 ReLU layers, scalar output	Synthetic: edge network with embedding 16 and width 64; gastrulation: time-indexed scalar table
Learning rate	$5 \cdot 10^{-4}$	Synthetic: $5 \cdot 10^{-4}$; gastrulation: 10^{-2}
Batch size	128	64
Maximum epochs	Ratio: 60; action: 100	100
Selection	Ratio validation loss; action observed-time validation TV	Observed-time validation TV
Early-stopping patience	Ratio: 10 epochs; action: none	15 epochs
Support treatment	Uniform-reference mixing $\delta = 0.01$	Pseudocount 0.5 per state; rollout mixture $0.98p + 0.02\pi$

The base setting is $(N, M, \beta) = (60, 1000, 0.1)$, where N is the number of states and M the samples per marginal in each split. One-factor sweeps vary $\beta \in \{0.01, 0.1, 0.2\}$, $N \in \{20, 40, 60, 80, 100\}$ or $M \in \{500, 1000, 2000, 5000, 10000\}$. Their union contains eleven settings, giving 363 graph/setting/seed combinations across seeds 42–44. Methods receive matched, independently sampled training, validation and test splits. Smaller sample budgets use prefixes of the same draws. Fitting and model selection use sampled observations only, without access to the generating probabilities, potentials or intermediate states.

Numerical generation. For $\beta = 0.1$, we use adaptive Dormand–Prince integration. For $\beta \in \{0.01, 0.2\}$, we use backward Euler in log-density coordinates with steps at most 10^{-4}, checked against $5 \cdot 10^{-5}$. Generated paths must have positive finite densities, mass error at most 10^{-8}, nonincreasing energy and observation-probability TV changes at most 10^{-4} under refinement.

Evaluation. We report Hellinger distance $H(p, q) = 2^{-1/2}\left[\sum_i (\sqrt{p_i} - \sqrt{q_i})^2\right]^{1/2}$ and total variation $\mathrm{TV}(p, q) = \frac{1}{2}\sum_i |p_i - q_i|$. All methods start from the training empirical marginal at $t = 0$, propagate without resets and are scored against independent test histograms at the four subsequent observation times. This measures reconstruction at observed times and includes test-sampling noise. We average errors over times within each run and over graphs/settings within each seed, then report means and sample standard deviations across three seeds.

DAM propagates marginals through frozen-generator matrix exponentials with eight substeps per observation interval for both validation and testing (Appendix B.3). Its softplus rate split uses $1/\varepsilon = 1$. FEGF propagates its predicted law in stationary-relative coordinates. Both comparisons use deterministic marginal propagation.

D.2 Training and baseline specification

Both DAM stages use the same smoothstep interpolant with uniform-reference mixing, including at the endpoints, and sample four neighbors per state with importance weights. Histogram DAM evaluates $U(t, i) = \log \widetilde{p}_i(t)$; learned-ratio DAM minimizes the local softplus loss. Both variants share action architectures, initial weights, training streams, update budgets and selection rules. Potentials have zero spatial mean. ReLU potentials are continuous and piecewise differentiable, and we use their almost-everywhere time derivatives. Table 2 gives the fixed settings, reused across gastrulation folds without a new tuning search.

FEGF adaptation. We adapt the free-energy gradient-flow method (FEGF) of Rancati et al. (2026). Empirical probabilities receive pseudocount 0.5 and are converted to $\rho = p/\pi$. Adjacent-interval targets use the actual time gap, a weighted-incidence least-squares solve with one pinned vertex, and the negative free-energy-difference convention. Logarithmic means use an input floor of 10^{-6}, and the continuity residual is checked. The last interval target is reused at the terminal observation.

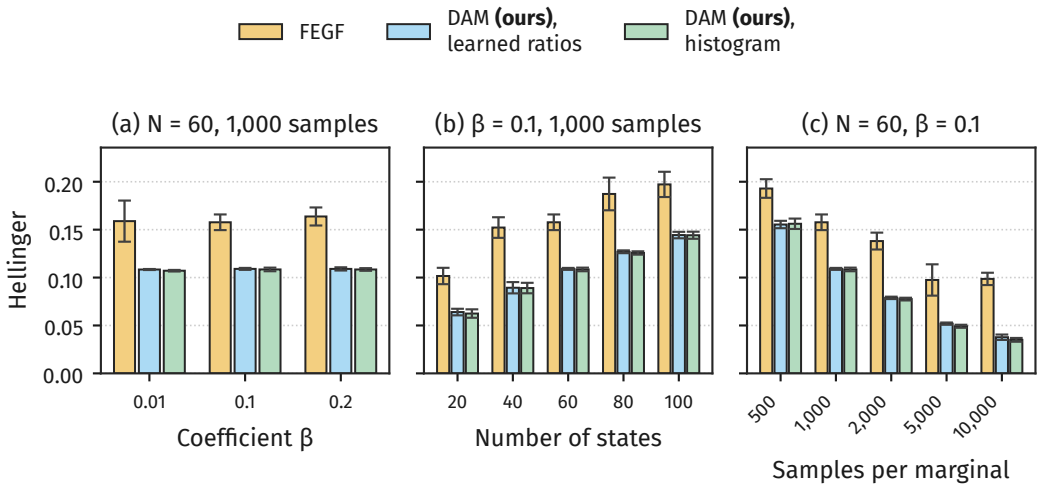


Figure 4: Synthetic sweeps with logarithmic mobility. Bars average Hellinger errors over eleven graph families; whiskers show sample SD across three seeds. Each panel varies one factor from the base setting. Smaller sample budgets share draws.

Synthetic FEGF learns edge differences and a positive entropy coefficient. Gastrulation FEGF instead learns a scalar value per state and observation time, interpolated in time, with no trained entropy coefficient. Both minimize the sample-based velocity-target loss and select checkpoints by observed-time validation TV under propagation without resets. The omitted atlas marginal is used only for scoring.

FEGF uses conductances $\pi_i K_{ij} \theta_{\log}(p_i/\pi_i, p_j/\pi_j)$, whereas logarithmic DAM uses $\pi_i K_{ij} \theta_{\log}(p_i, p_j)$. These geometries differ when π is nonuniform. The comparison evaluates marginal prediction under these respective geometries; fixed optimizer budgets do not imply equal prior tuning effort.

D.3 Synthetic results and ablations

Figure 2 reports the base setting, while Figure 4 varies one factor at a time. Both DAM variants achieve lower mean Hellinger errors than FEGF across all base-setting graph families and the reported sweep averages. Learned ratios and histograms perform comparably, with no consistent advantage for learned ratios across metrics.

Mobility and density-ratio comparisons. We refit each DAM variant with arithmetic, geometric, harmonic and logarithmic mobility, holding data, smoothing, interpolation, architecture and optimizer budgets fixed. Table 3 reports both metrics, and Figure 5 shows seed variability. Arithmetic action training is algebraically ratio-free, although its Markov sampler still uses density ratios. Harmonic mobility gives the lowest gastrulation means among the tested choices; it is a descriptive comparison, without independent mobility selection.

D.4 Mouse gastrulation

Data and graph. We use processed counts associated with the mouse gastrulation atlas of Pijuan-Sala et al. (2019): 89,267 cells, 34 cell types and nine stages from E6.5 to E8.5 at quarter-day intervals. These processed counts define the benchmark; the exact source release and filtering steps have not been fully recovered.

Atlas marginals use pseudocount smoothing, $p_{ki} = (n_{ki} + 0.5)/(n_k + 17)$. DAM additionally mixes each resampled empirical marginal with a uniform reference, since resampling can omit cell types. FEGF applies the separate support treatments in Table 2. The graph combines a connected

Table 3: Mean errors for separately fitted mobilities. Synthetic results average eleven base-setting graphs and three seeds; gastrulation results average seven omitted stages and three seeds. L: learned ratios; Hst: histograms. Lower is better; values are rounded to three decimals.

Dataset	Mobility	H (L)	H (Hst)	TV (L)	TV (Hst)
Synthetic	Arithmetic	0.110	0.110	0.106	0.107
	Geometric	0.109	0.108	0.106	0.106
	Harmonic	0.109	0.108	0.106	0.106
	Logarithmic	0.109	0.109	0.106	0.106
Gastrulation	Arithmetic	0.285	0.262	0.268	0.242
	Geometric	0.234	0.229	0.220	0.215
	Harmonic	0.220	0.203	0.207	0.188
	Logarithmic	0.250	0.251	0.232	0.233

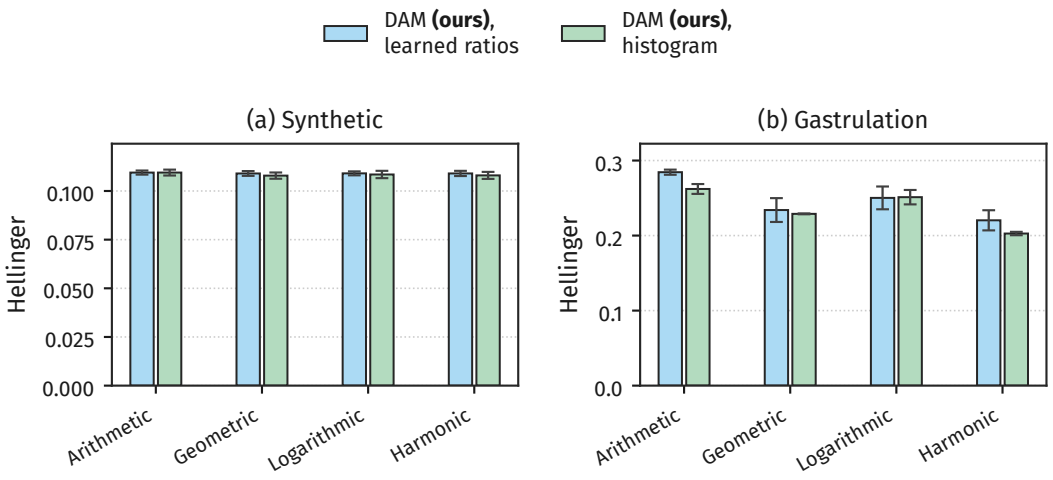


Figure 5: Hellinger errors for learned-ratio and histogram DAM, refitted for each mobility. Bars and whiskers show means and sample SD across three seeds after averaging graphs or omitted stages. Panels use separate vertical scales.

lineage prior with 44 unordered edges and a background weight of 0.01 added to every adjacency entry before normalization. All 561 unordered state pairs consequently support transport, including transitions outside the lineage prior.

Leave-one-stage-out interpolation. For each of the seven interior stages, all methods are fitted to the same eight remaining stages, including later observations. Times are normalized by $(t - 6.5)/2$. The omitted stage is removed before constructing training inputs and FEGF targets. Each retained stage supplies $10{,}000$ training, $5{,}000$ validation and $5{,}000$ independent test draws; checkpoint selection uses retained-time validation only.

Each prediction starts from the training empirical marginal at the immediately preceding retained stage and is scored against the omitted, pseudocount-smoothed atlas marginal. This evaluates interpolation on the measured atlas. The empirical initial law differs from the smoothed training interpolant, so the exact-initialization condition of the Markov realization need not hold.

Numerics and results. All methods propagate deterministic marginals. Without consulting target errors, we compare $128/256$ substeps and double up to $1024/2048$ until the prediction-TV change is at most 0.001. All 21 stage/seed cases pass. Scores average the seven stages equally; standard

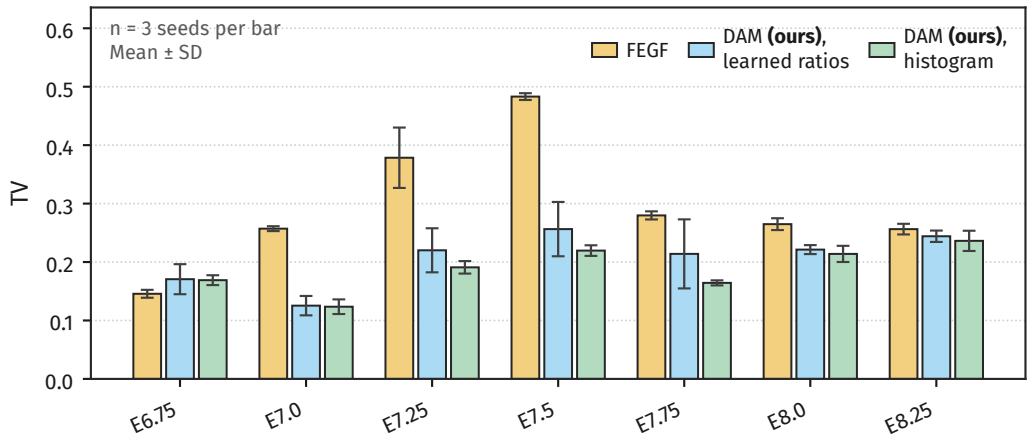


Figure 6: Gastrulation TV errors against omitted, smoothed atlas marginals. Each fold uses eight retained stages, including later observations. DAM uses harmonic mobility; FEGF uses its stationary-relative logarithmic geometry. Whiskers show sample SD across three sampling/optimization seeds, rather than biological replicates.

deviations use three stage-averaged seed values and reflect resampling and optimization of the same atlas.

With harmonic mobility, both DAM variants achieve lower mean Hellinger and TV errors than FEGF and persistence, which keeps the preceding empirical marginal unchanged. Figures 3 and 6 show stagewise Hellinger and TV errors. Differences between learned ratios and histograms depend on the stage. Discrepancy from the training interpolant also reflects initialization and fitting errors, which remain after numerical refinement.

D.5 Sphere and hand surface transport

We evaluate DAM on a sphere with 1020 vertices and a hand mesh with 1515 vertices, using marginal paths supplied by the ADMM surface-transport solver of Lavenant et al. (2018). The learner's graph geometry differs from the teacher's triangle-based geometry.

Table 4: Mean TV at 16 omitted times against independent test-particle marginals, seed 42; lower is better.

Method	Sphere	Hand
DAM, histogram	0.169	0.246
DAM, local + global ratios	0.192	0.250
Tuned FEGF	0.564	0.843

Each mesh uses 50,000 training, 25,000 validation and 25,000 independent test particles, with seed 42. Here DAM uses paired vertex trajectories derived from the teacher's continuity equation. Both DAM variants obtain lower omitted-time TV than tuned FEGF, with histogram DAM competitive or better (Table 4). These results concern two supplied paths and one seed; transfer to new endpoint pairs or unseen meshes remains untested.

D.6 Surface teacher, sampling and learning protocol

Teacher and paired sampling. The ADMM solver uses 31 time cells, congestion coefficient 0.02 and residual tolerance 10^{-4}. Its cell-centered marginals and endpoints form 33 saved times. Alternating times, including both endpoints, provide 17 DAM observations; the remaining 16 are excluded from DAM fitting and selection.

We convert finite-element momentum to minimum-norm triangle-edge flows preserving its divergence, then correct the finite ADMM continuity residual by a sparse graph projection. Conversion

requires residual TV at most 0.005 and an edge-correction ℓ^1 norm at most 10^{-4} plus 5% of the original flow norm on each interval. An area-reference mixture of 0.001 ensures positive sampling probabilities, and ODE verification checks accumulated sampling TV against 0.0005. This produces paired vertex Markov trajectories without identifying continuous surface characteristics. DAM receives no teacher velocities; all rollouts start from the same training empirical marginal and propagate without resets.

DAM training. Both variants use logarithmic mobility, smoothstep interpolation and reference mixing 0.01. Histogram DAM interpolates training histograms before taking logarithms and accesses this stored path at inference. Learned-ratio DAM combines local and uniformly proposed global-pair logistic losses with equal weights. Both train the potential using paired integration-by-parts samples, preserving the population objective.

Networks use 32-dimensional node embeddings, width-64 hidden layers, scalar time input and a zero-mean potential gauge. Ratio and action training use 60 and 100 epochs, respectively, with 200 updates per epoch. In the same order, learning rates are $2.2846 \cdot 10^{-3}$ and $9.4677 \cdot 10^{-4}$, weight decays $1.1117 \cdot 10^{-5}$ and $3.7445 \cdot 10^{-8}$, batch sizes 512 and 128, and clipping norms 0.25935 and 3.39333. The variants share potential architectures, initialization and budgets. Hand models reuse sphere settings with fresh weights.

Checkpoints minimize the respective simulation-free validation objectives on retained-time interpolants, using eight Gauss–Legendre nodes per interval and a final order-16 check. Validation uses independent samples; action losses under different frozen ratio models are not ranked against one another.

FEGF adaptation and selection. Surface FEGF uses a time-indexed scalar potential table, empirical geodesic targets and linear potential interpolation, with no trained entropy coefficient. Separate searches complete 93 sphere trials and 102 hand trials, with at most 200 epochs each. Every five epochs, selection evaluates TV at the 16 omitted validation times using 16 rollout substeps and 30-epoch patience. Thus FEGF and DAM use different validation information, geometries and tuning budgets; test data are excluded from selection.

Selected sphere and hand learning rates are 0.0587579 and 0.0582306, weight decays $2.87534 \cdot 10^{-8}$ and 0.000138841, batch sizes 256 and 128, clipping norms 1.75120 and 0.138872, updates per epoch 400 and 200, and mass floors 0.000144722 and 0.167839, respectively.

Rollouts and visualization. Step doubling must change predictions by at most 0.001 in TV, checking numerical consistency rather than absolute integration error. DAM uses 32 midpoint-frozen substeps per interval; tuned FEGF uses 1024 on sphere and 32 on hand, while fixed FEGF uses 256 on both. The accuracy table averages the 16 omitted times; diagnostics and costs use all 32 noninitial times.

Figures 7 and 8 show probability divided by normalized vertex area, with face colors averaging vertex densities. Each mesh uses a fixed camera and logarithmic color scale shared across methods and times. Quarter-time marginals are interpolated for display only; metrics use the original times.

D.7 Mesh estimator and integration diagnostics

In a cumulative sphere study, paired integration-by-parts training improves rollout accuracy. Adding global-pair ratio estimation improves density and endpoint errors but worsens mean trajectory TV. Better density estimates therefore do not uniformly improve rollouts. For the combined model, midpoint freezing reduces the substeps needed to meet the 0.001 refinement tolerance from 128 to 32, with little change in accuracy. These single-seed comparisons do not isolate independent causal effects, and numerical refinement does not control sampling or fitting error.

D.8 Accuracy versus computational cost

Figure 9 reports CPU times on an Intel Xeon Gold 6338 using float64 arithmetic and one thread. Medians and ranges cover three repetitions of configurations using seed 42, so whiskers measure runtime variation.

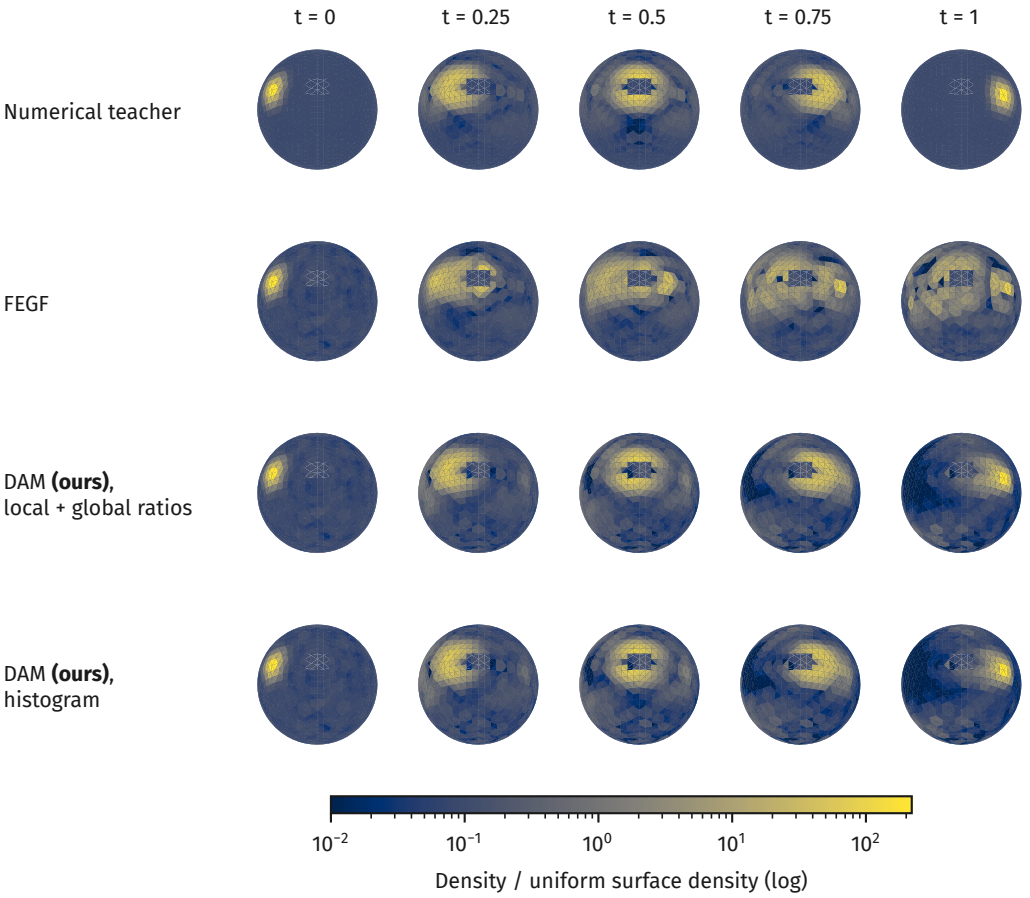


Figure 7: Sphere transport, seed 42. Rows: numerical teacher, tuned FEGF, DAM with local + global ratios, and histogram DAM. Columns: $t = 0, 0.25, 0.5, 0.75, 1$. Colors show area-normalized probability density.

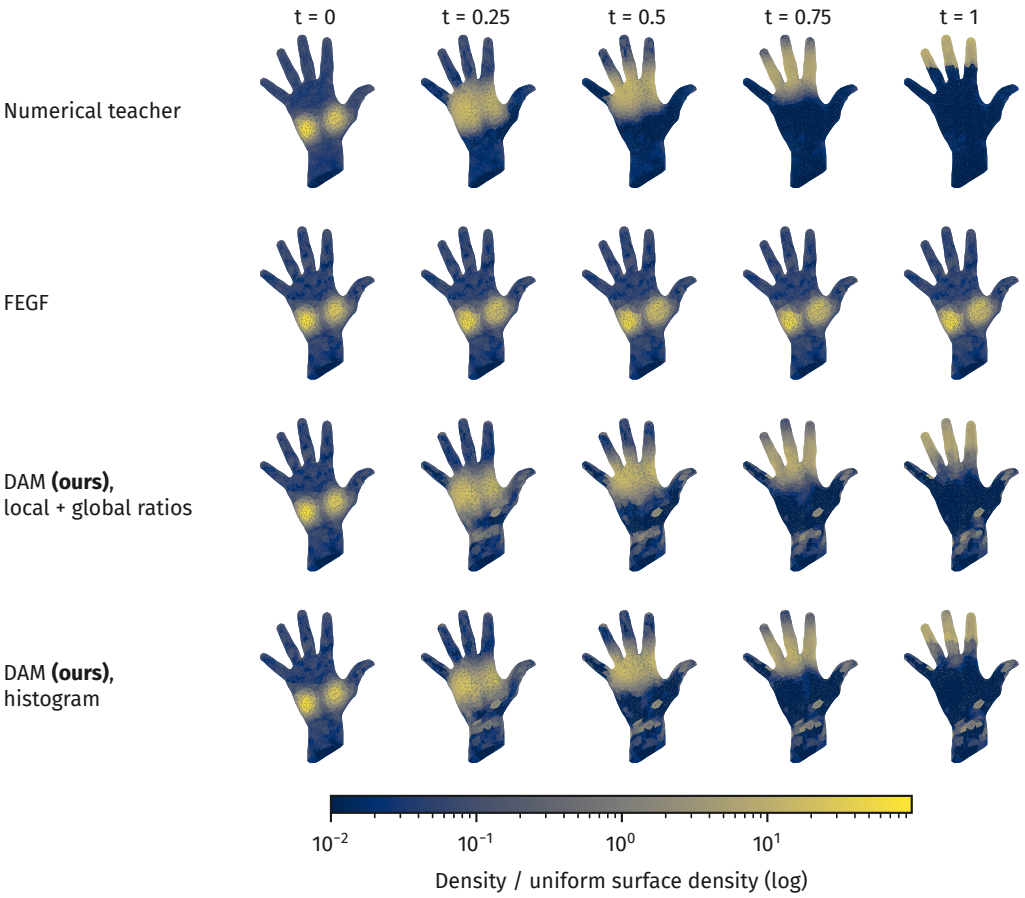


Figure 8: Hand transport, seed 42, using the row order, display times and density-rendering convention of Figure 7.

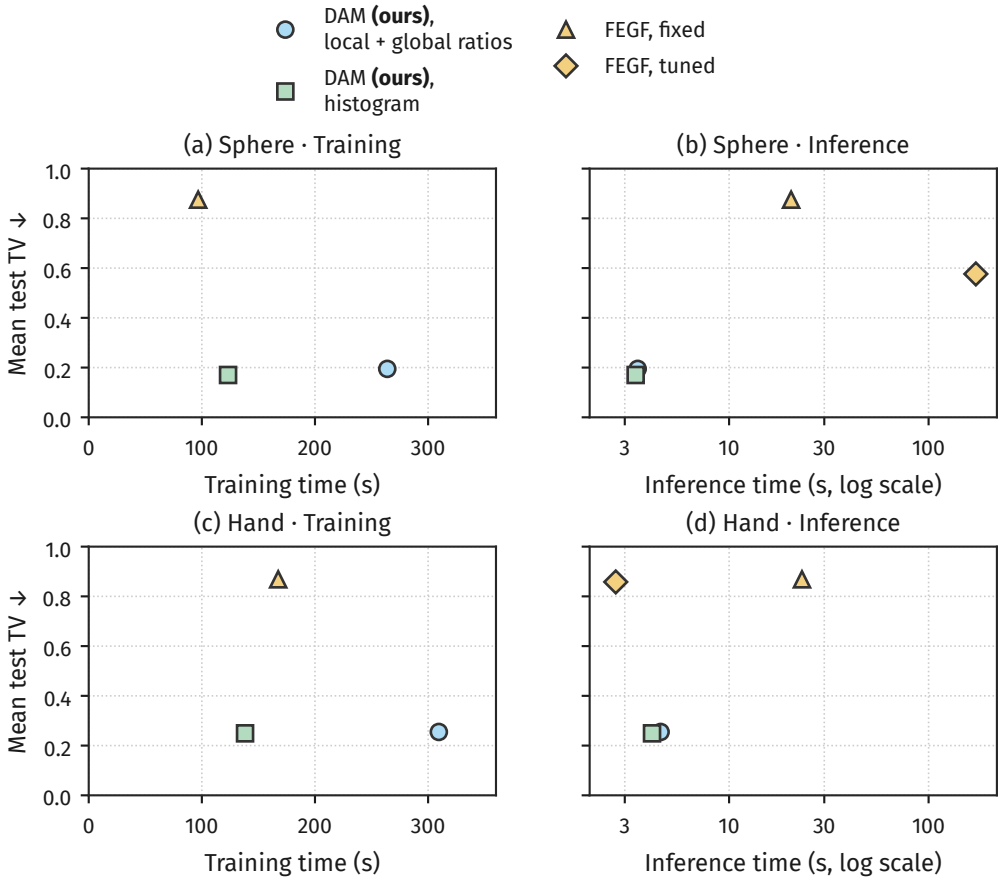


Figure 9: Accuracy versus CPU cost on sphere (top) and hand (bottom): training (left) and marginal inference (right, logarithmic time axis). TV averages 32 noninitial times. Whiskers span three runtime repetitions around median times. Tuned FEGF is inference-only.

DAM training includes setup and potential fitting, plus ratio fitting for the learned variant. Fixed FEGF includes target preparation, optimization and validation but excludes model setup. Both exclude teacher generation, particle-data creation and hyperparameter search. Tuned FEGF appears only for inference because its training time was not measured. Inference covers 33 output times, excluding model loading and output serialization.

Training exceeds a single teacher solve in these runs. Since the methods have unequal errors, histogram DAM accesses stored intermediate distributions, and endpoint-constrained transport differs from evaluating a learned path, these timings do not establish an equal-accuracy or end-to-end speedup on new transport problems.