

---

# Structured Denoising Diffusion Models in Discrete State-Spaces

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Denoising diffusion probabilistic models (DDPMs) [14] have shown impressive  
2 results on image and waveform generation in continuous state spaces. Here, we  
3 introduce Discrete Denoising Diffusion Probabilistic Models (D3PMs), diffusion-  
4 like generative models for discrete data that generalize the multinomial diffusion  
5 model of Hooeboom et al. [15], by going beyond corruption processes with uni-  
6 form transition probabilities. This includes corruption with transition matrices that  
7 mimic Gaussian kernels in continuous space, matrices based on nearest neighbors  
8 in embedding space, and matrices that introduce absorbing states. The third al-  
9 lows us to draw a connection between diffusion models and autoregressive and  
10 mask-based generative models. We show that the choice of transition matrix is an  
11 important design decision that leads to improved results in image and text domains.  
12 We also introduce a new loss function that combines the variational lower bound  
13 with an auxiliary cross entropy loss. For text, this model class achieves strong  
14 results on character-level text generation while scaling to large vocabularies on  
15 LM1B. On the image dataset CIFAR-10, our models approach the sample quality  
16 and exceed the log-likelihood of the continuous-space DDPM model.

## 17 1 Introduction

18 Generative modeling is a core problem in machine learning, useful both for benchmarking our ability  
19 to capture statistics of natural datasets and for downstream applications that require generating  
20 high-dimensional data like images, text, and speech waveforms. There has been a great deal of  
21 progress with the development of methods like GANs [12, 3], VAEs [19, 28], large autoregressive  
22 neural network models [37, 36, 38], normalizing flows [27, 9, 18, 26], and others, each with their  
23 own tradeoffs in terms of sample quality, sampling speed, log-likelihoods, and training stability.

24 Recently, diffusion models [30] have emerged as a compelling alternative for image [14, 33] and au-  
25 dio [5, 20] generation, achieving comparable sample quality to GANs and log-likelihoods comparable  
26 to autoregressive models with fewer inference steps. A diffusion model is a parameterized Markov  
27 chain trained to reverse a predefined forward process, which is a stochastic process constructed to  
28 gradually corrupt training data into pure noise. Diffusion models are trained using a stable objective  
29 closely related to both maximum likelihood and score matching [16, 39], and they admit faster  
30 sampling than autoregressive models by using parallel iterative refinement [24, 32, 34, 31].

31 Although diffusion models have been proposed in both discrete and continuous state spaces [30],  
32 most recent work has focused on Gaussian diffusion processes that operate in continuous state spaces  
33 (e.g. for real-valued image and waveform data). Diffusion models with discrete state spaces have  
34 been explored for text and image segmentation domains [15], but they have not yet been demonstrated  
35 as a competitive model class for large scale text or image generation.

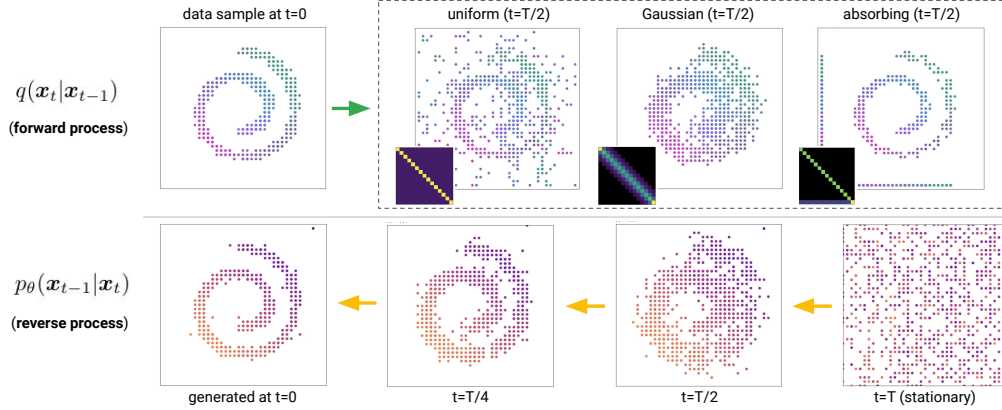


Figure 1: D3PM forward and (learned) reverse process applied to a quantized swiss roll. Each dot represents a 2D categorical variable. Top: samples from the uniform, discretized Gaussian, and absorbing state D3PM model forward processes, along with corresponding transition matrices  $Q$ . Bottom: samples from a learned discretized Gaussian reverse process.

36 Our aim in this work is to improve and extend discrete diffusion models by using a more structured  
 37 categorical corruption process to shape data generation, as illustrated in Figure 1. Our models do not  
 38 require relaxing or embedding discrete data (including images) into continuous spaces, and can embed  
 39 structure or domain knowledge into the transition matrices used by the forward process. We achieve  
 40 significantly improved results by taking advantage of this flexibility. We develop structured corruption  
 41 processes appropriate for text data, using similarity between tokens to enable gradual corruption  
 42 and denoising. Expanding further, we also explore corruption processes that insert [MASK] tokens,  
 43 which let us draw parallels to autoregressive and mask-based generative models. Finally, we study  
 44 discrete diffusion models for quantized images, taking inspiration from the locality exploited by  
 45 continuous diffusion models. This leads to a particular choice of discrete corruption process that  
 46 diffuses preferentially to more similar states and leads to much better results in the image domain.

47 Overall, we make a number of technical and conceptual contributions. Beyond designing several new  
 48 structured diffusion models, we introduce a new auxiliary loss which stabilizes training of D3PMs  
 49 and a family of noise schedules based on mutual information that lead to improved performance. We  
 50 strongly outperform various non-autoregressive baselines for text generation on character-level text  
 51 generation, and successfully scale discrete diffusion models to large vocabularies and long sequence  
 52 lengths. We also achieve strong results on the image dataset CIFAR-10, approaching or exceeding  
 53 the Gaussian diffusion model from Ho et al. [14] on log-likelihoods and sample quality.

## 54 2 Background: diffusion models

55 Diffusion models [30] are latent variable generative models characterized by a forward and a reverse  
 56 Markov process. The forward process  $q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1})$  corrupts the data  $\mathbf{x}_0 \sim$   
 57  $q(\mathbf{x}_0)$  into a sequence of increasingly noisy latent variables  $\mathbf{x}_{1:T} = \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T$ . The learned  
 58 reverse Markov process  $p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$  gradually denoises the latent variables  
 59 towards the data distribution. For example, for continuous data, the forward process typically adds  
 60 Gaussian noise, which the reverse process learns to remove.

61 In order to optimize the generative model  $p_\theta(\mathbf{x}_0)$  to fit the data distribution  $q(\mathbf{x}_0)$ , we typically  
 62 optimize a variational upper bound on the negative log-likelihood:

$$L_{vb} = \mathbb{E}_{q(\mathbf{x}_0)} \left[ \underbrace{D_{\text{KL}}[q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T)]}_{L_T} + \sum_{t=2}^T \underbrace{\mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)} [D_{\text{KL}}[q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)]]}_{L_{t-1}} \right. \\ \left. - \underbrace{\mathbb{E}_{q(\mathbf{x}_1|\mathbf{x}_0)} [\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)]}_{L_0} \right]. \quad (1)$$

63 When the number of time steps  $T$  goes to infinity, both the forward process and the reverse process  
 64 share the same functional form [10], allowing the use of a learned reverse process from the same  
 65 class of distributions as that of the forward process. Furthermore, for several choices of the forward  
 66 process the distribution  $q(\mathbf{x}_t|\mathbf{x}_0)$  converges to a stationary distribution  $\pi(\mathbf{x})$  in the limit  $t \rightarrow \infty$   
 67 independent of the value of  $\mathbf{x}_0$ . When the number of time steps  $T$  is large enough and we choose  
 68  $\pi(\mathbf{x})$  as the prior  $p(\mathbf{x}_T)$ , we can guarantee that the  $L_T$  term in (1) will approach zero regardless of  
 69 the data distribution  $q(\mathbf{x}_0)$ . (Alternatively, one can use a learned prior  $p_\theta(\mathbf{x}_T)$ .)

70 While  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$  can in theory be arbitrary, efficient training of  $p_\theta$  is possible when  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ :

- 71 1. Permits efficient sampling of  $\mathbf{x}_t$  from  $q(\mathbf{x}_t|\mathbf{x}_0)$  for an arbitrary time  $t$ , allowing us to  
 72 randomly sample timesteps and optimize each  $L_{t-1}$  term individually with stochastic  
 73 gradient descent,
- 74 2. Has a tractable expression for the forward process posterior  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ , which allows  
 75 us to compute the KL divergences present in the  $L_{t-1}$  term of (1).

76 The majority of recent work in continuous spaces [14, 31, 5, 24] defines the forward  
 77 and reverse distributions as  $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t|\sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$  and  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) =$   
 78  $\mathcal{N}(\mathbf{x}_{t-1}|\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t))$ , respectively. The aforementioned properties hold in the case of  
 79 these Gaussian diffusion models: the forward process  $q(\mathbf{x}_t|\mathbf{x}_0)$  converges to a stationary distribution,  
 80 motivating the choice  $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T|\mathbf{0}, \mathbf{I})$ , and both  $q(\mathbf{x}_t|\mathbf{x}_0)$  and  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  are tractable  
 81 Gaussian distributions for which the KL divergence can be computed analytically.

### 82 3 Diffusion models for discrete state spaces

83 Diffusion models with discrete state spaces were first introduced by Sohl-Dickstein et al. [30], who  
 84 considered a diffusion process over binary random variables. Hoogeboom et al. [15] extended  
 85 the model class to categorical random variables with transition matrices characterized by uniform  
 86 transition probabilities. In their supplementary material, Song et al. [31] also derived this extension,  
 87 although no experiments were performed with this model class. Here, we briefly describe a more  
 88 general framework for diffusion with categorical random variables which includes these models as  
 89 special cases.

90 For scalar discrete random variables with  $K$  categories  $x_t, x_{t-1} \in 1, \dots, K$  the forward transition  
 91 probabilities can be represented by matrices:  $[\mathbf{Q}_t]_{ij} = q(x_t = j|x_{t-1} = i)$ . Denoting the one-hot  
 92 version of  $x$  with the row vector  $\mathbf{x}$ , we can write

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \text{Cat}(\mathbf{x}_t; \mathbf{p} = \mathbf{x}_{t-1}\mathbf{Q}_t), \quad (2)$$

93 where  $\text{Cat}(\mathbf{x}; \mathbf{p})$  is a categorical distribution over the one-hot row vector  $\mathbf{x}$  with probabilities given  
 94 by the row vector  $\mathbf{p}$ , and  $\mathbf{x}_{t-1}\mathbf{Q}_t$  is to be understood as a row vector-matrix product. We assume  
 95 that  $\mathbf{Q}_t$  is applied to each pixel of an image or each token in a sequence independently, and that  
 96  $q$  factorizes over these higher dimensions as well; we thus write  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$  in terms of a single  
 97 element. Starting from  $\mathbf{x}_0$ , we obtain the following  $t$ -step marginal and posterior at time  $t-1$ :

$$q(\mathbf{x}_t|\mathbf{x}_0) = \text{Cat}(\mathbf{x}_t; \mathbf{p} = \mathbf{x}_0\bar{\mathbf{Q}}_t), \quad \text{with} \quad \bar{\mathbf{Q}}_t = \mathbf{Q}_1\mathbf{Q}_2 \dots \mathbf{Q}_t$$

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0)q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)} = \text{Cat}\left(\mathbf{x}_{t-1}; \mathbf{p} = \frac{\mathbf{x}_t\mathbf{Q}_t^\top \odot \mathbf{x}_0\bar{\mathbf{Q}}_{t-1}}{\mathbf{x}_0\bar{\mathbf{Q}}_t\mathbf{x}_t^\top}\right). \quad (3)$$

98 Note that due to the Markov property of the forward process  $q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0) = q(\mathbf{x}_t|\mathbf{x}_{t-1})$ . As-  
 99 suming that the reverse process  $p_\theta(\mathbf{x}_t|\mathbf{x}_{t-1})$  is also factorized as conditionally independent over  
 100 the image or sequence elements, the KL divergence between  $q$  and  $p_\theta$  can be computed by simply  
 101 summing over all possible values of each random variable; we thus satisfy criteria 1 and 2 discussed  
 102 in Section 2. Depending on  $\mathbf{Q}_t$ , the cumulative products  $\bar{\mathbf{Q}}_t$  can often be computed in closed form,  
 103 or simply precomputed for all  $t$ . However, for large  $K$  and large  $T$  this may be prohibitive. In  
 104 Appendix A.4 we discuss how to ensure  $\bar{\mathbf{Q}}_t$  can still be computed efficiently in this case, allowing  
 105 the framework to scale to a larger number of categories.

106 In the next section we discuss the choice of the Markov transition matrices  $\mathbf{Q}_t$  and corresponding  
 107 stationary distributions. From here on, we refer to the general class of diffusion models with discrete  
 108 state spaces as Discrete Denoising Diffusion Probabilistic Models (D3PMs).

### 109 3.1 Choice of Markov transition matrices for the forward process

110 An advantage of the D3PM framework described above is the ability to control the data corruption  
111 and denoising process by choosing  $Q_t$ , in notable contrast to continuous diffusion, for which only  
112 additive Gaussian noise has received significant attention. Besides the constraint that the rows of  $Q_t$   
113 must sum to one to conserve probability mass, the only other constraint in choosing  $Q_t$  is that the  
114 rows of  $\bar{Q}_t = Q_1 Q_2 \dots Q_t$  must converge to a known stationary distribution<sup>1</sup> when  $t$  becomes large,  
115 which can be guaranteed while imposing minimal restrictions on  $Q_t$  (see Appendix A.1).

116 We argue that for most real-world discrete data, including images and text, it makes sense to  
117 add domain-dependent structure to the transition matrices  $Q_t$  as a way of controlling the forward  
118 corruption process and the learnable reverse denoising process. Below we briefly discuss the uniform  
119 transition matrices that have been studied in prior work [15], along with a set of structured transition  
120 matrices we have explored for our image and text dataset experiments; see Appendix A.2 for more  
121 details on each matrix type. We also note that this set is not exhaustive, and many other transition  
122 matrices could also be used within the D3PM framework.

123 **Uniform (Appendix A.2.1).** Sohl-Dickstein et al. [30] considered a simple  $2 \times 2$  transition matrix for  
124 binary random variables. Hoogeboom et al. [15] later extended this to categorical variables, proposing  
125 a transition matrix  $Q_t = (1 - \beta_t)I + \beta_t/K \mathbb{1}\mathbb{1}^T$  with  $\beta_t \in [0, 1]$ . Since this transition matrix is  
126 doubly stochastic with strictly positive entries, the stationary distribution is uniform. Because the  
127 transition probability to any other state is uniform, in this paper we equivalently refer to this discrete  
128 diffusion instance as D3PM-uniform.

129 **Absorbing state (Appendix A.2.2).** Motivated by the success of BERT [8] and recent work on  
130 Conditional Masked Language Models (CMLMs) in text, we consider a transition matrix with an  
131 absorbing state (called [MASK]), such that each token either stays the same or transitions to [MASK]  
132 with some probability  $\beta_t$ . This does not impose particular relationships between categories, similar to  
133 uniform diffusion, but still allows corrupted tokens to be distinguished from original ones. Moreover,  
134 the stationary distribution is not uniform but has all the mass on the [MASK] token. For images, we  
135 reuse the grey pixel as the [MASK] absorbing token.

136 **Discretized Gaussian (Appendix A.2.3).** Instead of transitioning uniformly to any other state, for  
137 ordinal data we propose imitating a continuous space diffusion model by using a discretized, truncated  
138 Gaussian distribution. We choose a normalization such that the transition matrix is doubly stochastic,  
139 leading to a uniform stationary distribution. This transition matrix will transition between more  
140 similar states with higher probability, and is well suited for quantized ordinal data such as images.

141 **Token embedding distance (Appendix A.2.4).** Textual data does not have ordinal structure, but  
142 there may still be interesting semantic relationships. For instance, in a character level vocabulary  
143 vowels may be more similar to each other than they are to consonants. As a demonstration of the  
144 generality of the D3PM framework, we explore using similarity in an embedding space to guide the  
145 forward process, and construct a doubly-stochastic transition matrix that transitions more frequently  
146 between tokens that have similar embeddings while maintaining a uniform stationary distribution.

147 For uniform and absorbing-state diffusion, the cumulative products  $\bar{Q}_t$  can be computed in closed  
148 form (see Appendix A.4.1); the remainder can be precomputed.

### 149 3.2 Noise schedules

150 We consider several different options for the noise schedule of the forward process. For discretized  
151 Gaussian diffusion, we explore linearly increasing the variance of the Gaussian before discretizing  
152 it. (Note that a linear schedule for  $Q_t$  leads to a nonlinear amount of cumulative noise in  $\bar{Q}_t$ .) For  
153 uniform diffusion we use the cosine schedule which sets the cumulative probability of a transition to  
154 a cosine function, as introduced by Nichol and Dhariwal [24] and adapted by Hoogeboom et al. [15].  
155 For a general set of transition matrices  $Q_t$  (such as the one based on token embeddings), previously  
156 proposed schedules may not be directly applicable. We consider linearly interpolating the *mutual*  
157 *information* between  $x_t$  and  $x_0$  to zero, i.e.  $I(x_t; x_0) \approx (1 - \frac{t}{T})H(x_0)$ . Interestingly, for the  
158 specific case of absorbing-state D3PMs, this schedule reduces to exactly the  $(T - t + 1)^{-1}$  schedule

---

<sup>1</sup>If a stationary distribution is not known, we can instead append a rank-one matrix  $Q_{T+1}$  that ignores  $x_T$   
and produces a deterministic  $x_{T+1}$ ; we note that this is equivalent to introducing a learned prior  $p_\theta(x_T)$ .

159 proposed by Sohl-Dickstein et al. [30] for a Bernoulli diffusion process. See Appendix A.7 for more  
 160 details.

### 161 3.3 Parameterization of the reverse process

162 While it is possible to directly predict the logits of  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$  using a neural network  $\text{nn}_\theta(\mathbf{x}_t)$ ,  
 163 we follow Ho et al. [14] and Hoogeboom et al. [15] and focus on using a neural network  $\text{nn}_\theta(\mathbf{x}_t)$   
 164 to predict the logits of a distribution  $\tilde{p}_\theta(\tilde{\mathbf{x}}_0|\mathbf{x}_t)$ , which we combine with  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  and a  
 165 summation over one-hot representations of  $\mathbf{x}_0$  to obtain the following parameterization

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \sum_{\tilde{\mathbf{x}}_0} q(\mathbf{x}_{t-1}|\mathbf{x}_t, \tilde{\mathbf{x}}_0)\tilde{p}_\theta(\tilde{\mathbf{x}}_0|\mathbf{x}_t). \quad (4)$$

166 We note that under this  $\mathbf{x}_0$ -parameterization the KL divergence  $D_{\text{KL}}[q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)]$   
 167 will be zero if  $\tilde{p}_\theta(\tilde{\mathbf{x}}_0|\mathbf{x}_t)$  places all of its probability mass on the original value  $\mathbf{x}_0$ . The decomposition  
 168 of  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  in (3) also provides us with a motivation for this parameterization. According to  
 169 (3), in a given state  $\mathbf{x}_t$ , the optimal reverse process only takes into account transitions to states for  
 170 which  $q(\mathbf{x}_t|\mathbf{x}_{t-1})$  is non-zero. Therefore, the sparsity pattern of  $\mathbf{Q}_t$  determines the sparsity pattern  
 171 of the ideal reverse transition probabilities in  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ . The parameterization in (4) automatically  
 172 ensures that the learned reverse probability distribution  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$  has the correct sparsity pattern  
 173 dictated by the choice of the Markov transition matrix  $\mathbf{Q}_t$ . This parameterization also lets us perform  
 174 inference with  $k$  steps at a time, by predicting  $p_\theta(\mathbf{x}_{t-k}|\mathbf{x}_t) = \sum q(\mathbf{x}_{t-k}|\mathbf{x}_t, \tilde{\mathbf{x}}_0)\tilde{p}_\theta(\tilde{\mathbf{x}}_0|\mathbf{x}_t)$ .

175 Finally, when modeling ordinal discrete data, instead of predicting the logits of  $\tilde{p}_\theta(\tilde{\mathbf{x}}_0|\mathbf{x}_t)$  directly  
 176 with the output of a neural net, another option is to model the probabilities with a truncated discretized  
 177 logistic distribution (see Appendix A.8). This provides an extra ordinal inductive bias to the reverse  
 178 model and boosts FID and log-likelihood scores for images.

### 179 3.4 Loss function

180 While the original diffusion models introduced by Sohl-Dickstein et al. [30] were optimized with  
 181 the negative variational lower bound  $L_{\text{vb}}$  of (1), more recent diffusion models are optimized with  
 182 different objectives. For instance, Ho et al. [14] derive a simplified loss function ( $L_{\text{simple}}$ ) that  
 183 reweights the negative variational bound, and Nichol and Dhariwal [24] explore a hybrid loss  
 184  $L_{\text{hybrid}} = L_{\text{simple}} + \lambda L_{\text{vb}}$  (using one term to learn the predicted mean and the other to learn  
 185 predicted variance). Inspired by this recent work, we introduce an auxiliary denoising objective for  
 186 the  $\mathbf{x}_0$ -parameterization of the reverse process, which encourages good predictions of the data  $\mathbf{x}_0$  at  
 187 each time step. We combine this with the negative variational lower bound, yielding the following  
 188 alternative loss function:

$$L_\lambda = L_{\text{vb}} + \lambda \mathbb{E}_{q(\mathbf{x}_0)} \mathbb{E}_{q(\mathbf{x}_t|\mathbf{x}_0)} [-\log \tilde{p}_\theta(\mathbf{x}_0|\mathbf{x}_t)]. \quad (5)$$

189 Note that the auxiliary loss coincides with the cross entropy term  $L_0$  in (1) at  $t = 1$ . Fur-  
 190 thermore, due to the  $\mathbf{x}_0$ -parameterization of  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ , both the auxiliary loss term and  
 191  $D_{\text{KL}}[q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)]$  in  $L_{\text{vb}}$  are minimized exactly when  $\tilde{p}_\theta(\tilde{\mathbf{x}}_0|\mathbf{x}_t)$  has all its mass  
 192 on the datapoint  $\mathbf{x}_0$ . We find that training with this loss leads to improved quality of image samples.

## 193 4 Connection to existing probabilistic models for text

194 In this section we expand on interesting connections between the D3PM framework and several  
 195 existing probabilistic and language modeling approaches.

196 **BERT is a one-step diffusion model:** One possible D3PM transition matrix is a combination of a  
 197 uniform transition matrix and an absorbing state at the [MASK] token (i.e.  $\mathbf{Q} = \alpha \mathbb{1}e_m^T + \beta \mathbb{1}\mathbb{1}^T/K +$   
 198  $(1 - \alpha - \beta)I$ , where  $e_m$  is a one-hot vector on the [MASK] token). For a one-step diffusion process  
 199 in which  $q(\mathbf{x}_1|\mathbf{x}_0)$  replaces 10% of tokens with [MASK] and 5% uniformly at random, this leads  
 200 precisely to the BERT denoising objective, i.e.  $L_{\text{vb}} - L_T = -\mathbb{E}_{q(\mathbf{x}_1|\mathbf{x}_0)} [\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)] = L_{\text{BERT}}$ ,  
 201 since  $L_T$  is a constant independent of  $\theta$  (assuming a fixed prior).

202 **Autoregressive models are (discrete) diffusion models:** Consider a diffusion process that deter-  
 203 ministically masks tokens one-by-one in a sequence of length  $N = T$ :  $q([\mathbf{x}_t]_i | \mathbf{x}_0) = [\mathbf{x}_0]_i$  if  $i <$

204  $N - t$  else [MASK]. This is a deterministic forward process, so  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  is a delta distribution  
 205 on the  $\mathbf{x}_t$  sequence with one fewer mask:  $q([\mathbf{x}_{t-1}]_i|\mathbf{x}_t, \mathbf{x}_0) = \delta_{[\mathbf{x}_t]_i}$  if  $i \neq T - t$  else  $\delta_{[\mathbf{x}_0]_i}$ . While  
 206 this process is not applied independently to each token, it can be recast as an independently-applied  
 207 diffusion process on the product space  $[0 \dots N] \times \mathcal{V}$ , where each token is tagged with its position in  
 208 the sequence,  $\mathcal{V}$  is the vocabulary, and  $\mathbf{Q}$  is an  $N \times |\mathcal{V}| \times N \times |\mathcal{V}|$  sparse matrix.

209 Because all tokens except the  $i$ th token have deterministic posteriors, the KL divergence  
 210  $D_{KL}(q([\mathbf{x}_{t-1}]_j|\mathbf{x}_t, \mathbf{x}_0) || p_\theta([\mathbf{x}_{t-1}]_j|\mathbf{x}_t))$  is zero for all other positions. The only token for which  
 211 this is not true is the token at position  $i$ , for which  $D_{KL}(q([\mathbf{x}_{t-1}]_i|\mathbf{x}_t, \mathbf{x}_0) || p_\theta([\mathbf{x}_{t-1}]_i|\mathbf{x}_t)) =$   
 212  $-\log p_\theta([\mathbf{x}_0]_i|\mathbf{x}_t)$ , the standard cross entropy loss for an autoregressive model.

213 **(Generative) Masked Language-Models (MLMs) are diffusion models:** Generative Masked Lan-  
 214 guage Models ([11], [40]) are generative models that generate text from a sequence of [MASK]  
 215 tokens. They are usually trained by sampling a sequence  $\mathbf{x}_0$ , masking  $k$  tokens according to some  
 216 schedule, and learning to predict the masked tokens given context. It turns out that a D3PM absorbing  
 217 ([MASK]) model trained on the usual ELBO objective with the  $\mathbf{x}_0$ -parameterization from 3.3 reduces  
 218 to a reweighted version of this MLM objective (see Appendix A.3 for a detailed derivation).

## 219 5 Text generation

220 For text, we experiment with generation on two datasets: text8 [22], a character-level dataset extracted  
 221 from English-language Wikipedia, and the One Billion Word dataset (LM1B) [4], a large dataset of  
 222 shuffled English-language sentences. For both, we train a D3PM uniform model based on the work  
 223 by Hooeboom et al. [15] (D3PM uniform) and a model that masks tokens (D3PM absorbing). For  
 224 text8, we also consider a model that transitions uniformly to nearest neighbors in a token embedding  
 225 space (D3PM NN). We follow Hooeboom et al. [15] and use  $T = 1000$  timesteps, although we are  
 226 also able to evaluate on fewer due to the parameterization in Section 3.3.

### 227 5.1 Character-level generation on text8

228 text8 is a character-level text dataset consisting of a small vocabulary of 27 tokens: the letters ‘a’-‘z’  
 229 and the ‘\_’ whitespace token. We follow the convention of training and evaluating text8 in chunks  
 230 of length 256 without any preprocessing [15]. For nearest-neighbor D3PM, our nearest neighbor  
 231 graph in character-space is shown in Appendix B.2. D3PM uniform models were trained with a  
 232 cosine schedule from Hooeboom et al. [15] (ablations in Appendix B.2), while D3PM absorbing  
 233 and D3PM NN models were trained with a mutual information schedule.

Table 1: Quantitative results on text8. NLL is reported on the entire test set. Sample times are for generating a single example of length 256. Results are reported on two seeds. All models are standard 12-layer transformers unless otherwise noted. <sup>†</sup>Transformer XL is a 24-layer transformer, using a 784 context window. \*Results reported by [15] by running code from official repository.

Model	Model steps	NLL (bits/char) (↓)	Sample time (s) (↓)
Discrete Flow [35] ( $8 \times 3$ layers)	-	1.23	0.16
Argmax Coupling Flow [15]	-	1.80	$0.40 \pm 0.03$
IAF / SCF [41]*	-	1.88	$0.04 \pm 0.0004$
Multinomial Diffusion (D3PM uniform) [15]	1000	$\leq 1.72$	$26.6 \pm 2.2$
D3PM uniform [15] (ours)	1000	$\leq 1.61 \pm 0.02$	$3.6 \pm 0.4$
D3PM NN ( $L_{\lambda=0}$ ) (ours)	1000	$\leq 1.59 \pm 0.03$	$3.147 \pm 0.0002$
D3PM mask ( $L_{\lambda=0, .01}$ ) (ours)	1000	$\leq 1.45 \pm 0.02$	$3.4 \pm 0.3$
D3PM uniform [15] (ours)	256	$\leq 1.68 \pm 0.01$	$0.58 \pm 0.0001$
D3PM NN ( $L_{\lambda=0}$ ) (ours)	256	$\leq 1.64 \pm 0.02$	$0.81 \pm 0.002$
D3PM absorbing ( $L_{\lambda=0, .01}$ ) (ours)	256	$\leq 1.47 \pm 0.03$	$0.59 \pm 0.002$
Transformer decoder (ours)	256	1.23	$0.357 \pm 0.0002$
Transformer decoder [1]	256	1.18	-
Transformer XL [7] <sup>†</sup>	256	1.08	-
D3PM uniform [15] (ours)	20	$\leq 1.791 \pm 0.03$	$0.078 \pm 0.0001$
D3PM NN ( $L_{\lambda=0}$ ) (ours)	20	$\leq 1.75 \pm 0.02$	$0.111 \pm 0.0001$
D3PM absorbing ( $L_{\lambda=0, .01}$ ) (ours)	20	$\leq 1.56 \pm 0.04$	$0.078 \pm 0.0003$

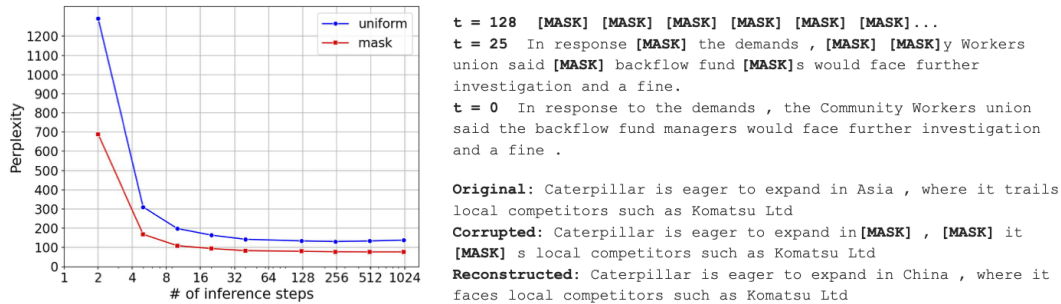


Figure 2: Left: perplexity v.s. sampling iterations for LM1B. Right: Using a trained D3PM absorbing model for LM1B to (top) generate new sentences and (bottom) reconstruct corrupted examples.

Table 2: Quantitative results on LM1B. Perplexity reported on the test set normalized by the number of words. Results are reported on two seeds. Context window has length 128 for all diffusion models, and all have 12 layers unless otherwise noted. <sup>†</sup> Transformer XL is a 24 layer transformer.

Metric:	Perplexity			Sample time		
	inference steps:	1000	128	64	1000	128
D3PM-uniform	137.9 ± 2.1	139.2 ± 1.2	145 ± 1.2	1.8 ± 0.02	0.21 ± 0.03	0.08 ± 0.01
D3PM-absorbing	76.9 ± 2.3	80.13 ± 1.2	83.62 ± 6.1	1.9 ± 0.01	0.19 ± 0.002	0.099 ± 0.01
Transformer (ours)	-	43.62549	-	-	0.256 ± 0.03	-
Transformer XL [7] <sup>†</sup>	-	21.8	-	-	-	-

234 Table 1 shows that for D3PM, the D3PM absorbing model performed the best, exceeding the  
235 uniform and NN diffusion models. We were able to improve upon the baseline result of [15] with  
236 hyperparameter tuning, and our uniform and NN results outperformed results from Hoogeboom  
237 et al. [15] across all inference steps, down to as few as 20. We found that  $L_{\lambda=0.01}$  worked best  
238 for D3PM absorbing, while  $L_{v_b}$  was better for D3PM uniform. Our model outperforms all non-  
239 autoregressive baselines except one, the Discrete Flow model [35] (for which unfortunately no  
240 open-source implementations exist), and is also faster than all but one method, the IAF/SCF model  
241 [41]. It is also nearly 20x faster than an autoregressive transformer of the same size. For further  
242 results, see Appendix B.2. D3PM with the mask absorbing token was by far the best performing  
243 model, which lends credibility to the use of masks in denoising auto-encoders. Nearest-neighbor  
244 diffusion only narrowly improves upon a D3PM-uniform model: this was a surprising negative result  
245 for us, suggesting that not all notions of structure are meaningful.

## 246 5.2 Text generation on LM1B

247 Text generation for large-scale text datasets and large vocabularies with discrete diffusion models has  
248 not been previously demonstrated. We include results from LM1B as a proof of concept, showing  
249 that these models can indeed scale (as discussed in Appendix A.4), and that the D3PM absorbing  
250 model continues to excel. We do not report D3PM-NN results on LM1B due to slow training and a  
251 lack of signal from the text8 results. All models were trained and evaluated on packed sequences of  
252 length 128, using a sentencepiece<sup>2</sup> vocabulary of size 8192.

253 Table 2 contains results from experiments on LM1B. We also include a plot of inference time as a  
254 function of iterations in Appendix B.3. Overall, mask diffusion (D3PM absorbing) does relatively  
255 well, approaching the performance of a comparable autoregressive model of the same size, and scaling  
256 to far fewer steps, while uniform diffusion performs significantly worse. We found the the  $L_{\lambda=0.01}$   
257 loss worked best for the mask absorbing model, but reduced performance for the other models. We  
258 note the surprising scaling in perplexity in Figure 2, achieving remarkably strong results with as few  
259 as 10 inference steps. We also show samples from our model, as well as completions from corrupted  
260 samples.

<sup>2</sup><https://github.com/google/sentencepiece>



Table 3: Inception scores (IS), Frechet Inception Distance (FID) and negative log-likelihood (NLL) on the image dataset CIFAR-10. The NLL is reported on the test set in bits per dimension. Where available, results are reported over models trained with two seeds.

Model	IS ( $\uparrow$ )	FID ( $\downarrow$ )	NLL ( $\downarrow$ )
Sparse Transformer [6]			2.80
NCSN [32]	$8.87 \pm 0.12$	25.32	
NCSNv2 [33]	$8.40 \pm 0.07$	10.87	
StyleGAN2 + ADA [17]	$9.74 \pm 0.05$	3.26	
Diffusion (original), $L_{vb}$ [30]			$\leq 5.40$
DDPM $L_{vb}$ [14]	$7.67 \pm 0.13$	13.51	$\leq 3.70$
DDPM $L_{simple}$ [14]	$9.46 \pm 0.11$	3.17	$\leq 3.75$
Improved DDPM $L_{vb}$ [24]		11.47	$\leq 2.94$
Improved DDPM $L_{simple}$ [24]		2.90	$\leq 3.37$
DDPM++ cont [34]		2.92	2.99
NCSN++ cont. [34]	9.89	2.20	
D3PM uniform $L_{vb}$	$6.09 \pm 0.09$	$48.065 \pm 0.83$	$\leq 5.11 \pm 0.01$
D3PM absorbing $L_{vb}$	$6.20 \pm 0.08$	$41.75 \pm 1.32$	$\leq 4.88 \pm 0.02$
D3PM absorbing $L_{\lambda=0.001}$	$6.70 \pm 0.11$	30.99	$\leq 4.45$
D3PM Gauss $L_{vb}$	$7.92 \pm 0.10$	15.72	$\leq 3.97$
D3PM Gauss $L_{\lambda=0.001}$	$8.54 \pm 0.16$	$8.23 \pm 0.18$	$\leq 3.985 \pm 0.007$
D3PM Gauss + logistic $L_{\lambda=0.001}$	$8.65 \pm 0.19$	$6.90 \pm 0.02$	$\leq 3.442 \pm 0.003$

## 261 6 Image generation

262 We evaluate the performance of several D3PM models on the task of unconditional image generation  
 263 with the dataset CIFAR-10 [21]. We follow Ho et al. [14] and use  $T = 1000$  timesteps for all models  
 264 and verify that for all models the forward process converges to the stationary distribution within  $T$   
 265 steps, yielding a value of at most  $L_T \approx 10^{-5}$  bits per dimension. We train three versions of D3PM  
 266 with different transition matrices: doubly stochastic matrices with uniform transition probabilities  
 267 (D3PM uniform) [15], transition matrices with an absorbing state located at R, G and B values of 128  
 268 (D3PM absorbing) and doubly stochastic discretized Gaussian transition matrices (D3PM Gauss). For  
 269 the D3PM uniform model we experimented with a linear  $\beta_t$  schedule as well as the cosine schedule  
 270 as proposed in [15], with the cosine schedule producing the best results. For D3PM absorbing we  
 271 use the schedule  $\beta_t = (T - t + 1)^{-1}$  as also proposed in [30], which corresponds to increasing the  
 272 probability of being in the absorbing state linearly over time. For D3PM Gauss we use the same  
 273 linear schedule as in [14]. See Appendix B.1 for more details on the experimental setup.

274 Table 3 shows that for D3PM models trained with the  $L_{vb}$  objective, D3PM Gauss performs better  
 275 than D3PM absorbing and uniform on all metrics: Inception score (IS), Frechet Inception Distance  
 276 (FID) and negative log-likelihood (NLL). The IS score of the uniform and absorbing D3PM models

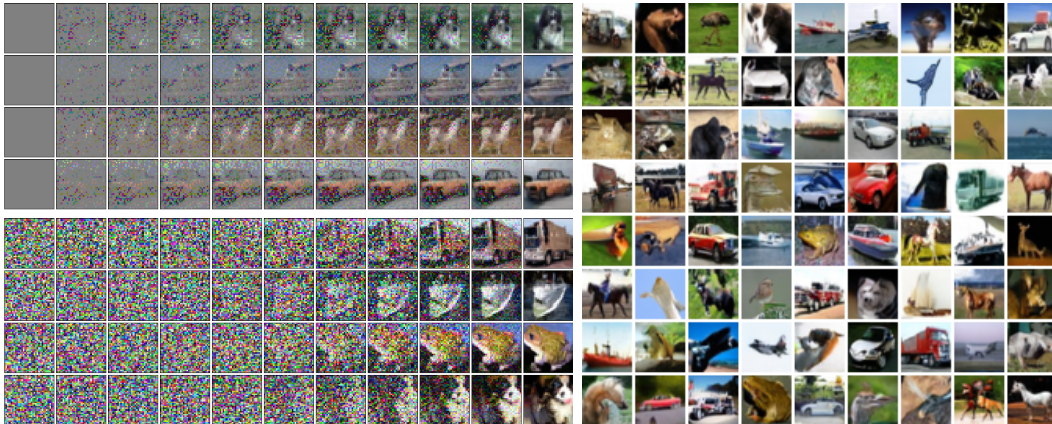


Figure 3: Left: progressive sampling at  $t = 1000, 900, 800, \dots, 0$  for D3PM absorbing (top) and D3PM Gauss + logistic (bottom), trained with  $L_{\lambda}$  loss on CIFAR-10. These samples were cherry picked. Right: (non cherry picked) samples from the D3PM Gauss + logistic model.



277 are comparable, while the FID score and NLL of the D3PM absorbing model are slightly better. We  
278 trained both D3PM absorbing and D3PM Gauss with the alternative loss function  $L_\lambda$  of (5), and  
279 we found  $\lambda = 0.001$  to work best. We have also experimented with larger values of  $\lambda$  and a model  
280 trained only with the auxiliary denoising term in (5). Although this led to a more rapid increase  
281 in performance early on in training, the NLL leveled off at higher values for larger  $\lambda$  and the FID  
282 even started increasing again. The results show that the models trained with  $L_\lambda$  perform significantly  
283 better than their counterparts trained with  $L_{vb}$ . One explanation for this boost in performance is that  
284 the cross entropy term leads to gradient noise that varies less with the time step  $t$ , which is in contrast  
285 to the large change in magnitude of the  $L_{t-1}$  terms in  $L_{vb}$  for smaller  $t$ , as demonstrated by Nichol  
286 and Dhariwal [24]. Finally, we achieve our best results by combining D3PM Gauss trained on  $L_\lambda$   
287 with a truncated logistic parameterization of the reverse process distribution  $p_\theta(\tilde{x}_0|\mathbf{x}_t)$  (D3PM Gauss  
288 + logistic). Figure 3 shows samples from our best model (D3PM Gauss + logistic), as well as the  
289 D3PM absorbing model.

## 290 7 Related Work

291 Diffusion generative models were first proposed by Sohl-Dickstein et al. [30] and have gained  
292 renewed attention recently due to strong results on image and waveform generation [14, 5]. Recent  
293 works have proposed improvements for diffusion model training, including importance sampling of  
294 the ELBO, better noise schedules [24] and implicit diffusion models [31]. Several works have also  
295 drawn connections to score matching [39, 16, 32], leading to improved sampling algorithms in the  
296 continuous-time limit [34].

297 While most works have considered continuous diffusion models, discrete diffusion-like models were  
298 described in [30] and applied to text generation and image segmentation data in [15]. Some works  
299 [25, 23] have dealt with discrete data by embedding it in continuous space and leveraging Gaussian  
300 diffusion, but have not applied this to text. Seff et al. [29] also considered generation of discrete  
301 structured objects using a diffusion-like Markov corruption process.

302 For text, denoising autoencoders have a long history both in representation learning [2, 8] and more  
303 recently as generative models [40]. These closely resemble our MASK diffusion variants for a  
304 particular schedule and transition matrix (see Section 4), although our framing allows us to compute  
305 log-likelihoods and experiment with alternative transition matrices. Other works have considered  
306 non-autoregressive translation via insertion and deletion [13] or masking [11].

## 307 8 Discussion

308 We have presented D3PMs, a class of models that improves diffusion models for discrete data by  
309 defining new kinds of discrete corruption processes. We achieve strong empirical results relative to  
310 previous work on discrete diffusion models, even surpassing performance of continuous diffusion  
311 models in terms of log-likelihoods for image generation. While these results are promising, one  
312 limitation is that—like much other work on non-autoregressive generative models—our models are  
313 still inferior to strong autoregressive models like Transformer XL for text generation, and continuous  
314 diffusion models still yield stronger results on image quality. We expect that D3PMs can benefit  
315 further from the rapid development of continuous diffusion models [34, 24]. For example, further  
316 research in alternative losses for D3PM’s can take inspiration from the reweighted  $L_{\text{simple}}$  objective  
317 used in [14], or the resampled variational bound in Nichol and Dhariwal [24]. Furthermore, D3PM’s  
318 might benefit from increasing the number of timesteps and a more optimized noise schedule, as  
319 discussed in Nichol and Dhariwal [24]. Another limitation comes from the choice of evaluation  
320 metrics that we use (and that are standard for evaluation of generative models). Inception score  
321 and Frechet Inception Distance are based on neural networks that have been trained on a particular  
322 distribution of data, which is not representative for all use-cases, and focusing on average quality  
323 metrics may not accurately reflect performance across the wide diversity of settings where these  
324 generative models may be applied. This creates a risk of negative social impacts where advances  
325 disproportionately favor a subset of the population. Going forward, we are excited about the space  
326 of possibilities that arise within the D3PM framework. We have found successes in leveraging the  
327 flexibility that comes from defining discrete corruption processes for discrete data, but we believe  
328 that there are many more possibilities that make use of richer forms of structure to define even more  
329 powerful discrete diffusion models.

## 330 References

- 331 [1] Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. Character-Level  
332 language modeling with deeper Self-Attention. August 2018.
- 333 [2] Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent. Generalized denoising Auto-  
334 Encoders as generative models. May 2013.
- 335 [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity  
336 natural image synthesis. In *International Conference on Learning Representations*, 2019.
- 337 [4] Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and  
338 Tony Robinson. One billion word benchmark for measuring progress in statistical language  
339 modeling. December 2013.
- 340 [5] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan.  
341 WaveGrad: Estimating gradients for waveform generation. September 2020.
- 342 [6] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with  
343 sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- 344 [7] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov.  
345 Transformer-XL: Attentive language models beyond a Fixed-Length context. January 2019.
- 346 [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of  
347 deep bidirectional transformers for language understanding. October 2018.
- 348 [9] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP.  
349 *arXiv preprint arXiv:1605.08803*, 2016.
- 350 [10] W Feller. On the theory of stochastic processes, with particular reference to applications. In  
351 *Proceedings of the [First] Berkeley Symposium on Mathematical Statistics and Probability*. The  
352 Regents of the University of California, 1949.
- 353 [11] Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. Mask-Predict: Parallel  
354 decoding of conditional masked language models. April 2019.
- 355 [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil  
356 Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural  
357 Information Processing Systems*, pages 2672–2680, 2014.
- 358 [13] Jiatao Gu, Changhan Wang, and Jake Zhao. Levenshtein transformer. May 2019.
- 359 [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In  
360 *Advances in Neural Information Processing Systems*, pages 6840–6851, 2020.
- 361 [15] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax  
362 flows and multinomial diffusion: Towards non-autoregressive language models. *arXiv preprint  
363 arXiv:2102.05379*, 2021.
- 364 [16] Aapo Hyvärinen, Juha Karhunen, and Erkki Oja. *Independent component analysis*, volume 46.  
365 John Wiley & Sons, 2004.
- 366 [17] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila.  
367 Training generative adversarial networks with limited data. *arXiv preprint arXiv:2006.06676v1*,  
368 2020.
- 369 [18] Diederik P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolu-  
370 tions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- 371 [19] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint  
372 arXiv:1312.6114*, 2013.
- 373 [20] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile  
374 diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.

- 375 [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.  
376 2009.
- 377 [22] Matt Mahoney. Text8 dataset. <http://mattmahoney.net/dc/textdata>, 2011. Accessed:  
378 2021-5-24.
- 379 [23] Gautam Mittal, Jesse Engel, Curtis Hawthorne, and Ian Simon. Symbolic music generation  
380 with diffusion models. March 2021.
- 381 [24] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. 2021.
- 382 [25] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon.  
383 Permutation invariant graph generation via score-based generative modeling. March 2020.
- 384 [26] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji  
385 Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint*  
386 *arXiv:1912.02762*, 2019.
- 387 [27] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In  
388 *International Conference on Machine Learning*, pages 1530–1538, 2015.
- 389 [28] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation  
390 and approximate inference in deep generative models. In *International Conference on Machine*  
391 *Learning*, pages 1278–1286, 2014.
- 392 [29] Ari Seff, Wenda Zhou, Farhan Damani, Abigail Doyle, and Ryan P Adams. Discrete object  
393 generation with reversible inductive construction. July 2019.
- 394 [30] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsuper-  
395 vised learning using nonequilibrium thermodynamics. In *International Conference on Machine*  
396 *Learning*, pages 2256–2265, 2015.
- 397 [31] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In  
398 *International Conference on Learning Representations*, 2021.
- 399 [32] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data  
400 distribution. In *Advances in Neural Information Processing Systems*, pages 11895–11907, 2019.
- 401 [33] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models.  
402 *arXiv preprint arXiv:2006.09011*, 2020.
- 403 [34] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon,  
404 and Ben Poole. Score-based generative modeling through stochastic differential equations.  
405 November 2020.
- 406 [35] Dustin Tran, Keyon Vafa, Kumar Agrawal, Laurent Dinh, and Ben Poole. Discrete flows:  
407 Invertible generative models of discrete data. In *Advances in Neural Information Processing*  
408 *Systems*, volume 32, 2019.
- 409 [36] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex  
410 Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. WaveNet: A generative  
411 model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- 412 [37] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural  
413 networks. *International Conference on Machine Learning*, 2016.
- 414 [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
415 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Informa-*  
416 *tion Processing Systems*, pages 5998–6008, 2017.
- 417 [39] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural*  
418 *Computation*, 23(7):1661–1674, 2011.
- 419 [40] Alex Wang and Kyunghyun Cho. BERT has a mouth, and it must speak: BERT as a markov  
420 random field language model. February 2019.
- 421 [41] Zachary M Ziegler and Alexander M Rush. Latent normalizing flows for discrete sequences.  
422 January 2019.

423 **Checklist**

- 424 1. For all authors...
- 425 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s  
426 contributions and scope? [Yes]
- 427 (b) Did you describe the limitations of your work? [Yes] See Section 8
- 428 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See  
429 Section 8
- 430 (d) Have you read the ethics review guidelines and ensured that your paper conforms to  
431 them? [Yes]
- 432 2. If you are including theoretical results...
- 433 (a) Did you state the full set of assumptions of all theoretical results? [Yes]
- 434 (b) Did you include complete proofs of all theoretical results? [Yes]
- 435 3. If you ran experiments...
- 436 (a) Did you include the code, data, and instructions needed to reproduce the main experi-  
437 mental results (either in the supplemental material or as a URL)? [No] The code cannot  
438 be made available at this time, but we will work to get it open sourced for a camera  
439 ready version.
- 440 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they  
441 were chosen)? [Yes] The training details can be found in Appendix B.
- 442 (c) Did you report error bars (e.g., with respect to the random seed after running experi-  
443 ments multiple times)? [Yes] We have included averages and standard deviations for as  
444 many models as possible for two seeds.
- 445 (d) Did you include the total amount of compute and the type of resources used (e.g., type  
446 of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix B.
- 447 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 448 (a) If your work uses existing assets, did you cite the creators? [Yes] We have included  
449 citations to all datasets.
- 450 (b) Did you mention the license of the assets? [N/A] These are standard datasets.
- 451 (c) Did you include any new assets either in the supplemental material or as a URL? [No]
- 452 (d) Did you discuss whether and how consent was obtained from people whose data you’re  
453 using/curating? [N/A]
- 454 (e) Did you discuss whether the data you are using/curating contains personally identifiable  
455 information or offensive content? [No] While it is valid to question if such information  
456 appears in large datasets like CIFAR-10 and LM1B, these are standard benchmarks,  
457 and it is beyond the scope of this work to do a thorough analysis.
- 458 5. If you used crowdsourcing or conducted research with human subjects...
- 459 (a) Did you include the full text of instructions given to participants and screenshots, if  
460 applicable? [N/A]
- 461 (b) Did you describe any potential participant risks, with links to Institutional Review  
462 Board (IRB) approvals, if applicable? [N/A]
- 463 (c) Did you include the estimated hourly wage paid to participants and the total amount  
464 spent on participant compensation? [N/A]