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=XsgHl54yO7



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







Can we extend guidance to discrete objects?

Protein Sequence

MYTWTGALITPCAAEESKLPINPLSNSLLRHH YDTRCFDSTVTESDIRVEESIYQCCDLAPEEA LTERLYIGGPLTNSKGQNCGYRRCRASGVLTT SCGNTLTCYLKATAACRAAKLQDCTMLVNGDD LVVICESAGTQEDAAALRAFTEAMTRYSAPPG



DNA Sequences

TGGGATTCAGGTCTCGGAGCCGTGGTGGTGCT GAACAACCGCGTGATTGGCGAAGGCTGGAACC GCGCGATTGGCCTGCATGATCCGACCGCGCAT GCGGAATTATGGCGCTGCGCCAGGCGGGCCTG GTGATGCAGAACTATCGCATTATTGATGCGAC

$$abla_x \log p(x \mid y) =
abla_x \log p(x) +
abla_x \log p(y \mid x)$$
Unconditional model 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.





Diffusion models can be guided to produce conditional samples





Guidance in discrete spaces seems intractable





Guidance in discrete spaces seems intractable





Guidance in discrete spaces seems intractable





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



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



Unlocking guidance in discrete spaces



Guidance in continuous and discrete spaces

Predictor-free guidance (PFG) Predictor-guidance (PG) unconditional conditional predictor unconditional Continuous $abla_{x_t} \log p^{(\gamma)}(x_t|y) =
abla_{x_t} \log p(x_t) + \gamma
abla_{x_t} \log p(y|x_t)$ $\nabla_{x_t} \log p^{(\gamma)}(x_t|y) = \gamma \nabla_{x_t} \log \frac{p(x_t|y)}{p(x_t|y)} + (1-\gamma) \nabla_{x_t} \log \frac{p(x_t)}{p(x_t)}$ state-space J. Sohl-Dickstein et al., ICML (2015). J. Ho and T. Salimans, NeurIPS - Workshop (2021). P. Dhariwal* and A. Nichol*, NeurIPS (2021). Y. Song *et al.*, ICLR (2021). unconditional conditional unconditional predictor Discrete $\log R_t^{(\gamma)}(x, \tilde{x}|y) = \log \frac{R_t(x, \tilde{x})}{R_t(x, \tilde{x})} + \gamma \Big(\log \frac{p(y|\tilde{x}, t)}{p(y|\tilde{x}, t)} - \log \frac{p(y|x, t)}{p(y|x, t)}\Big)$ $\log R_t^{(\gamma)}(x, \tilde{x}|y) = \gamma \log \frac{R_t(x, \tilde{x}|y)}{R_t(x, \tilde{x}|y)} + (1 - \gamma) \log \frac{R_t(x, \tilde{x})}{R_t(x, \tilde{x})}$ state-space H.N., J.X., S.A. and J.L., ICLR (2025). [This work] 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:
- $D = 32 \times 32 \times 3, S = 256 \Longrightarrow D \times S \approx 7.8 \times 10^5$
- Solution: Taylor-Approximated Guidance (TAG) requires one forward pass (but still D × 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 $^{12}\;$ Kevin Swersky $^2\;$ Milad Hashemi $^2\;$ David Duvenaud $^{12}\;$ Chris J. Maddison $^1\;$



Conditional generation tasks in the sciences



Small molecules



Proteins

Enhancer DNA



Stability-conditioned inverse folding

 $P(\text{sequence} \mid \text{structure}, \Delta\Delta G \geq 0) \propto P(\text{sequence} \mid \text{structure}) P(\Delta\Delta G \geq 0 \mid \text{sequence}, \text{structure})$







Stability-conditioned inverse folding

P(sequence | structure, $\Delta\Delta G \ge 0) \propto$ P(sequence | structure) P($\Delta\Delta G \ge 0$ | sequence, structure)





Thank You!



