Entering the Era of Discrete Diffusion Models: A Benchmark for Schrödinger Bridges and Entropic Optimal Transport

---

Xavier Aramayo Carrasco[1] **Grigoriy Ksenofontov**[1,2]
Aleksei Leonov[2,3] **Iaroslav Koshelev**[3] **Alexander Korotin**[1,4]

[1]Applied AI Institute
[2]Moscow Independent Research Institute of Artificial Intelligence
[3]AI Foundation and Algorithm Lab
[4]AXXX

2026

Background and Motivation

Background: EOT and SB Problems

Given p_0, p_1 and a reference process q^{ref} on a discrete space $\mathcal{X}$ we have following **equivalent** problems:

Dynamic Schrödinger Bridge (SB) problem

Seeks a **stochastic process** $q(x_0, x_{\text{in}}, x_1)$ via

$$\min_{q \in \Pi_N(p_0, p_1)} \text{KL}\left(q(x_0, x_{\text{in}}, x_1) \| q^{\text{ref}}(x_0, x_{\text{in}}, x_1)\right).$$

Here, $\Pi_N(p_0, p_1)$ is the set of **stochastic processes** with marginals p_0, p_1 at $t=0$, $t=1$.

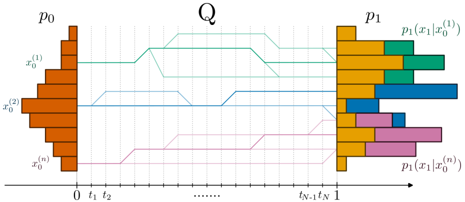

Static Schrödinger Bridge (SB) and Entropic Optimal Transport (EOT) problems

Seeks a **joint distribution** $q(x_0, x_1)$ via

$$\min_{q \in \Pi(p_0, p_1)} \text{KL}\left(q(x_0, x_1) \| q^{\text{ref}}(x_0, x_1)\right),$$

$$\min_{q \in \Pi(p_0, p_1)} \mathbb{E}_{q(x_0, x_1)} \underbrace{\left[-\log q^{\text{ref}}(x_1|x_0)\right]}_{\text{cost function}} - H(q(x_0, x_1)).$$

Here $\Pi(p_0, p_1)$ is the set of **joint distributions** with marginals p_0, p_1.

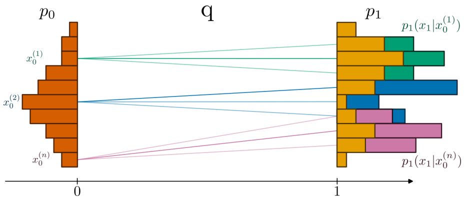

Motivation: Discrete and Continuous spaces

In continuous spaces, the state space is $\mathcal{X} = \mathbb{R}^D$ with real-valued variables, while in discrete spaces, $\mathcal{X} = \mathbb{S}^D$ with each component taking values in a finite set $\mathbb{S} = \{0, \dots, S-1\}$.

Continuous Domain

EOT/SB methods are well-established:

- Many solvers exist, e.g., DSBM, DSB;
- Evaluation protocols are standardized.

Example: Image Style Translation

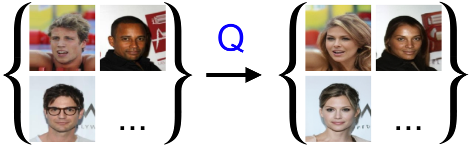

Discrete Domain

EOT/SB methods are still under development:

- CSBM is the only existing solver;
- Lack of standard evaluation protocols.

Example: Molecule Property Translation

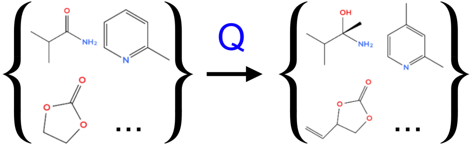

Goal: Establish the first benchmark for EOT/SB solvers in discrete spaces and expand the selection of available solvers.

Benchmark Construction

Benchmark Idea

Goal: Construct an analytically known ground-truth joint distribution $q^*(x_0, x_1)$ such that we can sample $(x_0, x_1) \sim q^*(x_0, x_1)$ where $x_0 \sim p_0$ and $x_1 \sim p_1$ are marginals.

Key idea is to choose p_0 and construct $q^*(x_1|x_0)$, instead of fixing (p_0, p_1) and solving EOT/SB directly. This defines $q^*(x_0, x_1) = p_0(x_0)q^*(x_1|x_0)$ through the following theorem.

Theorem (Benchmark Pair Construction for Discrete-Space EOT/SB)

Let $p_0 \in \mathcal{P}(\mathcal{X})$ be a given initial distribution on a discrete space $\mathcal{X}$ and $v^* : \mathcal{X} \to \mathbb{R}$ be a given scalar-valued function. Consider a joint distribution $q^* \in \mathcal{P}(\mathcal{X}^2)$ such that $q^*(x_0) = p_0(x_0)$ and $q^*(x_1|x_0) \propto v^*(x_1)q^{\text{ref}}(x_1|x_0)$ define $p_1(x_1) := q^*(x_1)$ as its second marginal. Then q^* together with q^{ref} defines the discrete-space EOT/SB between p_0 and p_1.

In general, sampling from $q^*(x_1|x_0)$ is computationally hard, since the discrete state space grows exponentially with dimension.

Canonical Polyadic Decomposition and Factorizable Reference Processes

The tractability of $q^*(x_1|x_0)$ relies on two following key ingredients:

We parametrize v^* via **the Canonical Polyadic (CP) decomposition**:

$$v^*(x_1) = \sum_{k=1}^{K} \beta_k \prod_{d=1}^{D} r_k^d[x_1^d].$$

Universal approximation: On a finite space, CP approximates $v^*(x_1)$ arbitrarily well as K grows.

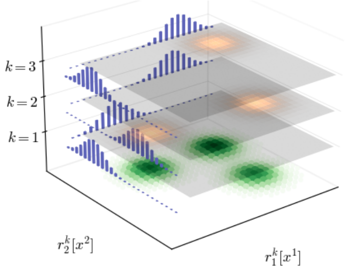

We use **factorizable** reference processes q^{ref}:

$$q^{\text{ref}}(x_1|x_0) = \prod_{d=1}^{D} q^{\text{ref}}(x_1^d|x_0^d).$$

- **Gaussian q^{ref}:** favors nearby transitions; used for ordered spaces.
- **Uniform q^{ref}:** assigns equal mass; used when no structure is assumed.

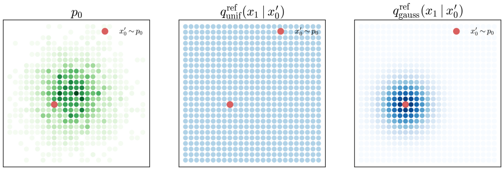

Resulting Parameterization of q^*

Given CP decomposition of v^* and factorizable q^{ref}, we obtain the following tractable forms.

Conditional (one-step) distribution

For static SB and EOT problems we have:

$$q^*(x_1|x_0) = \frac{1}{c^*(x_0)} \sum_{k=1}^{K} \beta_k \prod_{d=1}^{D} \left[r_k^d[x_1^d]\, q^{\text{ref}}(x_1^d|x_0^d) \right],$$

$$c^*(x_0) = \sum_{k=1}^{K} \beta_k \prod_{d=1}^{D} \left(\sum_{x_1^d} r_k^d[x_1^d]\, q^{\text{ref}}(x_1^d|x_0^d) \right).$$

Transition (multi-step) distributions

For dynamic SB problem we have:

$$q^*(x_{t_n}|x_{t_{n-1}}) \propto q^{\text{ref}}(x_{t_n}|x_{t_{n-1}}) \sum_{k=1}^{K} \beta_k \prod_{d=1}^{D} u_{k,t_n}^d[x_{t_n}^d],$$

$$u_{k,t_n}^d[x^d] = \sum_{x_1^d} q^{\text{ref}}(x_1^d|x^d)\, r_k^d[x_1^d].$$

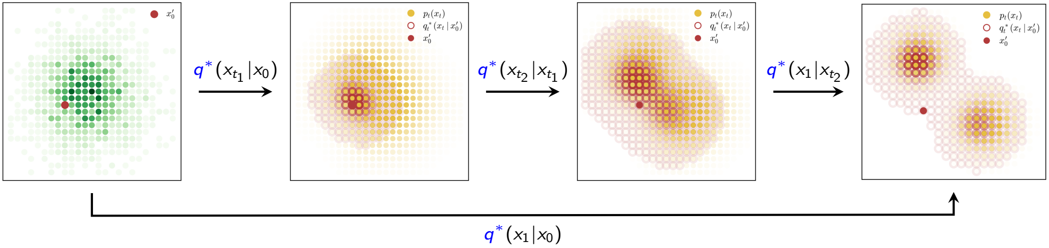

High-Dimensional Gaussian Mixture Benchmark Construction

The following setup is used to construct benchmark.

Parameter	Value
Number of dimensions (D)	$\{2, 16, 64\}$
Number of categories (S)	50
Number of CP components (K)	5
Gaussian reference (γ)	$\{0.02, 0.05\}$
Uniform reference (γ)	$\{0.005, 0.01\}$
Number of timesteps ($N+1$)	128

- **Initial distribution (p_0):** discretized $\mathcal{N}(0, I)$.
- **Component weights (β_k):** sampled from $\mathcal{U}([0, 1])$.
- **CP cores (r_k):** discretized $\mathcal{N}(\mu, \sigma I)$ with μ sampled on a sphere of radius 5 and $\sigma \in \{1.5, 1.5, 2.5\}$.

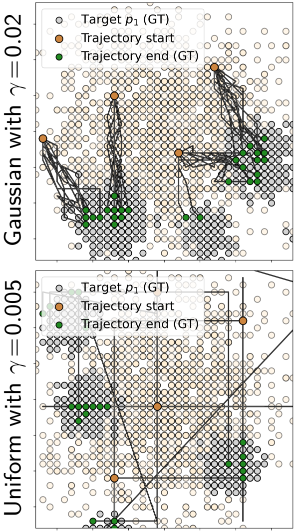

Discrete-Space EOT/SB Solvers

Discrete-Space EOT/SB Solvers

Existing Solvers

Categorical Schrödinger Bridge Matching (CSBM)

CSBM uses the fact that the SB is both Markov and reciprocal:

1. Start from an initial process $q^0(x_0, x_1)q^{\text{ref}}(x_{\text{in}}|x_0, x_1)$.
2. At iteration l, first project onto the Markov set, $q^{2l+1} = \text{proj}_{\mathcal{M}}(q^{2l})$, and then onto the reciprocal set, $q^{2l+2} = \text{proj}_{\mathcal{R}^{\text{ref}}}(q^{2l+1})$.
3. The sequence $\{q^l\}$ converges to the SB q^*.

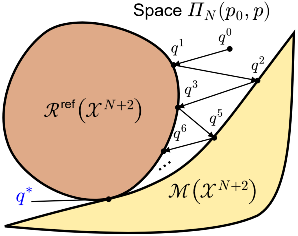

Loss. To perform projections, $proj_{\mathcal{R}^{\text{ref}}}(proj_{\mathcal{M}}(\cdot))$, the following objective is minimized:

$$\arg\min_{\theta} \mathcal{L}(\theta) = \arg\min_{\theta} \mathbb{E}_{q^{2l+1}(x_0,x_1)}\Bigg[\sum_{n=1}^{N} \mathbb{E}_{q^{\text{ref}}(x_{t_{n-1}}|x_0,x_1)} \text{KL}\big(q^{\text{ref}}(x_{t_n}|x_{t_{n-1}}, x_1) \,\|\, q_\theta(x_{t_n}|x_{t_{n-1}})\big) - \mathbb{E}_{q^{\text{ref}}(x_{t_N}|x_0,x_1)}[\log q_\theta(x_1|x_{t_N})]\Bigg].$$

Note: In practice, CSBM is implemented bidirectionally, alternating forward/backward model updates at each double projection. Moreover, KL can be replaced with MSE.

Discrete-Space EOT/SB Solvers

New Solvers

α-Categorical Schrödinger Bridge Matching (α-CSBM)

Problem. Bidirectional CSBM improves performance, but adds the overhead of training separate forward/backward models.

Solution. α-CSBM relaxes CSBM by replacing exact alternating projections with relaxed updates:

1. Start from an initial process $q^0(x_0, x_1)\, q^{\text{ref}}(x_{\text{in}} \mid x_0, x_1)$.
2. At iteration l, take a relaxed double projection:
$$q^{l+1} = (1-\alpha) q^l + \alpha \ \text{proj}_{\mathcal{R}^{\text{ref}}}\big(\text{proj}_{\mathcal{M}}(q^l)\big).$$
3. The sequence $\{q^l\}$ converges to the SB q^*.

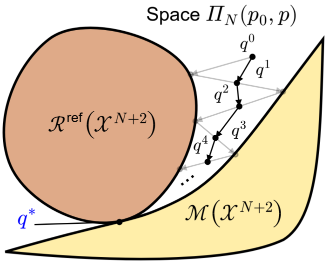

Loss. Since only a relaxed step is needed, forward/backward models are jointly optimized[1]:

$$\arg\min_{\theta} \mathcal{L}(\theta) = \arg\min_{\theta} \tfrac{1}{2}\Big(\text{KL}\left(\overrightarrow{r_{\text{sg}}}(x_0, x_{\text{in}}, x_1) \| \overleftarrow{q_\theta}(x_0, x_{\text{in}}, x_1)\right) + \text{KL}\left(\overleftarrow{r_{\text{sg}}}(x_0, x_{\text{in}}, x_1) \| \overrightarrow{q_\theta}(x_0, x_{\text{in}}, x_1)\right)\Big),$$

where $\rightleftarrows$ denote Markov representation directions[2], and $r_{\text{sg}} = \text{proj}_{\mathcal{R}^{\text{ref}}}(q_\theta)$ with stop-gradient.

[1] The KL terms in α-CSBM correspond to the loss used in CSBM.

[2] Direction is typically implemented by conditioning a neural network on a direction variable.

Discrete Light Schrödinger Bridge (Matching) (DLightSB(-M))

Both DLightSB(-M) solvers use the same benchmark parameterization for v_θ, with learnable parameters $\theta = \{\beta_k, r_k^d\}$, ensuring that q_θ lies in the set of all SBs.

Under this parametrization, solving the SB problem reduces to matching only the target marginal p_1, which can be achieved via a single projection.

Target: $q_\theta(x_1|x_0)$

Loss. The projection is defined as $\mathrm{KL}\left(q^* \| q_\theta\right)$ and minimized via the surrogate objective:

$$\arg\min_\theta \mathbb{E}_{p_0(x_0)}\left[\log c_\theta(x_0)\right] - \\ - \mathbb{E}_{p_1(x_1)}\left[\log v_\theta(x_1)\right].$$

Target: $q_\theta(x_{t_n}|x_{t_{n-1}})$

Loss. The projection of reciprocal process r defined as $\mathrm{KL}\left(r \| q_\theta\right)$ and minimized via:

$$\arg\min_\theta \mathbb{E}_{r(x_0,x_1)}\Bigg[\sum_{n=1}^{N} \mathbb{E}_{q^{\mathrm{ref}}(x_{t_{n-1}}|x_0,x_1)} \\ \mathrm{KL}\left(q^{\mathrm{ref}}(x_{t_n}|x_{t_{n-1}}, x_1) \,\|\, q_\theta(x_{t_n}|x_{t_{n-1}})\right) - \\ - \mathbb{E}_{q^{\mathrm{ref}}(x_{t_N}|x_0,x_1)}[\log q_\theta(x_1|x_{t_N})]\Bigg].$$

Evaluation of Solvers

Baselines and Metrics

Baselines

Independent: sample x_1 directly from p_1; ignores joint distribution.

Reference: sample x_1 from q^{ref}; follows prior transport.

Feature-wise SB: solve a 1D SB independently for each dimension using analytical CSBM and empirical marginals p_0, p_1; ignores cross-dimensional dependence.

Metrics

Shape Score (SSM) and **Trend Score (TSM)** are computed by averaging over d and (d_m, d_n), respectively.

$$\text{SSM}^d = 1 - \frac{1}{2} \sum_{x^d=0}^{S-1} \left| \tilde{q}^*(x^d) - \tilde{q}_\theta(x^d) \right|,$$

$$\text{TSM}^{d_m,d_n} = 1 - \frac{1}{2} \sum_{x^{d_m}=0}^{S-1} \sum_{x^{d_n}=0}^{S-1} \left| \tilde{q}^*(x^{d_m}, x^{d_n}) - \tilde{q}_\theta(x^{d_m}, x^{d_n}) \right|$$

where $\tilde{q}(x) = \frac{1}{|I|} \sum_{i \in I} \delta_{x^{(i)}}$ is the empirical distribution.

Conditional variants. We average SSM and TSM over sampled $x_0 \sim p_0$ and multiple generated x_1 per x_0.

Qualitative Results on the Conditional Trend Score

Method	Loss	$N+1$	$D=2$ gaussian 0.02	$D=2$ gaussian 0.05	$D=2$ uniform 0.005	$D=2$ uniform 0.01	$D=16$ gaussian 0.02	$D=16$ gaussian 0.05	$D=16$ uniform 0.005	$D=16$ uniform 0.01	$D=64$ gaussian 0.02	$D=64$ gaussian 0.05	$D=64$ uniform 0.005	$D=64$ uniform 0.01
Independent	–	–	0.47	0.78	0.52	0.55	0.57	0.68	0.37	0.46	0.48	0.51	0.35	0.43
Reference	–	–	0.09	0.28	0.23	0.27	0.11	0.08	0.09	0.10	0.20	0.12	0.14	0.11
Feature-wise SB	–	–	0.59	0.63	0.70	0.74	0.69	0.53	0.35	0.37	0.84	0.40	0.53	0.46
DLightSB	–	–	**0.91**	**0.85**	**0.87**	**0.87**	0.76	**0.84**	**0.84**	**0.84**	**0.85**	**0.75**	**0.82**	**0.77**
CSBM	KL	16	0.64	0.61	0.80	0.81	0.67	0.61	0.52	0.51	0.76	0.69	0.52	0.54
		64	<u>0.84</u>	0.79	0.84	0.85	0.80	0.74	0.57	0.62	<u>0.79</u>	0.70	0.49	0.54
	MSE	16	0.44	0.63	0.69	0.73	0.66	0.57	0.39	0.42	0.59	0.66	0.45	0.46
		64	0.23	0.75	0.69	0.67	0.73	0.72	0.48	0.55	0.50	<u>0.71</u>	0.48	0.48
α-CSBM	KL	16	0.59	0.62	0.82	0.81	0.61	0.66	0.55	0.56	0.77	0.66	0.48	0.52
		64	0.73	0.79	0.83	0.84	**0.82**	0.80	0.56	0.64	<u>0.79</u>	0.68	0.49	0.51
	MSE	16	0.47	0.61	0.77	0.78	0.60	0.65	0.45	0.51	0.59	0.62	0.45	0.46
		64	<u>0.84</u>	0.78	0.76	0.79	0.59	0.76	0.49	0.56	0.49	0.69	0.47	0.49
DLightSB-M	KL	16	0.82	**0.85**	<u>0.86</u>	<u>0.86</u>	0.70	0.82	<u>0.83</u>	<u>0.82</u>	0.73	0.62	0.47	0.50
		64	0.80	<u>0.84</u>	0.85	<u>0.86</u>	<u>0.79</u>	<u>0.83</u>	<u>0.83</u>	<u>0.82</u>	0.73	<u>0.71</u>	0.49	<u>0.67</u>
	MSE	16	0.62	<u>0.84</u>	0.76	0.84	0.45	0.82	0.78	0.79	0.53	0.68	<u>0.65</u>	0.33
		64	0.60	0.83	0.72	0.82	0.40	0.81	0.79	0.76	0.39	0.68	0.33	0.60

- *Reference* degrades because v^* strongly reweights the reference process.
- *Independent* worsens at low stochasticity.
- *Feature-wise SB* degrades as dimension increases.

These failures highlight the key challenges that our benchmarks are designed to test in EOT/SB solvers.

[2] Color code: vermillion for < 0.5, orange for $[0.5, 0.75)$, yellow for $[0.75, 0.85)$, and bluish-green for ≥ 0.85.

Qualitative Results on the Conditional Trend Score

Method	Loss	$N+1$	$D=2$ gaussian 0.02	$D=2$ gaussian 0.05	$D=2$ uniform 0.005	$D=2$ uniform 0.01	$D=16$ gaussian 0.02	$D=16$ gaussian 0.05	$D=16$ uniform 0.005	$D=16$ uniform 0.01	$D=64$ gaussian 0.02	$D=64$ gaussian 0.05	$D=64$ uniform 0.005	$D=64$ uniform 0.01
Independent	–	–	0.47	0.78	0.52	0.55	0.57	0.68	0.37	0.46	0.48	0.51	0.35	0.43
Reference	–	–	0.09	0.28	0.23	0.27	0.11	0.08	0.09	0.10	0.20	0.12	0.14	0.11
Feature-wise SB	–	–	0.59	0.63	0.70	0.74	0.69	0.53	0.35	0.37	0.84	0.40	0.53	0.46
DLightSB	–	–	**0.91**	**0.85**	**0.87**	**0.87**	0.76	**0.84**	**0.84**	**0.84**	**0.85**	**0.75**	**0.82**	**0.77**
CSBM	KL	16	0.64	0.61	0.80	0.81	0.67	0.61	0.52	0.51	0.76	0.69	0.52	0.54
		64	0.84	0.79	0.84	0.85	0.80	0.74	0.57	0.62	0.79	0.70	0.49	0.54
	MSE	16	0.44	0.63	0.69	0.73	0.66	0.57	0.39	0.42	0.59	0.66	0.45	0.46
		64	0.23	0.75	0.69	0.67	0.73	0.72	0.48	0.55	0.50	0.71	0.48	0.48
α-CSBM	KL	16	0.59	0.62	0.82	0.81	0.61	0.66	0.55	0.56	0.77	0.66	0.48	0.52
		64	0.73	0.79	0.83	0.84	**0.82**	0.80	0.56	0.64	0.79	0.68	0.49	0.51
	MSE	16	0.47	0.61	0.77	0.78	0.60	0.65	0.45	0.51	0.59	0.62	0.45	0.46
		64	0.84	0.78	0.76	0.79	0.59	0.76	0.49	0.56	0.49	0.69	0.47	0.49
DLightSB-M	KL	16	0.82	**0.85**	0.86	0.86	0.70	0.82	0.83	0.82	0.73	0.62	0.47	0.50
		64	0.80	0.84	0.85	0.86	0.79	0.83	0.83	0.82	0.73	0.71	0.49	0.67
	MSE	16	0.62	0.84	0.76	0.84	0.45	0.82	0.78	0.79	0.53	0.68	0.65	0.33
		64	0.60	0.83	0.72	0.82	0.40	0.81	0.79	0.76	0.39	0.68	0.33	0.60

- DLightSB achieves the strongest performance across all setups.
- DLightSB-M attains comparable results, with a slight drop likely due to variance from the KL minimization loss.
- Both methods benefit from the inductive bias used to construct the benchmark.

DLightSB(-M) can be viewed as oracle-like in this benchmark.

[2] Color code: vermillion for < 0.5, orange for $[0.5, 0.75)$, yellow for $[0.75, 0.85)$, and bluish-green for ≥ 0.85.

Qualitative Results on the Conditional Trend Score

			$D=2$				$D=16$				$D=64$			
			gaussian		uniform		gaussian		uniform		gaussian		uniform	
Method	Loss	$N+1$	0.02	0.05	0.005	0.01	0.02	0.05	0.005	0.01	0.02	0.05	0.005	0.01
Independent	–	–	0.47	0.78	0.52	0.55	0.57	0.68	0.37	0.46	0.48	0.51	0.35	0.43
Reference	–	–	0.09	0.28	0.23	0.27	0.11	0.08	0.09	0.10	0.20	0.12	0.14	0.11
Feature-wise SB	–	–	0.59	0.63	0.70	0.74	0.69	0.53	0.35	0.37	0.84	0.40	0.53	0.46
DLightSB	–	–	**0.91**	**0.85**	**0.87**	**0.87**	0.76	**0.84**	**0.84**	**0.84**	**0.85**	**0.75**	**0.82**	**0.77**
CSBM	KL	16	0.64	0.61	0.80	0.81	0.67	0.61	0.52	0.51	0.76	0.69	0.52	0.54
		64	<u>0.84</u>	0.79	0.84	0.85	0.80	0.74	0.57	0.62	<u>0.79</u>	0.70	0.49	0.54
	MSE	16	0.44	0.63	0.69	0.73	0.66	0.57	0.39	0.42	0.59	0.66	0.45	0.46
		64	0.23	0.75	0.69	0.67	0.73	0.72	0.48	0.55	0.50	<u>0.71</u>	0.48	0.48
α-CSBM	KL	16	0.59	0.62	0.82	0.81	0.61	0.66	0.55	0.56	0.77	0.66	0.48	0.52
		64	0.73	0.79	0.83	0.84	**0.82**	0.80	0.56	0.64	<u>0.79</u>	0.68	0.49	0.51
	MSE	16	0.47	0.61	0.77	0.78	0.60	0.65	0.45	0.51	0.59	0.62	0.45	0.46
		64	<u>0.84</u>	0.78	0.76	0.79	0.59	0.76	0.49	0.56	0.49	0.69	0.47	0.49
DLightSB-M	KL	16	0.82	**0.85**	<u>0.86</u>	<u>0.86</u>	0.70	0.82	<u>0.83</u>	<u>0.82</u>	0.73	0.62	0.47	0.50
		64	0.80	<u>0.84</u>	0.85	<u>0.86</u>	<u>0.79</u>	<u>0.83</u>	<u>0.83</u>	<u>0.82</u>	0.73	<u>0.71</u>	0.49	<u>0.67</u>
	MSE	16	0.62	<u>0.84</u>	0.76	0.84	0.45	0.82	0.78	0.79	0.53	0.68	<u>0.65</u>	0.33
		64	0.60	0.83	0.72	0.82	0.40	0.81	0.79	0.76	0.39	0.68	0.33	0.60

- CSBM and α-CSBM perform worse than DLightSB(-M).
- α-CSBM matches CSBM quality at lower computational cost.
- Increasing N generally improves the metrics.
- KL consistently outperforms MSE, likely because MSE produces over-smoothed solutions.

Our benchmark shows that further advances in discrete-space EOT/SB methods are still needed.

[2] Color code: vermillion for < 0.5, orange for $[0.5, 0.75)$, yellow for $[0.75, 0.85)$, and bluish-green for ≥ 0.85.

Thank you

Entering the Era of Discrete Diffusion Models:
A Benchmark for Schrödinger Bridges and Entropic Optimal Transport

A novel benchmark for discrete EOT/OT solvers.

GitHub

https://github.com/gregkseno/catsbench