

PLUG-AND-PLAY CONTROLLABLE GENERATION FOR DISCRETE MASKED MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

This article makes discrete masked models for the generative modeling of discrete data controllable. The goal is to generate samples of a discrete random variable that adheres to a posterior distribution, satisfies specific constraints, or optimizes a reward function. This methodological development enables broad applications across downstream tasks such as class-specific image generation and protein design. Existing approaches for controllable generation of masked models typically rely on task-specific fine-tuning or additional modifications, which can be inefficient and resource-intensive. To overcome these limitations, we propose a novel plug-and-play framework based on importance sampling that bypasses the need for training a conditional score. Our framework is agnostic to the choice of control criteria, requires no gradient information, and is well-suited for tasks such as posterior sampling, Bayesian inverse problems, and constrained generation. We demonstrate the effectiveness of our approach through extensive experiments, showcasing its versatility across multiple domains, including protein design.

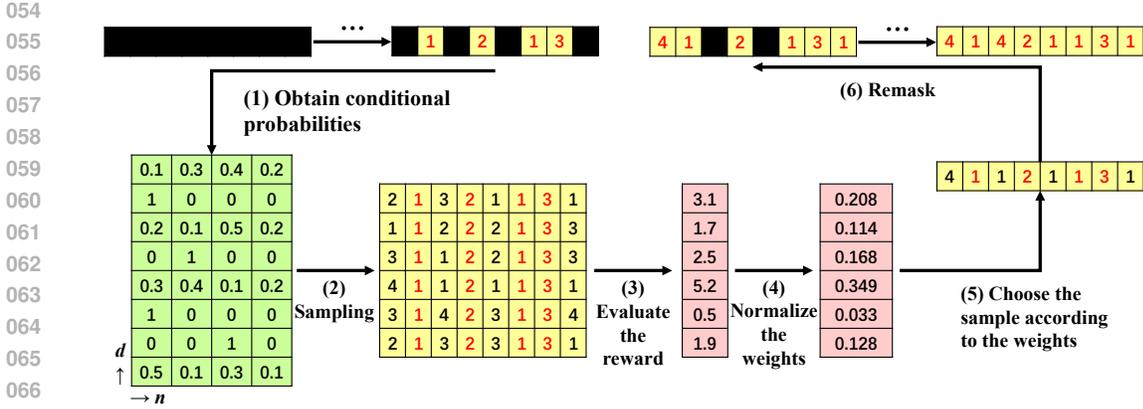
1 INTRODUCTION

Modeling complex discrete probability distributions in high-dimensional spaces is a crucial challenge across multiple domains in generative AI, including language, vision, audio, and biology. Among the various approaches, discrete masked models have emerged as powerful tools, offering robust solutions for generating and understanding discrete data. Notable examples include BERT for language modeling (Devlin et al., 2019), MaskGIT for image synthesis (Chang et al., 2022), DNABERT for DNA modeling (Ji et al., 2021; Zhou et al., 2023), the ESM series for protein generation (Rives et al., 2021; Lin et al., 2023; Hayes et al., 2024), and the more recent masked discrete diffusion models (see, e.g., Lou et al. (2024); Ou et al. (2024); Sahoo et al. (2024); Shi et al. (2024); Zheng et al. (2024)). These models typically learn the conditional distribution of each masked position given a partially masked data sequence, allowing for iterative decoding to generate a full sequence during inference.

In many practical applications of masked models, the objective extends beyond generating realistic samples from the data distribution to doing so in a *controlled* manner. This involves generating samples that align with specific constraints, conditions, or prompts, often by sampling from a conditional data distribution or maximizing a reward function (Zhang et al., 2023). Controlled generation is crucial in tasks such as (1) posterior sampling, where samples are drawn from a posterior distribution conditioned on observed data; (2) constrained generation, where the aim is to produce samples that meet predefined constraints. These applications highlight the growing need for generative models that not only capture the underlying data distribution, but also allow for flexible control over the generated outputs.

For controllable generation, it is desirable to establish a framework of *plug-and-play samplers* that allows for efficient sampling from the desired controlled distribution without requiring further training or fine-tuning of the underlying pre-trained model, given any suitable control criterion. This approach is computationally advantageous, as it avoids the costly and time-consuming process of retraining the model for each new task, making it a highly scalable and adaptable solution for real-world applications.

However, existing work on plug-and-play controllable generation primarily focused on the continuous domain, such as continuous diffusion models (Chung et al., 2023; Song et al., 2023; Huang et al.,



068 Figure 1: A demonstration of Algorithm 1 with vocabulary size $N = 4$, sequence length $D = 8$,
069 and number of Monte Carlo estimate $K = 6$.

070
071
072 2024), and often requires the differentiability in the control criteria. In contrast, controllable generation for discrete generative models often relies on learning-based approaches (Dathathri et al., 2020; Nisonoff et al., 2024; Li et al., 2024), and to the best of our knowledge, there is no plug-and-play sampler in this domain.

073
074
075 In this paper, we address this gap by developing a plug-and-play framework for controllable generation using discrete masked models, eliminating the need for task-specific fine-tuning. Our algorithm operates through iterative unmasking and remasking. At each step, given the unmasked positions, we apply a mean-field approximation to estimate the conditional distribution of the masked positions, sample from it via Monte Carlo, and employ importance sampling to filter the most likely samples. We then remask a portion of the newly generated positions, and repeat this unmasking-remasking process for several times until all positions are unmasked. Since the complexity of querying the masked model to obtain conditional probabilities is typically much higher than querying the reward function in most of the real-world applications, the Monte Carlo estimation introduces minimal computational overhead, keeping the sampling and filtering process efficient. In our experiments, high-quality samples can be obtained with approximately 10 queries to the masked model and around 1000 Monte Carlo samples, demonstrating the effectiveness of our proposed algorithm.

076
077
078
079
080
081
082
083
084
085
086
087 In summary, our contributions are as follows:

- 088
089
- 090 • We tackle the problem of plug-and-play controllable generation for discrete masked models, introducing an efficient and economical paradigm for sampling from these models.
 - 091
092
 - 093 • We propose a novel framework based on mean-field approximation of multi-dimensional discrete distributions and iterative masking-unmasking. This fine-tuning-free approach enables sample generation that satisfies control criteria for any (potentially non-differentiable) reward functions. To the best of our knowledge, this is the first plug-and-play controllable sampler for discrete masked models.
 - 094
095
096
097
 - 098 • We demonstrate the versatility of our method through multiple experiments, including sampling sequence of integers with equality constraint and designing protein sequences, highlighting its adaptability and effectiveness across diverse tasks.
- 099
100
101
102

103 **Notations.** The indicator function 1_A of a statement A is 1 if A is true and 0 if otherwise. We will use superscripts for indexing a vector, e.g., $x = (x^1, \dots, x^D)$. Given $x \in \mathbb{R}^D$ and a set $\Omega \subset \{1, \dots, D\}$, the slicing x^Ω is defined as $(x^d : d \in \Omega)$. Given a partition (Ω_1, Ω_2) of $\{1, \dots, D\}$ (i.e., Ω_1, Ω_2 are disjoint and their union is $\{1, \dots, D\}$), and two vectors $x = (x^d : d \in \Omega_1)$ and $y = (y^d : d \in \Omega_2)$, their concatenation $x \oplus y$ is a D -dimensional vector whose d -th entry is $x^d 1_{d \in \Omega_1} + y^d 1_{d \in \Omega_2}$. Finally, $f(x) \propto_x g(x)$ means that $f(x) = c \cdot g(x)$ for some constant $c > 0$ that does not depend on x .

104
105
106
107

2 PRELIMINARIES AND PROBLEM FORMULATION

2.1 DISCRETE MASKED MODELS

Discrete masked models (or simply **masked models**) are designed to learn the distribution of sequential discrete data. Let $\{1, \dots, N\}$ represent a vocabulary of N tokens, and consider a random sequence $X = (X^1, \dots, X^D) \in \{1, \dots, N\}^D$ of length D . For example, X can be a sentence with D words, a protein composed of D amino acids, or a discrete latent representation of images tokenized by a vector quantized variational autoencoder (VQ-VAE) (van den Oord et al., 2017). To learn the probability distribution p of X given i.i.d. samples of $X \sim p$, masked models are trained by randomly replacing certain positions in the sequence with a masked token \mathbf{M} , and learning the probability of the masked positions conditional on the unmasked portion of the sequence.

Throughout this paper, for a partially observed sequence $x \in \{1, \dots, N, \mathbf{M}\}^D$, we use $\Omega := \{d : x^d \neq \mathbf{M}\}$ and $\mathcal{M} := \{d : x^d = \mathbf{M}\}$ to denote the unmasked and masked positions, respectively. Formally, a masked model \hat{p} takes such a sequence as input and output a matrix $\hat{p}(x) \in \mathbb{R}_{\geq 0}^{D \times N}$ that approximates the following probability distribution:

$$\hat{p}(x)_{d,n} \approx p(X^d = n | X^\Omega = x^\Omega).$$

The rows in \hat{p} corresponding to the unmasked positions in Ω are trivial, as the probabilities are either 0 or 1. To learn the remaining entries, masked models are typically trained by minimizing the cross-entropy loss:

$$\min_{\hat{p}} \mathbb{E}_{X \sim p} \mathbb{E}_{\text{random subset } \Omega \subset \{1, \dots, D\}} \left(-\log \sum_{d \notin \Omega} \hat{p}(X^\Omega \oplus \mathbf{M})_{d, X^d} \right),$$

where $X^\Omega \oplus \mathbf{M}$ represents the sequence obtained by replacing all entries not in Ω with \mathbf{M} .

During inference, a direct approach is to initialize with the fully-masked sequence, select an order σ (a permutation of $\{1, \dots, D\}$), and autoregressively sample $X^{\sigma(t)}$ based on $X^{\sigma(<t)}$ for $t = 1, 2, \dots, D$. However, this scheme requires D queries to the masked model, which is inefficient for long sequences. To balance accuracy and efficiency, instead of unmasking one position at a time, one may introduce a decreasing function $\gamma : [0, 1] \rightarrow [0, 1]$ known as an **unmasking schedule**, which determines the number of remaining masked tokens at each step (Chang et al., 2022). At the t -th step ($t = 1, 2, \dots, T$) out of T total steps, given x with observed positions Ω , the following steps are implemented:

1. Predict the probability of the masked positions $p(X^d = \cdot | X^\Omega = x^\Omega) \approx \hat{p}(x)_{d,\cdot}$, $d \notin \Omega$ using the masked model.
2. Independently sample the masked values $x^d \sim \hat{p}(x)_{d,\cdot}$, $d \notin \Omega$.
3. Remask $\lfloor \gamma(t/T)D \rfloor$ newly-generated positions with the lowest predicted probability $\hat{p}(x)_{d,x^d}$.

Such steps will also be utilized in the design of our algorithm.

There are several equivalent formulation of masked models. First, as demonstrated in the sampling procedure outlined above, they can be interpreted as any-order autoregressive models (Hoogeboom et al., 2022; Shih et al., 2022), where the joint distribution is factorized using any arbitrary order σ , and the model learns the conditional distributions $p(X^{\sigma(d)} | X^{\sigma(<d)})$, $d \in \{1, \dots, D\}$. Second, recent research (Ou et al., 2024; Sahoo et al., 2024; Shi et al., 2024; Zheng et al., 2024) has shown that masked models are equivalent to the masked discrete diffusion model, which involves reversing a continuous-time Markov chain that transform any data distribution into the distribution concentrated on the sequence with all masked states.

2.2 CONTROLLABLE GENERATION

This paper focuses on the following task of **controllable generation** for masked models: given a pretrained masked model $\hat{p}(\cdot)$ that generates from a data distribution p , sample from a modified

distribution

$$q(x) = \frac{1}{Z} r(x) p(x), \quad x \in \{1, \dots, N\}^D,$$

where $r(x)$ is a non-negative **reward function**, and the normalizing constant $Z = \sum_x r(x) p(x)$ is unknown. This formulation encompasses various applications, including:

- **Posterior sampling for Bayesian inference:** Suppose $X \sim p(x)$ and we have a conditional distribution $p(y|x)$ for a related random variable Y , where $Y|X = x \sim p(y|x)$ serves as the reward function. Given an observation $Y = y$, the posterior distribution of X is $p(x|y) \propto_x p(y|x)p(x)$. For instance, if X is a tokenized image and Y is its corresponding class label given by a classifier, then sampling from $X|Y = y$ would generate an image belonging to a specific class y .
- **Constrained generation:** Given a specific subset $S \subset \{1, \dots, N\}^D$ of interest, we define the reward function as the indicator function $r(x) = 1_{x \in S}$. In this case, q represents the distribution p truncated to the set S . For example, X represents DNA sequences, and S is the set of DNA sequences whose percentage of A is less than 30%.

We also assume that evaluating the reward function $r(\cdot)$ is significantly less computationally expensive than querying the mask model $\hat{p}(\cdot)$, which is a common scenario in most of the real-world applications, and serves as an important starting point for our algorithm’s design.

3 CONTROLLABLE GENERATION FOR MASKED MODELS

In this section, we present our framework for plug-and-play controllable generation of masked models. We begin by drawing insights from the plug-and-play conditional generation for continuous diffusion models in Section 3.1, and then propose our novel algorithms tailored specifically for discrete masked models in Section 3.2.

3.1 EXISTING PLUG-AND-PLAY SAMPLERS FOR CONTINUOUS DIFFUSION MODELS

In continuous diffusion model, the data distribution is $X_0 \sim p(x_0)$, and a diffusion process $(X_t \sim p_t)_{t \in [0, T]}$ transforms data to noise following $X_t | X_0 = x_0 \sim \mathcal{N}(x_0, \sigma_t^2 I)$. Leveraging the data samples and the diffusion process, one can learn the score function $\nabla_{x_t} \log p_t(x_t)$ for all $t > 0$. Given a condition variable Y with distribution $p(y|x_0)$ conditional on $X_0 = x_0$, the task of sampling from the posterior distribution $p(x_0|y)$ boils down to estimating $\nabla_{x_t} \log p_t(y|x_t)$ for all $t > 0$, where $p_t(y|x_t)$ is the predicted distribution of Y given *noisy* sample $X_t = x_t$.

To estimate this quantity in a training-free way, Chung et al. (2023) observed that $p_t(y|x_t) = \mathbb{E}_{p(x_0|x_t)} p(y|x_0)$, and proposed to approximate the unknown distribution $p(x_0|x_t)$ by the point mass at $\hat{x}_0(x_t) := \mathbb{E}(X_0|X_t = x_t) = x_t + \sigma_t^2 \nabla_{x_t} \log p_t(x_t)$, resulting in $p_t(y|x_t) \approx p(y|\hat{x}_0(x_t))$. Song et al. (2023) later proposed a Gaussian approximation centered at $\hat{x}_0(x_t)$ with a suitable variance, estimating the expectation by the empirical mean over i.i.d. Gaussian samples. Finally, one step of backward propagation yields the gradient of $\log p(x_0|x_t)$ with respect to x_t . Throughout this process, only one query to the score model $\nabla_{x_t} \log p_t(x_t)$ is required, while the conditional probability $p(y|x_0)$ is generally easy to obtain, minimizing computational overhead.

To sum up, the key ingredient of this approach is the approximation of the distribution $p(x_0|x_t)$, which involves predicting the clean data X_0 given a noisy sample $X_t = x_t$.

3.2 CONTROLLABLE GENERATION FOR MASKED MODELS

We first continue the exploration of conditional generation, and then extend our methodology to general controllable generation problems.

In the discrete case, unlike the Gaussian noise $X_t|X_0 = x_0 \sim \mathcal{N}(x_0, \sigma_t^2 I)$, we now have a masking process $X_t|X_0 = x_0$ obtained by independently masking each position with a probability that depends on t . As the required entity is not the score function here, we cannot directly apply the continuous diffusion model’s strategy. Moreover, the discrete state space lacks a Gaussian distribution,

and the posterior mean $\mathbb{E}(X_0|X_t = x_t)$ is meaningless. However, a simple yet effective alternative exists: a mean-field approximation.

Our goal is to sample from $p(x_0|x_t, y)$. By conditional independence of X_t and Y given X_0 ,

$$\begin{aligned} p(x_0|x_t, y) &\propto_{x_0} p(x_0, x_t, y) = p(x_0)p(x_t|x_0)p(y|x_0) = p(x_0|x_t)p(x_t)p(y|x_0) \\ \implies p(x_0|x_t, y) &\propto_{x_0} \mathbb{E}_{p(x_0|x_t)} p(y|x_0). \end{aligned}$$

As X_t is by independently masking each position in X_0 , we only need to predict the conditional probability of the masked positions in x_t given the observed ones. As the masked model only predicts one-dimensional probability distributions, a straightforward approach is to iteratively unmasking one position at a time, requiring $|\mathcal{M}|$ queries to the masked model.

To avoid such computational overhead, we employ a **mean-field approximation**, i.e., assume that conditional on the observed part x^Ω of a sequence x , each remaining dimension in \mathcal{M} is independent. Formally,

$$p(X^\mathcal{M} = u|X^\Omega = x^\Omega) \approx \prod_{d \in \mathcal{M}} p(X^d = u^d|X^\Omega = x^\Omega),$$

which only requires one query to the masked model. While mean-field approximation introduces some error, it is effective for approximating multidimensional distributions when dependencies between different positions are relatively weak. Similar ideas have been incorporated to the sampling algorithm of MaskGIT (Chang et al., 2022) and the backward sampling of Continuous-Time Markov Chain (Lou et al., 2024; Ou et al., 2024; Zheng et al., 2024).

Building upon our insights into conditional generation, we are ready to present a solution to the general problem of controlled generation. Based on the problem settings in Section 2.2, we now illustrate our method for sampling from q in a plug-and-play manner.

Suppose we have a partially observed sequence x during the generation process, with observed and masked positions Ω and \mathcal{M} . We can calculate the conditional distribution of masked positions given the observed ones under q as follows: for all $u \in \{1, \dots, N\}^{|\mathcal{M}|}$,

$$\begin{aligned} q(X^\mathcal{M} = u|X^\Omega = x^\Omega) &\propto_u q(X^\mathcal{M} = u, X^\Omega = x^\Omega) \\ &\propto_u r(x^\Omega \oplus u)p(X^\mathcal{M} = u, X^\Omega = x^\Omega) \\ &\propto_u r(x^\Omega \oplus u)p(X^\mathcal{M} = u|X^\Omega = x^\Omega) \\ &\approx r(x^\Omega \oplus u) \prod_{d \in \mathcal{M}} p(X^d = u^d|X^\Omega = x^\Omega), \end{aligned} \quad (1)$$

where the last line is obtained by mean-field approximation. As the normalizing constant of this distribution is unknown, one way to sample a u from this distribution is by leveraging importance sampling. Let us first recall the following lemma.

Lemma 1 (Importance Sampling). *Suppose two probability masses or densities p, q are related through $q(x) = \frac{1}{Z}r(x)p(x)$. Then, with x_1, \dots, x_K i.i.d. samples from p , one can approximate q with the following weighted empirical distribution:*

$$q(x) \approx \tilde{q}(x) := \sum_{k=1}^K q_k \delta_{x_k}(x), \text{ where } q_k = \frac{r(x_k)}{\sum_{j=1}^K r(x_j)}.$$

Proof. Since $p(x) \approx \tilde{p}(x) := \frac{1}{K} \sum_{k=1}^K \delta_{x_k}(x)$, we have the following approximation of q :

$$q(x) = \frac{r(x)p(x)}{\mathbb{E}_p r} \approx \frac{r(x)\tilde{p}(x)}{\mathbb{E}_{\tilde{p}} r} = \frac{r(x) \sum_{k=1}^K \delta_{x_k}(x)}{\sum_{k=1}^K r(x_k)} = \tilde{q}(x).$$

□

Thanks to Lemma 1, we can approximately sample from the distribution in Equation (1) by independently sampling each dimension $d \in \mathcal{M}$ conditional on the observed sequence x^Ω and then obtain their corresponding weights, using which we can sample a u .

Algorithm 1: Plug-and-play controllable sampler of discrete masked models

Input: Vocabulary size N , length of sequence D , reward function $r(\cdot)$, masked model $\widehat{p}(\cdot)$, number of unmasking steps T , unmasking schedule $\gamma : [0, 1] \rightarrow [0, 1]$, number of Monte Carlo samples K .

Output: An approximate sample x from the distribution $q(x) \propto r(x)p(x)$.

- 1 Initialize $x = (\mathbf{M}, \dots, \mathbf{M})$, the mask positions $\mathcal{M} = \{1, \dots, D\}$, and the observed positions $\Omega = \emptyset$;
- 2 **for** $t = 1$ **to** T **do**
- 3 Compute the probabilities $\{\widehat{p}(x)_{d,n} \approx p(X^d = n | X^\Omega = x^\Omega) : d \in \mathcal{M}, n \in \{1, \dots, N\}\}$ from the masked model;
- 4 **for** $d \in \mathcal{M}, k \in \{1, \dots, K\}$ (in parallel) **do**
- 5 Independently sample $u_{(k)}^d \sim \widehat{p}(x)_{d,\cdot}$;
- 6 **end**
- 7 **for** $k \in \{1, \dots, K\}$ (in parallel) **do**
- 8 Assign the sample $u_{(k)} = (u_{(k)}^d : d \in \mathcal{M})$ with a weight $w_{(k)} = r(x^\Omega \oplus u_{(k)})$;
- 9 **end**
- 10 **for** $k \in \{1, \dots, K\}$ (in parallel) **do**
- 11 Normalize the weights $w_{(k)}$ by $w_{(k)} \leftarrow \frac{w_{(k)}}{\sum_{j=1}^K w_{(j)}}$;
- 12 **end**
- 13 Obtain one sample of $x^{\mathcal{M}}$ via selecting $u_{(k)}$ with probability $w_{(k)}$;
- 14 Remask $\lfloor \gamma(t/T)D \rfloor$ positions in $x^{\mathcal{M}}$ according to a defined rule (e.g., uniform remasking: sample a subset \mathcal{M}_0 of \mathcal{M} with $\lfloor \gamma(t/T)D \rfloor$ elements uniformly at random, and remask the positions $x^d, d \in \mathcal{M}_0$);
- 15 Update the masked positions $\mathcal{M} \leftarrow \{d : x^d = \mathbf{M}\}$ and the observed positions $\Omega \leftarrow \mathcal{M}^c$;
- 16 **end**
- 17 **return** x .

Although this process generates a full sequence x , its alignment to the target distribution q may be sub-optimal due to the lost of dependencies in mean-field approximation, especially when $|\mathcal{M}|$ is large (i.e., many positions are unmasked in a row). To address this, we further introduce a re-masking step as discussed in Section 2.1: remask some of positions in \mathcal{M} according to a remasking schedule γ (e.g., by sampling a random subset of \mathcal{M}), so that in the t -th step among the T total steps ($t = 1, \dots, T$), after remasking, the remaining number of masked tokens is $\lfloor \gamma(t/T)D \rfloor$. In our experiments, we found that this uniform remasking strategy performs well. However, exploring more advanced remasking strategies is a promising area for future research.

We summarize the entire procedure in Algorithm 1. Notably, our algorithm can be easily extended for the inpainting task, i.e., given a subset $\overline{\Omega}$ of $\{1, \dots, D\}$ that has known values $x^{\overline{\Omega}}$, we aim to sample the remaining positions according to $q(X^{\overline{\Omega}^c} = \cdot | X^{\overline{\Omega}} = x^{\overline{\Omega}})$. This modified version of the algorithm is provided in Algorithm 2.

4 RELATED WORK

Conditional generation based on guidance. As an important task in controllable generation, conditional generation aims to generate a random variable $X \sim p(x)$ (e.g., image) given another random variable Y known as the condition (e.g., the text description of an image). A popular approach is guidance: for example, in continuous diffusion model, the classifier guidance (Dhariwal & Nichol, 2021) learns a classifier $p(y|x)$ that predicts the condition Y given a *noisy* sample of X and leverages its information to generate X conditional on $Y = y$. Classifier-free guidance (Ho & Salimans, 2022) trains a score model that approximates both the conditional and unconditional score functions, using a combination of them for conditional generation. These approaches can be extended to the discrete diffusion model (see Nisonoff et al. (2024)).

Controllable generation for discrete generative models. The study of controllable generation is an emerging area in language modeling (see, e.g., Zhang et al. (2023) for a review). A notable work, Dathathri et al. (2020), proposed applying gradient updates to the key-value cache in transformers, a task-agnostic approach but requiring fine-tuning during inference. For diffusion models, a recent work Li et al. (2024) introduced soft value-based decoding, a derivative-free algorithm that requires pre-sampled trajectories x_0, x_1, \dots, x_T of discrete diffusion model to estimate a conditional expectation. This method does not exploit the special properties of the masking process. To the best of our knowledge, there are no fine-tuning-free samplers for controllable generation in discrete masked models.

5 EXPERIMENTAL RESULTS

5.1 TOY EXPERIMENT ON SAMPLING EQUALITY-CONSTRAINED SEQUENCES

We first apply our controllable generation framework to constrained sampling, demonstrating its outstanding efficiency in challenging tasks where the constraint set is extremely sparse.

Consider the state space $\{1, \dots, N\}^D$ comprising sequences of integers $x = (x^1, \dots, x^D)$, where we fix the sequence length D to 10. Let the unconditional distribution be the uniform distribution. In this case, the masked model has a closed form, requiring no training:

$$\hat{p}(x)_{d,n} = p(X^d = n | X^\Omega = x^\Omega) = \begin{cases} 1/N, & d \in \mathcal{M} \\ \delta_{x^d, n}, & d \in \Omega \end{cases}.$$

We study sampling from the uniform distribution restricted to the following set:

$$S_N = \{x \in \{1, \dots, N\}^D : \phi(x) := x^1 - x^2 x^3 - x^4 + x^5 x^6 x^7 + x^8 + x^9 - x^{10} = 0\}.$$

The ratio between the cardinalities of S_N and the entire state space $\{1, \dots, N\}^D$ is extremely low: approximately 1.07% when $N = 10$, 0.20% when $N = 20$, and 0.07% when $N = 30$.

We use Algorithm 1 to sample from the uniform distribution constrained on S_N , and present the results in Figure 2. We observe that the number of Monte Carlo samples, K , significantly impacts the quality of generated samples, while the number of steps of unmasking T has a less pronounced effect, likely due to the short sequence length. Further experimental details are provided in Appendix B.1.

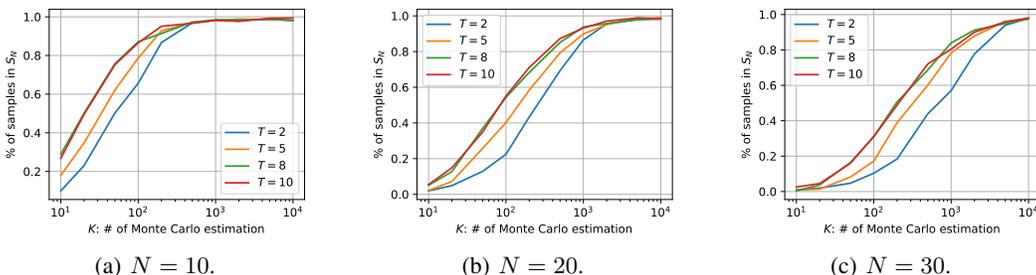


Figure 2: Result for sampling equality-constrained sequences.

5.2 CONTROLLABLE PROTEIN GENERATION

We employ ESM3 (Hayes et al., 2024) as the underlying protein generation model, which is a pioneering generative model for protein, capable of reasoning simultaneously over multiple modalities including sequence, structure, and function, achieving state-of-the-art performance in multiple protein generation tasks. In our experiments, we focus on generation in the *sequence* domain, and fix the length of sequence to 50. The vocabulary considered is the set of 20 standard amino acids¹. We will consider three tasks of controlled generation: generating hydrophilic and hydrophobic proteins,

¹They are: A, C, D, E, F, G, H, I, K, L, M, N, P, Q, R, S, T, V, W, Y.

sampling proteins with a propensity for alpha-helices, and protein inpainting for higher percentage of alpha-helices. In all tasks, we impose a constraint based on the instability index (Guruprasad et al., 1990), a positive number that estimates the stability of a protein sequence. A protein with value between 0 and 40 is predicted as stable. To evaluate all the required metrics of protein sequences, we use the Biopython package (Cock et al., 2009), a widely used tool in computational molecular biology.

Design of reward function. To begin with, we first introduce a flexible way of designing the reward function $r(\cdot)$ for a given task. Suppose we have M metrics $m_i : \{1, \dots, N\}^D \rightarrow \mathbb{R}, i = 1, 2, \dots, M$ that we are interested in, and our constraint is expressed as

$$\bigcap_{i=1}^M \{x : m_i(x) \in A_i\},$$

where $A_i \subset \mathbb{R}$ is an interval. We propose the following reward function:

$$r(x) = \exp\left(-\sum_{i=1}^M w_i \text{dist}(m_i(x), A_i)^{\alpha_i}\right),$$

where $\text{dist}(a, A)$ is the distance from $a \in \mathbb{R}$ to the interval $A \subset \mathbb{R}$, $w_i > 0$ is the weight of the i -th constraint ($m_i(x) \in A_i$), and $\alpha_i > 0$ determines the shape of penalty (e.g., linear or quadratic).

Sampling hydrophilic/hydrophobic proteins. Solubility is one of the key traits of protein and it is a joint consequence of amino acid sequence composition as well as 3D structure. Protein solubility is often quantified through the hydropathy index, with high hydropathy values indicating water-repelling (hydrophobic) and low hydropathy value indicating water-attracting (hydrophilic). Designing hydrophilic or hydrophobic proteins have numerous important applications, yet the tasks also pose unique challenges (Qing et al., 2022). Low hydropathy value is often desirable for therapeutic purposes as it’s central to the protein expression and purification process. Hydrophilicity is also critical for a longer circulation time in the human bloodstream, potentially indicating a better therapeutic efficacy (Garidel, 2013). Proteins with high hydropathy value are often designed for transmembrane proteins, important examples of which include cell surface receptors. Transmembrane proteins are often targets of drugs, especially receptors like G-protein-coupled receptors. Designing these proteins can help create more effective disease treatments by improving the capability to modulate cellular signals (Yang et al., 2021).

In the following, we consider the challenging task of generating hydrophilic and hydrophobic protein sequences. We set the length of protein sequence to 50, and use GRAVY (Grand Average of Hydropathy) value (Kyte & Doolittle, 1982) to quantify the hydropathy level of the generated sequence, which is defined as the average hydropathy value of a peptide or protein. Negative GRAVY value indicates hydrophilic, while positive one means hydrophobic.

We compare the GRAVY values of the samples generated by ESM3 model in both uncontrolled and controlled way in Table 1. For controlled generation, we fix the hyperparameters w_2 and A_2 for promoting low instability index, and experimented with multiple choices of w_1 and A_1 (see further details in Appendix B). The table presents the best generation result. Using our proposed controllable generation method, we successfully sample protein sequences with the desired values of GRAVY while maintaining protein stability. We also observe a significant reduction in standard deviation among the controlled generated samples, highlighting the reliability of our method.

Table 1: Controlled generation for high and low GRAVY values. The metrics are presented in the form of mean \pm std.

| Task | $m_1 = \text{GRAVY}$ | $m_2 = \text{instability}$ | GRAVY | Instability \downarrow |
|------------------------|---|--|--------------------|--------------------------|
| Uncontrolled | / | / | -0.207 ± 0.552 | 41.467 ± 17.906 |
| Controlled, high GRAVY | $w_1 = 30, \alpha_1 = 1, A_1 = [1, \infty)$ | $w_2 = 5, \alpha_2 = 2, A_2 = [0, 40]$ | 1.209 ± 0.223 | 25.176 ± 10.094 |
| Controlled, low GRAVY | $w_1 = 35, \alpha_1 = 1, A_1 = (-\infty, -1]$ | $w_2 = 5, \alpha_2 = 2, A_2 = [0, 40]$ | -1.261 ± 0.260 | 31.569 ± 8.204 |

Sampling alpha-helix-rich protein. Proteins fold naturally into unique three-dimensional structures based on their amino acid sequence composition. This spatial conformation further determines

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

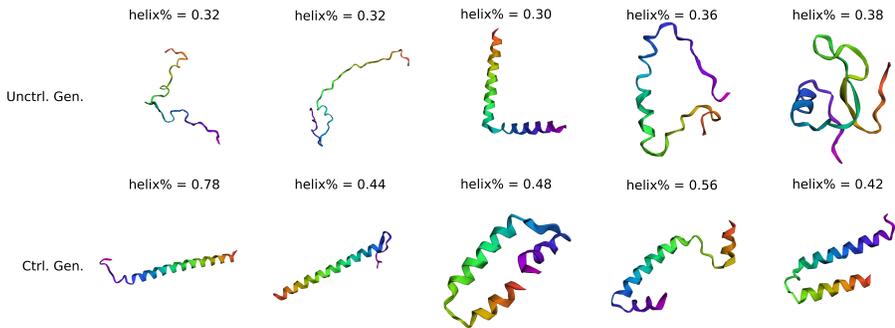


Figure 3: Structure of protein sequences predicted by ESM3. The upper row are the uncontrolled generated sequences, and the lower row are the controlled generated sequences. The sequences are randomly chosen.

its molecular and cellular function. Among the possible spatial structures, alpha helix is one of the most common secondary structures in protein, where the amino acids in the polypeptide chain form into a coil (helix) through hydrogen bonding. Particularly, designing proteins rich in alpha helices is of special interest for engineering certain protein functions (Kortemme, 2024), such as binding and self-assembly (Sakuma et al., 2024). These functions are central to applications such as therapeutic protein discovery (Walsh & Jefferis, 2006) and alpha-helical nanofiber design (Zhang, 2003).

Motivated by the versatile purpose of protein with rich alpha-helix structures, in the following, we focus on generating protein sequences with secondary structure being all or predominant alpha-helices. We compare the metrics of the generated proteins by our proposed method with those generated by the unconditional model in Table 2. Our controlled generation method successfully produces sequences with a high predicted percentage of alpha-helices (helix%). To verify the structural accuracy of the generated sequences, we use the folding algorithm provided by ESM3 to predict their 3D structures given the sequences. The result is displayed in Figure 3, where we randomly sample five generated sequences from both the uncontrolled and controlled generated sets and predict their structure. The controlled generated samples exhibit a higher frequency of alpha-helices in the predicted structure, confirming the effectiveness of our algorithm. For further details on the figures, please refer to Appendix B.3.

We also investigate the influence of the hyperparameters w_i and A_i in the design of reward function $r(\cdot)$. Our empirical findings suggest an optimal choice for these parameters, providing valuable insights into designing reward function for general controlled generation problems. Additional details are presented in Appendix B.2.

Table 2: Controlled generation for high alpha-helix percentage. The metrics are presented in the form of mean \pm std.

| Task | $m_1 = \text{helix\%}$ | $m_2 = \text{instability}$ | Helix% \uparrow | Instability \downarrow |
|--------------|---|--|-------------------|--------------------------|
| Uncontrolled | / | / | 0.315 ± 0.072 | 41.467 ± 17.906 |
| Controlled | $w_1 = 50, \alpha_1 = 1, A_1 = [0.8, \infty)$ | $w_2 = 5, \alpha_2 = 2, A_2 = [0, 40]$ | 0.600 ± 0.116 | 27.387 ± 12.025 |

Protein inpainting. We consider the following inpainting problem: using a protein chain with high beta-sheet percentage as a prompt to generate the remaining positions in a controlled manner to maximize the percentage of alpha-helices. In particular, the prompt is chosen as a 35-amino-acid slice from the SM-related protein of P. Abyssis (PDB ID: 1H64), chain A, which consists of 71 amino acids and has a secondary structure composed almost entirely of beta-sheets. The cyan motif in the subfigure wrapped in red box in Figure 4 highlights the prompt. We use Algorithm 2 to generate alpha-helix-rich proteins based on this prompt, and display the predicted 3D structure

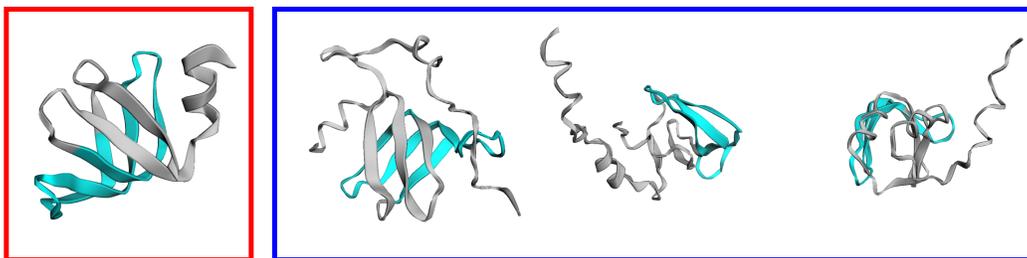


Figure 4: Visualization of the generated sequence from protein inpainting.

of three randomly generated sequences in the blue box in Figure 4, demonstrating the presence of alpha-helices in the generated portions. Further experimental details can be found in Appendix B.4.

6 CONCLUSION AND FUTURE WORK

In this paper, we study the task of controllable generation for discrete masked models and introduce an efficient paradigm for plug-and-play sampling. Our approach is based on mean-field approximation and iterative masking and remasking, demonstrating promising potential for real-world applications. Future research directions include exploring alternative remasking strategies beyond uniform remasking, rigorously analyzing the error bounds of our method, and extending its application to other domains such as image, molecule, and DNA.

REFERENCES

- Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. MaskGIT: Masked generative image transformer. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11305–11315, 2022. doi: 10.1109/CVPR52688.2022.01103.
- Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=OnD9zGAGT0k>.
- Peter JA Cock, Tiago Antao, Jeffrey T Chang, Brad A Chapman, Cymon J Cox, Andrew Dalke, Iddo Friedberg, Thomas Hamelryck, Frank Kauff, Bartek Wilczynski, et al. Biopython: freely available python tools for computational molecular biology and bioinformatics. *Bioinformatics*, 25(11):1422, 2009.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text generation. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1edEyBKDS>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=AAWuCvzaVt>.

- 540 Patrick Garidel. Protein solubility from a biochemical, physicochemical and colloidal perspective.
541 *Am Pharm Rev*, 2(5):26–28, 2013.
- 542
- 543 Kunchur Guruprasad, B.V.Bhasker Reddy, and Madhusudan W. Pandit. Correlation between stabil-
544 ity of a protein and its dipeptide composition: a novel approach for predicting in vivo stability of a
545 protein from its primary sequence. *Protein Engineering, Design and Selection*, 4(2):155–161, 12
546 1990. ISSN 1741-0126. doi: 10.1093/protein/4.2.155. URL [https://doi.org/10.1093/
547 protein/4.2.155](https://doi.org/10.1093/protein/4.2.155).
- 548 Tomas Hayes, Roshan Rao, Halil Akin, Nicholas J Sofroniew, Deniz Oktay, Zeming Lin, Robert
549 Verkuil, Vincent Q Tran, Jonathan Deaton, Marius Wiggert, et al. Simulating 500 million years
550 of evolution with a language model. *bioRxiv*, pp. 2024–07, 2024.
- 551 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint*
552 *arXiv:2207.12598*, 2022.
- 553
- 554 Emiel Hooeboom, Alexey A. Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, and
555 Tim Salimans. Autoregressive diffusion models. In *International Conference on Learning Rep-*
556 *resentations*, 2022. URL <https://openreview.net/forum?id=Lm8T39vLDTE>.
- 557 Yujia Huang, Adishree Ghatare, Yuanzhe Liu, Ziniu Hu, Qinsheng Zhang, Chandramouli Shama
558 Sastry, Siddharth Gururani, Sageev Oore, and Yisong Yue. Symbolic music generation with non-
559 differentiable rule guided diffusion. In *Forty-first International Conference on Machine Learning*,
560 2024. URL <https://openreview.net/forum?id=g8AigOTNXL>.
- 561 Yanrong Ji, Zhihan Zhou, Han Liu, and Ramana V Davuluri. DNABERT: pre-trained bidirectional
562 encoder representations from transformers model for dna-language in genome. *Bioinformatics*, 37
563 (15):2112–2120, 02 2021. ISSN 1367-4803. doi: 10.1093/bioinformatics/btab083. URL <https://doi.org/10.1093/bioinformatics/btab083>.
- 564
- 565 Tanja Kortemme. De novo protein design—from new structures to programmable functions. *Cell*,
566 187(3):526–544, 2024.
- 567
- 568 Jack Kyte and Russell F. Doolittle. A simple method for displaying the hydropathic character of
569 a protein. *Journal of Molecular Biology*, 157(1):105–132, 1982. ISSN 0022-2836. doi: [https://doi.org/10.1016/0022-2836\(82\)90515-0](https://doi.org/10.1016/0022-2836(82)90515-0). URL [https://www.sciencedirect.com/sc
570 ience/article/pii/0022283682905150](https://www.sciencedirect.com/science/article/pii/0022283682905150).
- 571
- 572 Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Bian-
573 calani, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-free guidance in continuous
574 and discrete diffusion models with soft value-based decoding. *arXiv preprint arXiv:2408.08252*,
575 2024.
- 576
- 577 Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin,
578 Robert Verkuil, Ori Kabeli, Yaniv Shmueli, Allan dos Santos Costa, Maryam Fazel-Zarandi, Tom
579 Sercu, Salvatore Candido, and Alexander Rives. Evolutionary-scale prediction of atomic-level
580 protein structure with a language model. *Science*, 379(6637):1123–1130, 2023. doi: 10.1126/sc
581 ience.ad2574. URL [https://www.science.org/doi/abs/10.1126/science.ad
582 e2574](https://www.science.org/doi/abs/10.1126/science.ad2574).
- 583 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios
584 of the data distribution. In *Forty-first International Conference on Machine Learning*, 2024. URL
585 <https://openreview.net/forum?id=CNicRIVIPA>.
- 586 Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking guidance
587 for discrete state-space diffusion and flow models. *arXiv preprint arXiv:2406.01572*, 2024.
- 588
- 589 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan
590 Li. Your absorbing discrete diffusion secretly models the conditional distributions of clean data.
591 *arXiv preprint arXiv:2406.03736*, 2024.
- 592 Rui Qing, Shilei Hao, Eva Smorodina, David Jin, Arthur Zalevsky, and Shuguang Zhang. Protein
593 design: From the aspect of water solubility and stability. *Chemical Reviews*, 122(18):14085–
14179, 2022.

- 594 Alexander Rives, Joshua Meier, Tom Sercu, Siddharth Goyal, Zeming Lin, Jason Liu, Demi Guo,
595 Myle Ott, C. Lawrence Zitnick, Jerry Ma, and Rob Fergus. Biological structure and function
596 emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of*
597 *the National Academy of Sciences*, 118(15):e2016239118, 2021. doi: 10.1073/pnas.2016239118.
598 URL <https://www.pnas.org/doi/abs/10.1073/pnas.2016239118>.
- 599 Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T
600 Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
601 models. *arXiv preprint arXiv:2406.07524*, 2024.
- 602 Koya Sakuma, Naohiro Kobayashi, Toshihiko Sugiki, Toshio Nagashima, Toshimichi Fujiwara,
603 Kano Suzuki, Naoya Kobayashi, Takeshi Murata, Takahiro Kosugi, Rie Tatsumi-Koga, et al.
604 Design of complicated all- α protein structures. *Nature Structural & Molecular Biology*, 31(2):
605 275–282, 2024.
- 606 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K Titsias. Simplified and gener-
607 alized masked diffusion for discrete data. *arXiv preprint arXiv:2406.04329*, 2024.
- 608 Andy Shih, Dorsa Sadigh, and Stefano Ermon. Training and inference on any-
609 order autoregressive models the right way. In S. Koyejo, S. Mohamed, A. Agar-
610 wal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Pro-*
611 *cessing Systems*, volume 35, pp. 2762–2775. Curran Associates, Inc., 2022. URL
612 https://proceedings.neurips.cc/paper_files/paper/2022/file/123fd8a56501194823c8e0dca00733df-Paper-Conference.pdf.
- 613 Jiaming Song, Qinsheng Zhang, Hongxu Yin, Morteza Mardani, Ming-Yu Liu, Jan Kautz, Yongxin
614 Chen, and Arash Vahdat. Loss-guided diffusion models for plug-and-play controllable generation.
615 In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and
616 Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*,
617 volume 202 of *Proceedings of Machine Learning Research*, pp. 32483–32498. PMLR, 23–29 Jul
618 2023. URL <https://proceedings.mlr.press/v202/song23k.html>.
- 619 Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning.
620 In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett
621 (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,
622 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/7a98af17e63a0ac09ce2e96d03992fbc-Paper.pdf.
- 623 Gary Walsh and Roy Jefferis. Post-translational modifications in the context of therapeutic proteins.
624 *Nature biotechnology*, 24(10):1241–1252, 2006.
- 625 Dehua Yang, Qingtong Zhou, Viktorija Labroska, Shanshan Qin, Sanaz Darbalaei, Yiran Wu,
626 Elita Yuliantie, Linshan Xie, Houchao Tao, Jianjun Cheng, et al. G protein-coupled receptors:
627 structure-and function-based drug discovery. *Signal transduction and targeted therapy*, 6(1):7,
628 2021.
- 629 Hanqing Zhang, Haolin Song, Shaoyu Li, Ming Zhou, and Dawei Song. A survey of controllable
630 text generation using transformer-based pre-trained language models. *ACM Comput. Surv.*, 56
631 (3), October 2023. ISSN 0360-0300. doi: 10.1145/3617680. URL <https://doi.org/10.1145/3617680>.
- 632 Shuguang Zhang. Fabrication of novel biomaterials through molecular self-assembly. *Nature*
633 *biotechnology*, 21(10):1171–1178, 2003.
- 634 Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Masked
635 diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical
636 sampling. *arXiv preprint arXiv:2409.02908*, 2024.
- 637 Zhihan Zhou, Yanrong Ji, Weijian Li, Pratik Dutta, Ramana Davuluri, and Han Liu. Dnabert-
638 2: Efficient foundation model and benchmark for multi-species genome. *arXiv preprint*
639 *arXiv:2306.15006*, 2023.
- 640
641
642
643
644
645
646
647

A ALGORITHM FOR INPAINTING

See Algorithm 2. Its difference with Algorithm 1 is marked in blue.

Algorithm 2: Plug-and-play controllable sampler of discrete masked models, the inpainting version.

Input: Vocabulary size N , length of sequence D , reward function $r(\cdot)$, masked model $\widehat{p}(\cdot)$, number of unmasking steps T , unmasking schedule $\gamma : [0, 1] \rightarrow [0, 1]$, number of Monte Carlo samples K , **inpainted indices** $\overline{\Omega} \subset \{1, \dots, D\}$, **inpainted values** $x^{\overline{\Omega}} \in \{1, \dots, N\}^{|\overline{\Omega}|}$.

Output: An approximate sample x from the distribution $q(x) \propto r(x)p(x)$.

```

1 Initialize  $x = x^{\overline{\Omega}} \oplus \mathbf{M}$ , the mask positions  $\mathcal{M} = \overline{\Omega}^c$ , and the observed positions  $\Omega = \overline{\Omega}$ ;
2 for  $t = 1$  to  $T$  do
3   Compute the probabilities  $\{\widehat{p}(x)_{d,n} \approx p(X^d = n | X^\Omega = x^\Omega) : d \in \mathcal{M}, n \in \{1, \dots, N\}\}$ 
4   for  $d \in \mathcal{M}, k \in \{1, \dots, K\}$  (in parallel) do
5     | Independently sample  $u_{(k)}^d \sim \widehat{p}(x)_{d,\cdot}$ ;
6   end
7   for  $k \in \{1, \dots, K\}$  (in parallel) do
8     | Assign the sample  $u_{(k)} = (u_{(k)}^d : d \in \mathcal{M})$  with a weight  $w_{(k)} = r(x^\Omega \oplus u_{(k)})$ ;
9   end
10  for  $k \in \{1, \dots, K\}$  (in parallel) do
11    | Normalize the weights  $w_{(k)}$  by  $w_{(k)} \leftarrow \frac{w_{(k)}}{\sum_{j=1}^K w_{(j)}}$ ;
12  end
13  Obtain one sample of  $x^{\mathcal{M}}$  via selecting  $u_{(k)}$  with probability  $w_{(k)}$ ;
14  Remask  $\lfloor \gamma(t/T)(N - |\overline{\Omega}|) \rfloor$  positions in  $x^{\mathcal{M}}$  according to a defined rule (e.g., uniform
15  remasking: sample a subset  $\mathcal{M}_0$  of  $\mathcal{M}$  with  $\lfloor \gamma(t/T)(N - |\overline{\Omega}|) \rfloor$  elements uniformly at
16  random, and remask the positions  $x^d, d \in \mathcal{M}_0$ ;
17  Update the masked positions  $\mathcal{M} \leftarrow \{d : x^d = \mathbf{M}\}$  and the observed positions  $\Omega \leftarrow \mathcal{M}^c$ ;
18 end
19 return  $x$ .

```

B SUPPLEMENTARY EXPERIMENTAL RESULTS

B.1 IMPLEMENTATION DETAILS OF THE TOY EXAMPLE

We first demonstrate how we choose the reward function $r(\cdot)$. Recall that the equality constraint is $\phi(x) = 0$. Thus, we propose to use the reward function $r(x) = \exp(-w \max(|\phi(x)|, m))$, where the weight $w > 0$ and the truncation threshold $m > 0$. We choose $w = 5$ and $m = 10$ throughout all of our experiments.

We fix the remasking schedule $\gamma(t) = \cos(\frac{\pi}{2}t)$ as suggested by Chang et al. (2022).

The values of K experimented in Figure 2 are 10, 20, 50, 100, 200, 500, 1000, 2000, 5000, 10000, while the values of T experimented are 2, 5, 8, 10.

B.2 TUNING THE HYPERPARAMETERS FOR CONTROLLED PROTEIN GENERATION

In this section, we study how the hyperparameters w_i and A_i in the reward function influences the quality of controllable generation. In particular, we focus on the task of sampling alpha-helix rich protein, fix the hyperparameters for $m_2 = \text{instability}$, and vary the choice of w_1, A_1 while fixing $\alpha_1 = 1$ throughout the experiments. We experiment $w_1 \in \{10, 15, \dots, 45, 50\}$ and $A_1 = [a_1, \infty)$ for $a_1 \in \{0.5, 0.6, \dots, 1.3, 1.4\}$, and for each pair of (w_1, A_1) , generate 16 protein sequences and

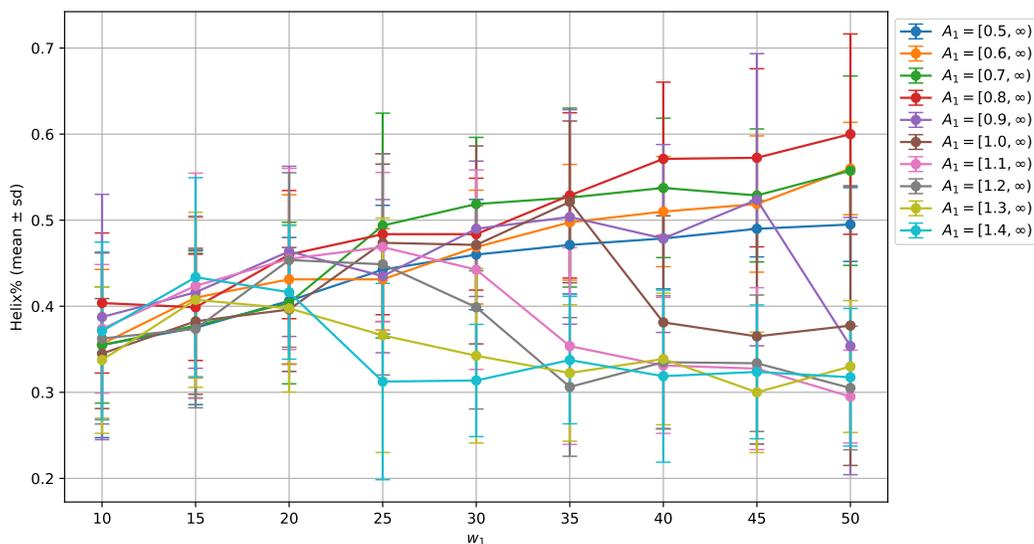


Figure 5: Influence of w_1 and A_1 on helix% in sampling alpha-helix rich proteins.

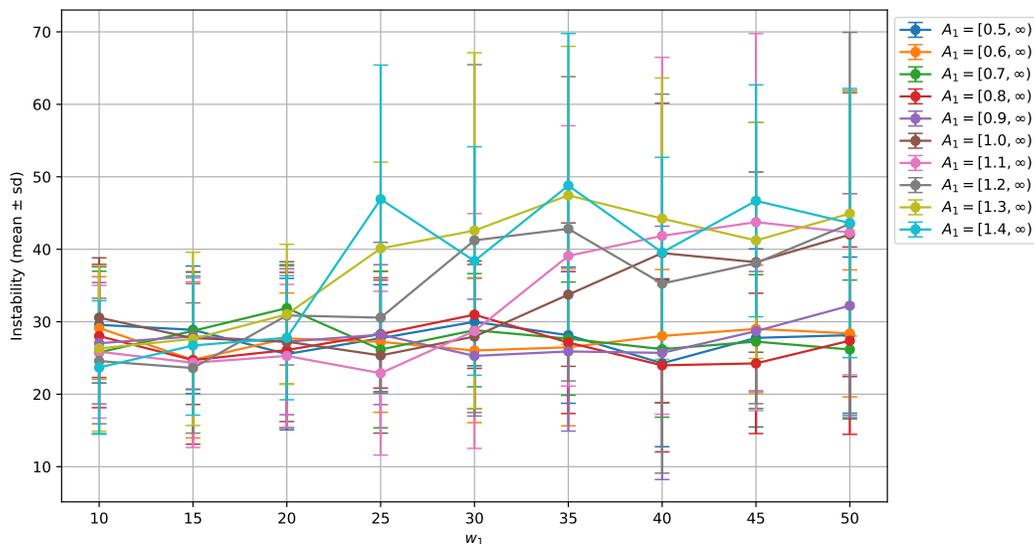


Figure 6: Influence of w_1 and A_1 on instability in sampling alpha-helix rich proteins.

evaluate the metrics helix% and instability. The results are displayed in Figures 5 and 6. We find that

- For a fixed a_1 , as w_1 grows, the helix% of the generated sequences would grow at first, but may decline when w_1 is larger than some threshold. This threshold becomes smaller when a_1 gets larger.
- For small a_1 , the instability index of the generated sequences does not vary significantly among different choices of w_1 . But for large a_1 , the generated sequences is prone to become more unstable when w_1 is large.

These results show that there may exist an optimal choice of the hyperparameters w_1 and A_1 that maximizes the helix% of the generated sequences while keeping low instability index, which may provide insights in designing the reward function $r(\cdot)$ for general controlled generation problems.

756 B.3 PROTEIN SEQUENCES IN FIGURE 3.
757758 From left to right, the uncontrolled generated sequences are:
759

- 760 1. RAGPRAPPRSDAGRTRGVGRKGQLLVGTGKLDAPTLLSLPAAVKSTGATRS
- 761 2. MGFPNVPATAAPCPAAPTYEDYAAAARGGSLPQVIQHALPVIFTAPLRKST
- 762 3. MNQQSTADIRMLIEIGSFMNDPNMMTLINLLILSNVFIILLIVIIYYRWRS
- 763 4. MKFLARSTAKTEQLRERYLKTIDIQILVYETIQGDFESIRLLPASVYNVSL
- 764 5. MLPSPAFAISEAQATVESGSIAGPELLAVAVEAPSTQDHRVFAGEETYGV

766 From left to right, the controlled generated sequences are
767

- 768 1. MINAEAADKDECRLADLLEAKELEMLELKALYLRLEENKALKELARAMA
- 769 2. TACVEKPTHGNPTLHLAAKAFNAEIIIDLAF LGQKREKLTQSNLRVISEK
- 770 3. MNNDLWVWWSKSKGLINKDKYKNLDNTIMYMKQNMKDIKELKGLLETILNA
- 771 4. MNDQTREKLLKPGAAEVFAKKYRREKEAIESRAIARVADIDEALKLAAQL
- 772 5. MLAEPIGNIVTYAYVILSILLLVKGLAENMETSVALTTLLFSNIWQLR

773
774
775 These sequence are randomly sampled from the batch of generated sequences and are not cherry-
776 picked.
777778 B.4 MORE DETAILS OF PROTEIN INPAINTING
779780 The whole sequence of the protein is ERPLDVIHRSLDKDVLVILKKGFEFRGRLIGYDIHLNVVLA
781 DAEMIQDGEVVKRYGKIVIRGDNVLAISPT, where the cyan part with 35 amino acids is the prompt
782 we use for inpainting. We choose the same hyperparameters as in Table 2. From left to right, the
783 three sequences displayed in the blue box in Figure 4 are

- 784 1. MSLAVLKNSEDTLVKAELKGDVSVRGRLIGYDIHLNVVLA DAEMIQDGEVVKRYGKIVIR
- 785 GDSVVTVHLLTALESQIHEIEDEKAKADRAVKARTKAIKA
- 786 2. MEGIALKALMDFQVVMKLGKELRGRLIGYDIHLNVVLA DAEMIQDGEVVKRYGKIVIR
- 787 GTVITLIHIPEEVDFEAALKLLEKKPKKRIRRLKAEKSKK
- 788 3. SMSLAMQNLGKEMKIRLAGGMCMRGRLIGYDIHLNVVLA DAEMIQDGEVVKRYGKIVIR
- 789 GNCIVYLDLPDSLKDELQSHERVHQYRGLKGAHAVKEKKR

791
792 C CODES FOR ALGORITHMS 1 AND 2
793

```

794 1 import tqdm
795 2 import torch
796 3 import numpy as np
797 4
798 5
798 6 @torch.no_grad()
799 7 def ctrl_gen(observe_logits, obtain_reward, device,
800 8             B, D, N, K, T,
801 9             mask_state: int = None,
802 10            invalid_ids: list = None,
803 11            inpainting_pos: list = None,
804 12            inpainting_values: list = None,
805 13            gamma=lambda r: np.cos(r * np.pi / 2),
806 14            prob_low_threshold=1e-10,
807 15            disable_tqdm=False):
808 16     """
809 17     B: batch size
810 18     D: sequence length
811 19     N: vocabulary size including the masked token
812 20     K: number of samples for Monte Carlo estimation

```

```

810 21
811 22 Sample from  $q(x) \propto r(x) p(x)$ ,  $p(x)$  comes from the masked model,
812 23 and  $r(x)$  is the reward function.
813 24
814 25 obtain_logits: [B,D] -> [B,D,N], return the logits  $P(X^d|X^{UM})$ 
815 26 obtain_reward: [*,D] -> [*], return the reward  $r(X)$ 
816 27 return: [B,D], the sampled sequence
817 28
818 29 By default, mask_state is N-1, and the valid_ids (i.e., real tokens)
819 30 are from 0 to N-2.
820 31 If there are other invalid tokens, one can specify the invalid_ids (
821 32 default is [mask_state]).
822 33 The logits for invalid tokens will be set to -inf.
823 34 when inpainting_pos is not None, we will always inpaint the values at
824 35 these positions with inpainting_values.
825 36 """
826 37
827 38 if mask_state is None:
828 39     mask_state = N-1
829 40
830 41 if invalid_ids is None:
831 42     invalid_ids = [mask_state]
832 43
833 44 elif mask_state not in invalid_ids:
834 45     invalid_ids = invalid_ids.append(mask_state)
835 46
836 47 if inpainting_pos is not None:
837 48     assert len(inpainting_values) == len(inpainting_pos)
838 49     assert all([0 <= pos < D for pos in inpainting_pos])
839 50     assert all([0 <= val < N and val !=
840 51                mask_state for val in inpainting_values])
841 52     inpainting_values = torch.tensor(
842 53         inpainting_values, dtype=torch.int).to(device)
843 54
844 55 def inpaint(x):
845 56     """*,D -> *,D inpaint"""
846 57     if inpainting_pos is None:
847 58         return x
848 59     else:
849 60         shape = x.shape
850 61         x = x.reshape(-1, D)
851 62         x[:, inpainting_pos] = inpainting_values
852 63         return x.reshape(shape)
853 64
854 65 D_essential = D - len(set(inpainting_pos))
855 66 # only the dimensions in dims_to_sample will be sampled. The rest will
856 67 be inpainted.
857 68
858 69 x = torch.full((B, D), mask_state, dtype=torch.int).to(device)
859 70 x = inpaint(x)
860 71 # B,D, initialize with mask_state and inpainted values
861 72
862 73 for t in tqdm(range(1, T+1), desc="Unmasking steps", disable=
863 74             disable_tqdm):
864 75     if int(D_essential*gamma(t/T)) == int(D_essential*gamma((t-1)/T)):
865 76         continue # no more tokens to unmask in this step
866 77
867 78     # predict all the masked tokens
868 79     logits = obtain_logits(x) # B,D,N,  $p(x^d | x^{UM})$ 
869 80     logits[:, :, invalid_ids] = -np.inf
870 81     masked = x == mask_state # B,D
871 82     masked_logits = logits[masked] # ?,N
872 83     samples = torch.distributions.Categorical(logits=masked_logits).
873 84         sample(
874 85             (K,)).to(dtype=torch.int, device=device) # K,?
875 86     x = x.repeat(K, 1).reshape(K, B, D).to(device) # K,B,D
876 87     x[:, masked] = samples

```

```
864 79     x = x.transpose(0, 1) # B,K,D
865 80     x = inpaint(x)
866 81     probs = obtain_reward(x) # B,K, p(y|x)
867 82     # avoid numerical instability during division
868 83     probs[probs < prob_low_threshold] = prob_low_threshold
869 84     weights = probs / probs.sum(dim=1, keepdim=True) # B,K, normalized
870 85     selected = torch.distributions.Categorical(probs=weights).sample()
871 86     # B
872 87     x = x[torch.arange(B), selected] # B,D
873 88
874 89     # remask the tokens based on their confidence scores
875 90     confidence = torch.ones_like(x, dtype=torch.float64).to(device)
876 91     confidence[masked] = torch.rand_like(
877 92         confidence[masked], dtype=torch.float64).to(device)
878 93     low_k_values, low_k_indices = torch.topk(
879 94         confidence, k=int(D_essential*gamma(t/T)), dim=-1, largest=False
880 95     )
881     x[torch.arange(B).unsqueeze(1), low_k_indices] = mask_state
882     return x
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
```