

Unlocking Guidance for Discrete State-Space Diffusion and Flow Models


Hunter Nisonoff*, Junhao (Bear) Xiong*, Stephan Allenspach*, Jennifer Listgarten
University of California, Berkeley

ICLR (2025)


*Equal Contribution

<https://openreview.net/forum?id=XsgHI54yO7>

Conditional generation tasks are ubiquitous in the sciences



Designing plastic-degrading enzymes




SARS-CoV-2 Main Protease Inhibitor



DNA promoter engineering

Diffusion models can be guided to produce conditional samples



$$p(\mathbf{x} | y) = \frac{p(y | \mathbf{x})p(\mathbf{x})}{p(y)}$$

$$\nabla_{\mathbf{x}} \log p_t(\mathbf{x} | y) = \underbrace{\nabla_{\mathbf{x}} \log p_t(\mathbf{x})}_{\text{Unconditional model}} + \underbrace{\nabla_{\mathbf{x}} \log p_t(y | \mathbf{x})}_{\text{Predictor}}$$


P. Dhariwal* and A. Nichol*, NeurIPS (2021).

Can we extend guidance to discrete objects?

Protein Sequence

MYTWTGALITPCAAEESKLPINPLSNSLLRHH
YDTRCFDSTVTESDIRVEESIYQCCDLAPEEA
LTERLYIGGPLTNSKGQNCGYRRCRASGVLT
SCGNTLTCYLKATAACRAAKLQDCTMLVNGDD
LVVICESAGTQEDAAALRAFTEAMTRYSAAPP

Small molecule Graphs



DNA Sequences

TGGGATTCAGGTCTCGGAGCCGTGGTGGTGCT
GAACAACCGCGTATTGGCGAAGGCTGGAACC
GCGCGATTGGCCTGCATGATCCGACCGCGCAT
GCGGAATTATGGCGCTGCGCCAGGCGGGCCTG
GTGATGCAGAACTATCGCATTATTGATGCGAC



$$\nabla_x \log p(x | y) = \underbrace{\nabla_x \log p(x)}_{\text{Unconditional model}} + \underbrace{\nabla_x \log p(y | x)}_{\text{Predictor}}$$

Challenge: for **sequences, graphs, text**, etc., $\nabla_x \log p_\theta(x)$ not useful.

Consequence I: Standard diffusion/score doesn't work. (other work)


Consequence II: Cannot use guidance, $\nabla_x \log p_\phi(y|x)$. (our work)

Discrete time diffusion on discrete spaces




J. Ingraham et al. Nature 623.7989 (2023): 1070-1078.

Mutation-
based
noise
process



Austin et al. NeurIPS (2021).

Diffusion models can be guided to produce conditional samples



$$p(x | y) = \frac{p(y | x)p(x)}{p(y)}$$


$$\nabla_x \log p(x | y) = \underbrace{\nabla_x \log p(x)}_{\text{Unconditional model}} + \underbrace{\nabla_x \log p(y | x)}_{\text{Predictor}}$$

Guidance in discrete spaces seems intractable

$$\nabla_x \log p(x | y) = \underbrace{\nabla_x \log p(x)}_{\text{Unconditional model}} + \underbrace{\nabla_x \log p(y | x)}_{\text{Predictor}}$$

Bayes' theorem:

$$p(x_t | y, x_{t+\Delta t}) = \frac{p(y | x_t, x_{t+\Delta t}) p(x_t | x_{t+\Delta t})}{\sum_{x'_t} p(y | x'_t, x_{t+\Delta t}) p(x'_t | x_{t+\Delta t})}$$



denoising process toward $t = 0$


Guidance in discrete spaces seems intractable

~~$$\nabla_x \log p(x | y) = \underbrace{\nabla_x \log p(x)}_{\text{Unconditional model}} + \underbrace{\nabla_x \log p(y | x)}_{\text{Predictor}}$$~~

Bayes' theorem:

$$p(x_t | y, x_{t+\Delta t}) = \frac{p(y | x_t, x_{t+\Delta t}) p(x_t | x_{t+\Delta t})}{\sum_{x'_t} p(y | x'_t, x_{t+\Delta t}) p(x'_t | x_{t+\Delta t})}$$

Mutate with probability, $p \propto \Delta t$




Guidance in discrete spaces seems intractable

~~$$\nabla_x \log p(x | y) = \underbrace{\nabla_x \log p(x)}_{\text{Unconditional model}} + \underbrace{\nabla_x \log p(y | x)}_{\text{Predictor}}$$~~


Bayes' theorem:

$$p(x_t | y, x_{t+\Delta t}) = \frac{p(y | x_t, x_{t+\Delta t}) p(x_t | x_{t+\Delta t})}{\sum_{x'_t} p(y | x'_t, x_{t+\Delta t}) p(x'_t | x_{t+\Delta t})}$$

Mutate with probability, $p \propto \Delta t$




Diffusion models for discrete spaces with continuous-time Markov chains (CTMCs)



$$p(x_{t+dt} = \tilde{x} | x_t = x) = \delta_{x, \tilde{x}} + R_t(x, \tilde{x})dt$$


Campbell *et al.*, NeurIPS (2022).
Campbell*, J. Yim* *et al.*, ICML (2024).

Unlocking guidance in discrete spaces



Bayes' theorem:

$$p(x_t | y, x_{t+\Delta t}) = \frac{p(y | x_t, x_{t+\Delta t}) p(x_t | x_{t+\Delta t})}{\sum_{x'_t} p(y | x'_t, x_{t+\Delta t}) p(x'_t | x_{t+\Delta t})}$$



In continuous time diffusion:

- Only single dimensional changes can occur in infinitesimal time!
- This enables tractable guidance!

Guidance in continuous and discrete spaces

Predictor-guidance (PG)

**Continuous
state-space**

$$\nabla_{x_t} \log p^{(\gamma)}(x_t|y) = \nabla_{x_t} \log \overset{\text{unconditional}}{p(x_t)} + \gamma \nabla_{x_t} \log \overset{\text{predictor}}{p(y|x_t)}$$

J. Sohl-Dickstein *et al.*, ICML (2015).
P. Dhariwal* and A. Nichol*, NeurIPS (2021).
Y. Song *et al.*, ICLR (2021).

Predictor-free guidance (PFG)

$$\nabla_{x_t} \log p^{(\gamma)}(x_t|y) = \gamma \nabla_{x_t} \log \overset{\text{conditional}}{p(x_t|y)} + (1 - \gamma) \nabla_{x_t} \log \overset{\text{unconditional}}{p(x_t)}$$

J. Ho and T. Salimans, NeurIPS - Workshop (2021).

**Discrete
state-space**

$$\log R_t^{(\gamma)}(x, \tilde{x}|y) = \log \overset{\text{unconditional}}{R_t(x, \tilde{x})} + \gamma \left(\log \overset{\text{predictor}}{p(y|\tilde{x}, t)} - \log \overset{\text{predictor}}{p(y|x, t)} \right)$$

H.N., J.X., S.A. and J.L., ICLR (2025). [This work]

$$\log R_t^{(\gamma)}(x, \tilde{x}|y) = \gamma \log \overset{\text{conditional}}{R_t(x, \tilde{x}|y)} + (1 - \gamma) \log \overset{\text{unconditional}}{R_t(x, \tilde{x})}$$

H.N., J.X., S.A. and J.L., ICLR (2025). [This work]

Accurate approximation for further speed-up

- *Predictor-guidance* requires $O(D \times S)$ forward passes of the predictor model for each sampling step.
- e.g., state-space of CIFAR10:
 $D = 32 \times 32 \times 3, S = 256 \Rightarrow D \times S \approx 7.8 \times 10^5$
- Solution: **Taylor-Approximated Guidance (TAG)** requires one forward pass (but still $D \times S$ computations):

$$\log p^\phi(y|\tilde{x}, t) - \log p^\phi(y|x, t) \approx (\tilde{x} - x)^T \nabla_x \log p^\phi(y|x, t)$$


Bayes' theorem:

$$p(x_t|y, x_{t+\Delta t}) = \frac{p(y|x_t, x_{t+\Delta t})p(x_t|x_{t+\Delta t})}{\sum_{x'_t} p(y|x'_t, x_{t+\Delta t})p(x'_t|x_{t+\Delta t})}$$


Oops I Took A Gradient: Scalable Sampling for Discrete Distributions

Will Grathwohl^{1,2} Kevin Swersky² Milad Hashemi² David Duvenaud^{1,2} Chris J. Maddison¹


Conditional generation tasks in the sciences



Small molecules




Enhancer DNA



Proteins



Stability-conditioned inverse folding

$$P(\text{sequence} \mid \text{structure}, \Delta\Delta G \geq 0) \propto \underbrace{P(\text{sequence} \mid \text{structure})}_{\text{Inverse Folding (FMIF)}} \underbrace{P(\Delta\Delta G \geq 0 \mid \text{sequence}, \text{structure})}_{\text{Stability predictor}}$$




Stability-conditioned inverse folding

$$P(\text{sequence} \mid \text{structure}, \Delta\Delta G \geq 0) \propto \underbrace{P(\text{sequence} \mid \text{structure})}_{\text{Inverse Folding (FMIF)}} \underbrace{P(\Delta\Delta G \geq 0 \mid \text{sequence}, \text{structure})}_{\text{Stability predictor}}$$




Thank You!

a Unconditional Generation



b Predictor



c Guided Generation

