Minimal-Action Discrete Schrödinger Bridge Matching for Peptide Sequence Design

Shrey Goel,[1] Pranam Chatterjee[2,†]

[1]Department of Computer Science, Duke University
[2]Department of Computer and Information Science, University of Pennsylvania
[3]Department of Bioengineering, University of Pennsylvania
[†]Correspondence to pranam@seas.upenn.edu

Duke

Penn

Motivation

Peptide sequence generation can be framed as a transport problem over discrete sequence space, but learning biologically plausible paths remains challenging:

- Discrete diffusion and autoregressive approaches rely on fixed corruption paths, forcing generation through **low-likelihood intermediate sequences**
- Classical Schrödinger bridge solvers require iterative projections, which are **computationally intractable** in high-dimensional, language-like biological sequences.
- Existing discrete SBM approaches **use a pre-trained base path measure, limiting their applicability** to de novo generation where no such reference process exists.

Discrete Schrödinger Bridge Problem

💡 **Find the stochastic process whose path measure minimizes the KL divergence between the path measures induced by the base and reference processes**

$$\mathbb{P}_{u^\star} = \arg\min_{\mathbb{P}_u \in \mathcal{P}} \left\{ \mathrm{KL}(\mathbb{P}_u \,\|\, \mathbb{P}_0) \mid (\mathbb{P}_u)_0 = \mu_0,\ (\mathbb{P}_u)_T = \mu_1 \right\}.$$

CTMC transition rates model transitions between discrete sequence states

$$\mathrm{KL}(\mathbb{P}_u \,\|\, \mathbb{P}_0) = \mathbb{E}_{\mathbb{P}_u}\left[\int_0^T \sum_{x' \neq X_t} \left(R_u(X_t, x') \log \frac{R_u(X_t, x')}{R_0(X_t, x')} - R_u(X_t, x') + R_0(X_t, x') \right) dt \right]$$

- We aim to learn a **control field** that tilts a given reference process
- **An exponential tilt of the base process represents the "minimal action"** required by the control field to yield the data distribution

$$R_{u,\theta}(x, x') = R_0(x, x')\, \exp\big(u_\theta(x, x', t)\big)$$

$$\longrightarrow \quad \log R_u(x, x') = \log R_0(x, x') + u_\theta(x, x')$$

💡 **Learning the log rates simplifies the learning problem by decoupling the reference process from the transition rates we aim to learn**

Training Algorithm

Log reparameterization simplifies learning the control field to maximizing the transition intensity toward the target token

$$\log p_\theta(\cdot \mid x_t, t) \propto \underbrace{u_\theta(x_t, \cdot, t)}_{\text{Learned tilt}} + \underbrace{(1-t)\, f_\phi(x_t)}_{\text{Reference } R_0}$$

💡 **Key insight:** use time-gated, pre-trained LM logits as reference rates

💡 **The loss function is equivalent to independently minimizing the cross-entropy objective over corrupted tokens**

$$\mathcal{L}(\theta) = -\sum_{i=1}^{L} \mathbf{1}_{\{x_t^{(i)} \neq x_1^{(i)}\}} \log p_\theta(x_1^{(i)} \mid x_t, t)$$

Define the **linear interpolant between marginals** as the trajectory connecting clean and noisy sequence states

$$p_t\big(x^{(i)} \mid x_1^{(i)}\big) = (1-t)\, \delta_{\mathcal{M}}(x^{(i)}) + t\, \delta_{x_1^{(i)}}(x^{(i)})$$

Unconditional Sequence Generation Quality

Steps	Model	PPL (↓)	pLDDT (↑)	% Invalid (↓)
$N = 32$	ESM	4.189 ± 6.531	—	89.130
	EvoFlow	8.192 ± 8.861	83.597 ± 10.092	0.000
	DPLM	7.887 ± 8.336	77.497 ± 9.651	6.957
	MadSBM	$\mathbf{6.415 \pm 7.837}$	$\mathbf{83.007 \pm 9.424}$	**0.000**
$N = 64$	ESM	3.974 ± 5.871	—	88.261
	EvoFlow	7.676 ± 9.221	$\mathbf{85.856 \pm 8.823}$	0.000
	DPLM	7.171 ± 8.491	79.453 ± 9.891	9.130
	MadSBM	$\mathbf{6.387 \pm 7.403}$	82.307 ± 9.603	**0.000**
$N = 128$	ESM	4.085 ± 6.185	—	88.261
	EvoFlow	7.304 ± 9.270	$\mathbf{86.742 \pm 8.687}$	0.000
	DPLM	6.864 ± 9.850	82.120 ± 11.266	6.957
	MadSBM	$\mathbf{6.445 \pm 7.898}$	83.921 ± 9.434	**0.000**
$N = 512$	ESM	5.486 ± 8.112	—	84.783
	EvoFlow	6.484 ± 9.246	$\mathbf{87.450 \pm 8.022}$	0.000
	DPLM	$\mathbf{5.884 \pm 9.619}$	84.456 ± 12.036	16.522
	MadSBM	6.447 ± 8.375	81.564 ± 9.755	**0.000**
Test Set		8.244 ± 6.090	76.617 ± 10.493	—

💡 **When given a restricted sampling budget at inference time, MadSBM consistently produces 1) higher-likelihood peptide sequences with plausible structures and 2) minimally invalid peptide sequences relative to SOTA discrete diffusion models like DPLM**

Discrete Classifier Guidance

	Binding Affinity (↑)			**Best ipTM** (↑)			**Docking Score** (kcal/mol) (↓)		
Target	Existing	Unconditional	Guided	Existing	Unconditional	Guided	Existing	Unconditional	Guided
5E1C	4.932	5.416	6.109	0.83	0.04	0.75	-4.3	-6.0	-6.3
4EZN	6.176	5.468	6.072	0.53	0.28	0.68	-4.1	-6.0	-6.8
1AYC	6.576	7.272	7.982	0.58	0.34	0.57	-5.3	-6.8	-7.5
5KRI	4.932	5.416	6.109	0.83	0.05	0.76	-3.5	-6.5	-6.7
3HVE	**	5.646	6.334	**	0.410	0.810	–	-6.3	-6.8
6A9P	**	5.288	5.917	**	0.100	0.190	–	-4.2	-4.3

- We augment MadSBM sampling with **Sampling-Importance-Resampling (SIR)** as a form of discrete, gradient-free classifier guidance
- This **improves the binding affinity of generated peptides** to targets with and without known binders

Ablation Study

	$\log R_0$	Test PPL (↓)
MadSBM	$(1-t) f_\phi(x_t)$	**4.503**
w/o gating	$f_\phi(x_t)$	4.987
& w/o ESM-2	0	4.750

Acknowledgments

We would like to thank the entire Chatterjee Lab for their guidance and support. We further thank Penn CETS for hardware support and Yinuo Zhang and Sophie Vincoff technical help.

Read the full paper!

Programmable Biology Group

ChatterjeeLab

Conclusions

Our results establish MadSBM as discrete Schrödinger bridge-based framework for discrete peptide generation that remains within high-likelihood regions of the sequence space. By learning a control field over a biologically informed reference process, MadSBM transports masked priors to realistic peptide distributions. Augmenting sampling with Sampling-Importance-Resampling (SIR) provides discrete, gradient-free classifier guidance, improving binding affinity across targets with and without known binders. Benchmarking against SOTA diffusion baselines (DPLM, EvoFlow), MadSBM demonstrates competitive or superior sample quality, particularly in low-step regimes. Overall, MadSBM offers a flexible and sample-efficient platform for programmable peptide design and therapeutic discovery.