SUNMASK: Mask Enhanced Control in Step Unrolled Denoising Autoencoders

Anonymous Author(s)
Affiliation
Address
email

Abstract

This paper introduces SUNMASK, an approach for generative sequence modeling based on masked unrolled denoising autoencoders. By explicitly incorporating a conditional masking variable, as well as using this mask information to modulate losses during training based on expected exemplar difficulty, SUNMASK models discrete sequences without direct ordering assumptions. The addition of masking terms allows for fine-grained control during generation, starting from random tokens and a mask over subset variables, then predicting tokens which are again combined with a subset mask for subsequent repetitions. This iterative process gradually improves token sequences toward a structured output, while guided by proposal masks. The broad framework for unrolled denoising autoencoders is largely independent of model type, and we utilize both transformer and convolution based architectures in this work. We demonstrate the efficacy of this approach both qualitatively and quantitatively, applying SUNMASK to generative modeling of symbolic polyphonic music, and language modeling for English at both the word level and the character level.

1 Introduction

Generative modeling is a widespread methodology, used across various domains for automatic content generation, dataset understanding, compression, feature learning, and expert aid for domain analysis. Due to improvements in quality, diversity, and control for generative models, these tools have begun to see commercial application and adoption in an ever growing set of real-world settings. Generative models as a whole can be further stratified into different factorizations of the modeling approach, with one of the most prominent categories being autoregressive models (AR).

1.1 Autoregressive Models

AR modeling with deep neural networks has been a dominant approach to generative modeling and feature learning [35, 66, 69, 36, 72, 70] which has many crucial advantages in both training and inference. One key concern is the necessity of defining a "dependency chain" in the form of a (typically) directed acyclic graph (DAG). By defining this chain of variables, sampling during inference can be accomplished in a straightforward manner using ancestral sampling - sampling from the first variable or variables in the DAG, using those to conditionally estimate a probability distribution for subsequent variables. Repeatedly applying this process to walk the chain will eventually result in a sampled estimate of the data, under the assumptions of the AR factorization used.

Many applications have straightforward orderings in which to define this chain of variables, based on domain knowledge. For example following the flow of time for timeseries modeling is often

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

a logical choice, allowing models to make predictions into the future from the past. However in many other domains, for example images, language, or music, the process of defining a dependency chain over input variables (e.g. pixels, characters, words, or notes) is far from straightforward, as for any arbitrary ordering there can frequently be examples where this ordering *creates* long-term dependencies, or otherwise makes satisfaction of dependencies during training and evaluation more difficult than another alternative ordering.

This divide becomes further compounded in many creative applications to these domains, as creators typically iterate repeatedly: forming a concept, applying an initial sequence of steps to create the framing of the concept, and seeing where the creative flow may lead to alterations in the original concept, thus altering future steps. Though the resulting output may be perceived in a time-ordered fashion (for example, reading a book or listening to a song), the initial creation was performed globally and holistically. This global view is often critical to creating elements such as foreshadowing and tension which make the resulting output interesting or enjoyable. This iterative process is directly at odds with a strict AR factorization, and requires well trained AR models to cope with a high degree of uncertainty and multi-modality for long range dependencies, which can lead to logical mistakes or other errors.

1.2 Non-Autoregressive Models

An alternative methodology for generative modeling is non-autoregression (NAR), broadly covering a large number of different modeling approaches which attempt to remove assumptions about variable ordering, instead either hand-defining per-exemplar orderings, or modeling variables jointly without resorting to chain rule factorization. One way to define an ordering over variables is via masking of inputs or intermediate network representations [20, 67, 68, 73, 69, 53], and indeed modern AR approaches such as transformers [71] use an autoregressive mask internally to define the chain of variables order. These masks can either be constant over all training (as in standard AR transformers and PixelCNN [69]) or dynamic per example (as in MADE [20]). When masks are dynamic per example, we begin to see the relationship between enforcing AR via masking and NAR methods, as although some ordering is assumed this ordering is no longer constant, and it becomes possible to use the same trained model to evaluate the probability of a particular output variable under *multiple* possible orderings.

Closely linked to masking methods are so called *diffusion models*, which relax the variable ordering problem through noise prediction [63, 65, 27]. Rather than predicting a new variable or variables given previous ones in an arbitrarily chosen DAG, diffusion models focus on predicting a less noisy version of many variables jointly, given a set of noisy input variables. Iteratively applying this learned denoising improvement operator should eventually result in predicting a fully clean output estimate, given either a noisy version of the target domain, or even starting from pure noise. Given this framing it is clear that diffusion models are closely linked to denoising methods in general, specifically denoising autoencoders, as well as modern density modeling approaches such as generative adversarial networks (GAN [21]), variational autoencoders (VAE [38]), flow-based models (NICE [13], RealNVP [14], Normalizing Flows [57], IAF [39], MAF [53]), iterative canvas sampling (DRAW [23]), and noise contrastive estimation (NCE [25]). Particular applications of this denoising philosophy such as BERT [12], WaveGrad [8], and GLIDE [52], have resulted in large quality improvements for feature learning and data generation for text, images, and audio [43, 62, 56, 26].

1.3 Trade-offs

The choice between AR and NAR methods is not clear-cut. For many domains, high-quality models exist using both approaches but we can define some crucial parameters. Some NAR methods such as GAN or VAE are capable of generating output in only one inference step, however they are typically hard to train on certain data modalities (e.g. text data) comparing to the AR counterparts. Other NAR methods such as diffusion models typically allow for choosing a diffusion length during inference, which is independent of that used at training. Choosing a low diffusion length can frequently lead to poor sample quality, and tuning this setting (among many others) is critical to high quality generation. However if the tuned diffusion length for a given sequence (of length T) is shorter than the length of those sequences, the NAR method has a computational advantage over the equivalent AR model (which would require T steps for a T length sequence). In addition, the ability to tune this diffusion

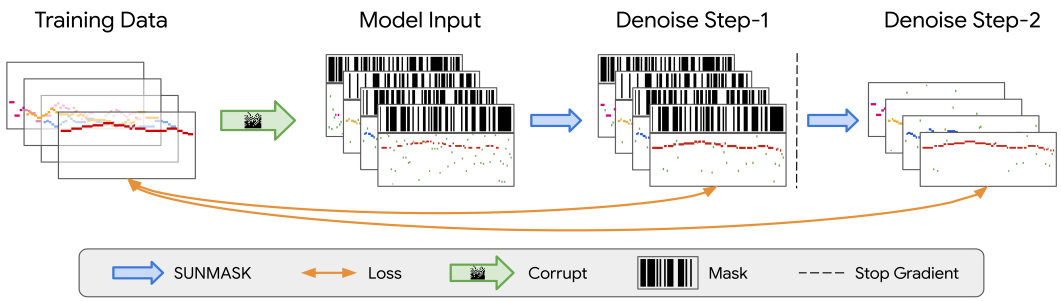

Figure 1: Step-unrolled denoising training procedure for SUNMASK on polyphonic music with step length 2. Training data (left) consists of four voices which are first corrupted by sampling a random mask for each voice and replacing the masked training data (red) with random pitches (green). SUNMASK then takes both mask and corrupted training data as input, predicting the original data as output. In the second step, the model takes a sampled version of the model predictions and the same mask as input, outputting another prediction of the original data.

length can be useful in interactive applications, or when a variety of output is desirable. This setting can also be a curse, as even well-trained models perform poorly with improper diffusion settings. Several branches of current research are focused on improving guarantees [32, 64, 37] and convergence speed for diffusion models [40, 41, 33, 74].

2 Method

The relationship between discrete diffusion and denoising autoencoders has been explored in previous work [28, 61, 3, 29]. We build upon this foundation, combined with many insights from prior orderless modeling approaches, crucially Orderless NADE [68], Coconet [30], and SUNDAE [61]. We approach generative modeling in the following manner: given a target data representation, learn a model which can map from a noisy version of input data to a corrected form of the input. In this work, we use multinomial noise - namely entries are corrupted to 1 of P possible values (for a given set size P), with the number of noised entries defining the relative noise level for the training example. This is similar to many diffusion approaches at a high level, and shown to be an effective tool in SUNDAE [61]. In addition to the use of multinomial noise, we also form a mask representing *where* the data was noised, feeding this mask alongside the input data to form a conditional probability distribution.

Like SUNDAE, our method is built around a process $x_t \sim f_\theta(\cdot|x_{t-1}; m)$ on a space $X = \{1, \ldots, v\}^N$ of arrays of categorical variables. However in our case the parametric transition function f_θ takes an additional argument $m \in 0, 1^N$. During training, m indicates variables that were not corrupted, and as a consequence we can use it during inference to tell f_θ which variables to trust. As in SUNDAE, the transition f_θ models the variables as conditionally independent of one another.

Given a sequence of masks $m_0, \ldots, m_{T-1}$, the generating distribution of our model derives from a prior p_0 (typically uniform noise) and repeated application of f_θ:

$$p_T(x_T; m_0, \ldots, m_{T-1}) = \left(\sum_{x_1, \ldots, x_{T-1} \in X} \prod_{t=1}^{T} f_\theta(x_t | x_{t-1}; m_{t-1}) \right) p_0(x_0)$$

In practice, p_0 is typically elementwise iid uniform noise, and the masks $m_0, \ldots, m_{T-1}$ are drawn according to a schedule and may be held constant for several steps.

To train f_θ, we take a training example $x \sim p_{\text{data}}$ and draw a mask m. We apply the corruption procedure $x_0 \sim q(\cdot|x; m)$ to obtain x_0 which equals x where the mask m is true and uniform random values elsewhere. Then we iterate $x_t \sim f_\theta(\cdot|x_{t-1}; m)$ with the aim of reconstructing x.

Similar to SUNDAE, our objective is to minimize $\frac{1}{2}(L^{(1)} + L^{(2)})$ where

$$L^{(t)}(\theta) = -\mathbb{E}_{m_0,\cdots,m_{t-1}}\mathbb{E}_{\substack{x \sim p_{\text{data}} \\ x_0 \sim q(\cdot|x, m_0) \\ x_1 \sim f_\theta(\cdot|x_0; m_0) \\ x_2 \sim f_\theta(\cdot|x_1; m_1) \\ \cdots \\ x_{t-1} \sim f_\theta(\cdot|x_{t-2}; m_{t-2})}} \left[\frac{\sum_i (1 - m_{t-1}^{(i)}) \log f_\theta^{(i)}(x^{(i)}|x_{t-1}; m_{t-1})}{\sum_i 1 - m_{t-1}^{(i)}} \right]$$

is the reconstruction loss for the elements of x that were corrupted. As in Orderless NADE, we weigh each term according to the size of the mask, to ensure that the overall weight on each conditional $f_\theta^{(i)}$ is uniform across i. In practice we choose $m_0 = \cdots = m_{t-1}$ – since we only go to $t = 2$, keeping the mask constant is a close enough approximation to the masking schedule used in inference.

Coupled with multi-step unrolling, this training scheme encourages learning complex relationships between the mask and the data, allowing the potential for multi-level trust over the input data: variables with a mask value of 1 which appear correct (given context), variables of mask value 1 which appear incorrect, variables of mask value 0 which appear correct, and variables of mask value 0 which appear incorrect. Denoising only methods (such as SUNDAE [61]) would need to form an internal, non-controllable mask in order to disentangle these states, and 0 mask models (such as Coconet [30]) have controllable input masks but combine both masked states.

This approach allows for direct control at inference using both proposal masks and noising of variables, combining elements of both SUNDAE and Coconet, which we name SUNMASK. We show a high level example of the unrolled training scheme, mask proposals, and input data processing in Figure 1.

This outer unrolled loop is largely independent of architecture choice, and as long as the internal architecture does not make any ordering assumption over the input data we can incorporate it into SUNMASK. We use two primary archetypes for the internal model in this paper: Attentional U-Net and Relative Transformer. Detailed description of the respective architectures can be seen in the Appendix.

2.1 Training

During training, the internal architecture is combined with a *step unrolled* training procedure, as highlighted by SUNDAE [61]. Rather than directly randomizing positions, we re-write this as a masking scheme, first sampling a mask (with 0 randomize, 1 keep, which we denote as 0-active format) then performing randomization to one of P possibilities, for the masked subset of K variables. This random masking procedure is equivalent to the approach from SUNDAE, but using a mask allows us to further combine the mask information with the input data, in order to form a conditional probability estimate. In addition, this 0-active masking scheme makes direct comparison to masking schemes with absorbing states (such as OrderlessNADE [68], Coconet [30], and OA-ARDM [28]) simpler, as the mask can be directly multiplied with the data in a 0-active format.

Convolutional SUNMASK incorporates mask information with a one-hot data representation by concatenation along the channel axis. Transformer SUNMASK uses a slightly different setting - assuming input transformer data (T, B), with T timesteps, B batch elements, and vocabulary size P is transformed to a (T, B, L) dimensional embedding, we repeat the mask along the embedding axis L times, downweighting the values in the mask by $\frac{1}{\sqrt{L}}$ for numerical reasons. Concatenating this reduced mask with the input embedding along the last axis is sufficient to form the desired conditional probability distribution. This stretched and reduced mask format provides more stable training than other schemes such as separately embedding the mask, then concatenating or summing with the input data embedding.

Each training batch is randomly sampled from the training dataset, and a corresponding noise value drawn from $rand(N)$ for N examples in the minibatch. This per-example noise value is then used to derive a per-step mask over T timesteps, by comparing noise $rand(N) < rand(N, T)$. During training, a variable difficulty mask is drawn for each minibatch, with some examples having a high per-example noise value (e.g. .99), and thus many values masked and noised in the training, while other examples may have a low noise value (e.g. .01) drawn instead. This means that, combined with a training loss which learns to denoise the input, the overall model will learn a chain to go from more noisy data to less noisy step-wise, resulting in a learned improvement operator.

This improvement operator, once learned, can be applied to noisy data or pure noise, and iterate toward a predictive sample from the training distribution. See Multinomial Diffusion [29] and SUNDAE [61] for more detail on this proof. In SUNMASK, we combine the mask used to noise the input with the input data itself, while modifying the loss to predict *only masked variables*. In addition, we downweight the loss by $\frac{1}{1+\sum 1-mask}$ for each batch element, meaning that losses for heavily masked entries are downweighted compared to losses on examples with little masking, in a form of curriculum weighting based on expected estimation difficulty.

While a one step denoising schemes can be sufficient for learning the data manifold [3, 1], *unrolling* this denoising scheme into a multi-step process can have performance benefits. SUNMASK directly uses the unrolled loop scheme described in [61], using a step value of 2. For a detailed description of the step unrolled training scheme, see Appendix or the overview description from SUNDAE [61]. The masked and unrolled training can be seen as a container for any internal model which does not make ordering assumptions, and we utilize both convolutional U-Net (a variant of GLIDE [52] U-Net) and Relative Transformer [11, 31, 55] models for various experiments, shown in Section 4.

2.2 Inference

Inference with the learned model can be an involved process, allowing many possible tasks. Because training was done with different masks, a well-trained model should be applicable to content generation, as well as a variety of partially conditional generative tasks such as infilling and human-in-the-loop creation. Basic sampling involves creating a set of variables, with all variables randomly set to 1 of P values in the domain (or partial randomization in the case of infilling) along with an accompanying mask, which is initially all 0 or full generation, or mixed 1s and 0s for partial generation tasks. Given this data and mask as input, trained model then predicts a probability distribution over all possible P values, for all variables. Despite the use of masked losses in training, we sample these prediction distributions for *all* variables. These predictions are then accepted or rejected from the original set, resulting in a new variable set. We then sample a new mask (based on a predefined schedule) and combine it with the initial mask, then iterate this overall process, updating at least some of the variables at each step.

During inference we use several key techniques to improve generative quality. We use typicality sampling [48] on the output probability distribution and a variable number of diffusion steps, on the order of 100 to 2000. Masks are randomly sampled using the schedule defined in [30] which linearly decreases the number of masked variables over time according to $\alpha_n = \max(\alpha_{\min}, \alpha_{\max} - \frac{n}{\eta N}(\alpha_{\max} - \alpha_{\min}))$ with $\alpha_{\min} = .001$, $\alpha_{\max} = .999$, and $\eta = {}^3\!/_4$ as in [30], along with a triangular linear ramp-up and ramp-down schedule for the probability of accepting predictions from the model into the current variable set at each step, as shown in [61]. Acceptance can be further driven to explicitly incorporate the update count by increasing the probability of updating variables which have been updated less often, rather than relying on random exploration which will eventually result in even coverage but may have limitations on small diffusion steps or early in the chain. Variables can also be re-noised at any step, randomly resetting any variable with a 0 value for the mask at that diffusion step. Some problems (primarily symbolic music modeling) showed increased variability from re-noising, but the usefulness of this option is task dependent.

We caution that tuning these hyperparameters for inference is critical to success, as improper settings can drastically lower the performance of SUNMASK, see Section 4 for variance over various inference settings in different tasks. For human-in-the-loop applications, the existence of these controls can allow a number of fine-grained workflows to emerge, driven by expert users to create and curate interesting output [16, 10], demonstrated in Figure 5.

3 Related Work

We state here some key related approaches, as well as how our method differentiates from these previous settings. A number of recent publications on diffusion models and feature learning have incorporated masks as part of their overall training scheme [28, 7, 26], however these papers to use masks for blanking, rather than as indicators over stochastic variables. Many infilling models [15, 34, 12, 45, 9], and masked image models [26] feature conditional modeling with a mask (blank) token, predicting the variables masked from the input for feature learning or generative modeling. XLNet [73] combines the infilling and autoregressive paradigms, learning arbitrary permuted orders

over masked out variables, using blank-out masking and randomly generated autoregressive ordering similar to OrderlessNADE and Coconet. Conditional diffusion generators [49, 59, 60] and GAN generators [18] have the combination of mask indicators as well as preserving stochasticity of the masked variables. However these methods do not use an unrolled training scheme, and generally target image related tasks, with the notable exception of maskGAN. Many models use a concept of a working canvas, and do repeated inference steps for generation or correction of data [23, 4, 42, 19, 51], SUNMASK differs from these models due to architecture choices, training scheme, and loss weighting, as well as application domain [50, 54, 58, 24, 52, 56].

3.1 Convolutional SUNMASK

SUNMASK is most closely related to coconet [30] and SUNDAE [61]. Coconet (as an instance of OrderlessNADE), trains by sampling a random mask per training example, using this mask to set part of the input (in one hot format) to zero. The mask is further concatenated to the zeroed data along the channel axis, and this combined batch is passed through a deep convolutional network with small 3×3 kernels. Coconet is trained to reconstruct the input in whole (categorical loss per variable, over all variables), downweighting losses based on the overall number of active variables in the mask. However, alternate versions of Coconet also experimented with partial prediction using loss only over missing variables. Convolutional SUNMASK has largely the same input structure as Coconet, and uses a downweighted loss over only variables masked in the input. However, SUNMASK additionally uses the unrolled training scheme, as well as a different inference procedure due to preserving the values of masked out variables during training and sampling.

Our best performing convolutional SUNMASK architecture takes hints from recent image transformer and vector quantized generators, exchanging the small kernels used in Coconet for extremely large kernels of shape $4 \times P$ over the time and feature dimensions, somewhat analogous to input patches, removing the model's translation invariance over the feature axis by setting kernel dimension equal to the total feature size. However this makes the number of parameters per convolutional layer extremely large. Convolutional SUNMASK adopts an attentional U-Net structure which reduces only across the time axis, modified from GLIDE [52], rather than the deep residual convolution network used by Coconet. Combined with the addition of step unrolled training, we are only able to train convolutional SUNMASK with a batch size of 1 (expanded to effective batch size 2 due to step unrolling) on commodity GPU hardware with 16GB VRAM.

Due to the design choice of extremely large kernel sizes which depend on the size of the domain, we only use convolutional SUNMASK for polyphonic music experiments, see Section 4 for more details. Exact specification of the convolutional U-Net architecture can be seen in the Appendix.

3.2 Transformer SUNMASK

Transformer SUNMASK relates closely to the transformer used in SUNDAE. The architecture uses a relative multi-head attention [11, 31] and has no autoregressive masking. SUNMASK transformer also uses larger batch sizes, typically 20 or larger. Though this is far smaller than the batch sizes seen in the experiments of SUNDAE. Sequence length and data iterator strategy were both a critical aspect for training transformer SUNMASK. We found short sequences (from 32 to 128) worked best, along with iteration strategies that were example based. In the language experiments, padding each example to some max length resulted in more stable training than the typical language modeling approach of treating the corpus as one long sequence and slicing into even sized chunks, or iterating sequentially. The stability gap between padded sequences and non-overlapping chunking became especially apparent at sequence lengths above 128 with transformer SUNMASK.

We list the hyperparameters for the transformer SUNMASK models in the Appendix. Transformer SUNMASK was trained on every dataset used in this paper, and we show performance in Section 4, as well as comparisons to convolutional SUNMASK on symbolic polyphonic music modeling. Both convolutional and transformer based SUNMASK use the Adam optimizer, with gradient clipping by value at 3. Inference hyperparameter types and general sampling strategies used are the same with both models, though specific hyperparameter values may change between datasets.

Table 1: Quantitative results from the Bach Mock grading function [17] evaluated on Bach ground truth data (Bach GT), BachMock Transformer (BachMock), Coconet, SUNDAE (SD), and SUNMASK convolutional (SMc) and SUNMASK transformer (SMt). Model variants are indicated with Active Balance (A), Typical Sampling (T), and Noise (N). For every grading function in Bach Mock grading function, we show the median value and (standard deviation) as well as the overall grade. Lower values represent better chorales.

Model	Note	Rhythm	Parallel Errors	Harmonic Quality	S Intervals	A Intervals	T Intervals	B Intervals	Repeated Sequence	**Overall**
Bach GT	0.24 (0.15)	0.23 (0.14)	0.0 (0.69)	0.41 (0.2)	0.47 (0.28)	0.49 (0.22)	0.53 (0.24)	0.69 (0.4)	1.29 (0.88)	4.91 (1.63)
BachMock	0.37 (0.22)	0.26 (0.14)	2.16 (3.22)	0.54 (0.31)	0.53 (0.35)	0.71 (0.34)	0.73 (0.38)	0.89 (0.68)	1.86 (2.81)	8.94 (4.64)
Coconet	0.44 (0.23)	1.85 (0.39)	2.61 (6.56)	1.38 (0.39)	0.70 (0.17)	0.86 (0.73)	0.86 (0.42)	1.02 (0.38)	6.07 (1.76)	17.00 (6.58)
SD	0.59 (1.82)	0.93 (0.84)	6.42 (4.11)	0.98 (0.67)	1.17 (5.09)	2.65 (4.08)	1.57 (5.68)	2.57 (3.28)	2.45 (2.39)	23.25 (21.45)
SD-A	0.61 (1.80)	0.93 (0.84)	5.85 (4.02)	0.98 (0.66)	1.28 (5.15)	2.69 (4.20)	1.57 (5.61)	3.28 (3.44)	2.30 (2.38)	22.88 (21.87)
SD-T	0.63 (2.40)	0.60 (0.96)	3.82 (4.98)	0.96 (0.64)	1.21 (5.03)	3.40 (4.99)	3.02 (5.02)	2.36 (3.90)	1.52 (3.43)	20.09 (23.88)
SD-AT	0.52 (2.42)	0.60 (0.95)	3.18 (5.10)	0.96 (0.64)	1.24 (5.00)	3.93 (5.03)	2.22 (5.04)	2.00 (3.91)	1.80 (3.39)	18.90 (24.15)
SMc	0.87 (2.05)	0.63 (0.77)	1.38 (6.00)	1.02 (0.49)	1.41 (5.28)	2.02 (4.36)	1.94 (5.72)	2.91 (4.94)	2.32 (2.31)	22.47 (20.80)
SMc-A	1.02 (2.22)	0.47 (0.77)	3.92 (3.91)	0.91 (0.55)	2.32 (5.23)	3.54 (4.98)	2.74 (5.30)	5.96 (4.59)	2.23 (3.82)	27.82 (18.82)
SMc-T	**0.57 (1.79)**	**0.69 (0.35)**	**1.28 (3.73)**	**0.93 (0.49)**	**0.80 (4.51)**	**0.99 (4.01)**	**1.20 (4.68)**	**1.40 (3.91)**	**1.81 (0.83)**	**13.43 (19.27)**
SMc-AT	0.66 (1.90)	0.55 (0.29)	2.76 (3.63)	0.94 (0.47)	0.91 (4.11)	1.10 (4.00)	1.26 (4.26)	1.45 (4.56)	2.05 (0.96)	16.50 (17.96)
SMc-N	2.07 (2.29)	0.63 (1.33)	7.24 (6.79)	1.06 (0.70)	5.97 (4.77)	7.62 (4.81)	7.37 (4.15)	6.73 (4.92)	1.78 (2.72)	41.67 (21.71)
SMc-AN	2.24 (2.68)	0.69 (0.35)	7.85 (4.28)	1.29 (0.47)	5.08 (4.96)	8.84 (5.41)	8.98 (4.16)	6.96 (4.02)	1.99 (1.27)	47.73 (19.19)
SMc-TN	2.36 (2.94)	0.74 (1.02)	4.42 (6.22)	1.33 (0.70)	6.79 (4.33)	8.19 (4.73)	4.83 (5.06)	6.45 (3.81)	2.36 (2.20)	47.00 (17.32)
SMc-ATN	2.24 (2.36)	0.58 (0.49)	6.82 (4.81)	1.56 (0.54)	6.46 (4.14)	8.51 (4.43)	7.21 (4.28)	7.60 (3.11)	1.47 (1.02)	43.85 (18.41)
SMt	3.00 (1.85)	0.74 (0.90)	0.00 (1.95)	1.64 (0.70)	8.94 (4.66)	6.49 (4.99)	8.47 (5.58)	7.72 (4.41)	3.10 (2.97)	42.87
SMt-A	3.00 (1.85)	0.74 (0.90)	0.00 (1.95)	1.64 (0.70)	8.94 (4.66)	6.49 (4.99)	8.47 (5.58)	7.72 (4.41)	3.10 (2.97)	42.87
SMt-T	3.74 (2.16)	0.58 (0.56)	0.00 (2.56)	1.73 (0.73)	8.75 (4.62)	6.22 (3.99)	8.05 (4.73)	7.95 (4.49)	2.35 (1.79)	46.21 (17.30)
SMt-AT	3.74 (2.16)	0.58 (0.56)	0.00 (2.56)	1.73 (0.73)	8.75 (4.62)	6.22 (3.99)	8.05 (4.73)	7.95 (4.49)	2.35 (1.79)	46.21 (17.30)

4 Experiments

4.1 Quantitative

We demonstrate the use of SUNMASK for polyphonic symbolic music modeling on the JSB dataset [2, 5]. The JSB dataset consists of 382 four-part chorales, originally written by Johann Sebastian Bach. These chorales are quantized at the 16th note interval, cut into non-overlapping chunks of length 128, skipping chunks which would cross the end of a piece. This processing results in a training dataset of 4956 examples, with each example being size $(4, 128)$. We train convolutional and transformer versions of both SUNMASK and SUNDAE for comparison, as well as the pretrained Coconet [30]. For polyphonic music, the quantized data was rasterized in soprano, alto, tenor, bass (SATB) order, as in Music Transformer [31] and BachBot [44], then chunked into non-overlapping training examples. Results are shown in Table 1.

The EMNLP 2017 News dataset is a common benchmark for word-level language modeling [6], containing a large number of news article sentences [46]. Preprocessing steps remove sentences from an initial corpus down to those containing the most common 5700 words, resulting in a training set of $200k$ sentences with a test set of $10k$. The overall maximum sentence length is 51. Common processing for this dataset includes padding all sentences up to this maximum length, which is somewhat different than the standard long sequence chunking commonly used in other language modeling tasks. We show the results of several SUNMASK models, compared to benchmarks using the Negative BLEU/Self-BLEU in generated corpora of 1000 sentences in Figure 2.

The text8 dataset [47] is derived from 1GB of Wikipedia text, cleaned to a 27 character alphabet (a-z and spaces), spelling out numbers as their letter equivalent. This cleaned equivalent is then further truncated to 100MB. This results in small dataset used extensively for testing compression, and more recently generative modeling of text [22]. Although the data doesn't fully match natural language, it is sufficient to test many properties of generative models for text. We train character level language models, using transformer SUNMASK and SUNDAE variants to compare learning on small datasets. Generated samples from SUNMASK can be seen in Table 2.

4.2 Masking For General Task Control

Given the flexibility of masking at inference, we perform a number of qualitative queries to inspect how the model adapts based on noise and mask value. Figures 3 and 5 demonstrate the use of SUNMASK for musical inpainting, holding the top voice (soprano) either fully or partially fixed to the well-known melody "Ode to Joy", by Ludwig van Beethoven. The task of infilling support parts around a central theme is broadly known as harmonization in music composition, and we can specify

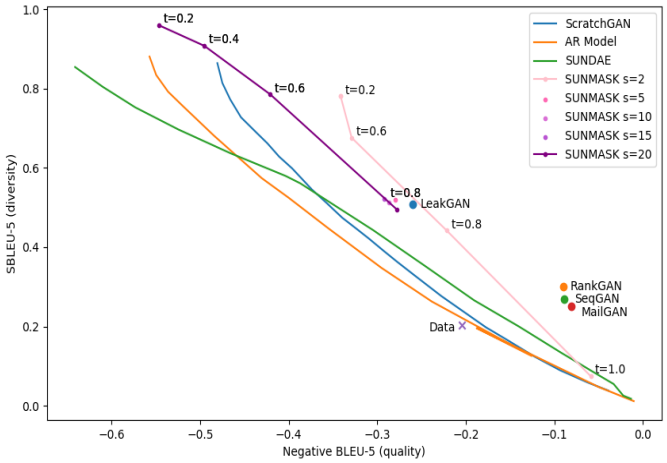

Figure 2: Negative BLEU/Self-BLEU scores on EMNLP2017 News. Left is better, lower is better. Quality/variation is controlled by changing the temperature, and varying diffusion schedule. For SUNMASK, *typical* sampling results [48] are shown. Sampling temperature and schedule are indicated by t and s, respectively.

Table 2: Samples from a Transformer SUNMASK trained on the text8 dataset.

N	Text
1	with a team means voices of word decembers organizes she is one
2	ae capital missions with favour to a place events of a human dol
3	s office and basefully called cases may be inmixed ideated piece
4	in germanic nine borders canada aulained in canadarian german re
5	factures since of the us large human north origin some houses of
6	short houses also remained and as a mathematic province concern
7	attempts despite death six five one k one that six one one nine
8	ed status like any higher language growth ratio battle is down a
9	ts and simply does and a new western and provinces after after s
10	that is often for distinct for a money and control legality the

a mask in order to perform harmonization of this melody. The melodic style of "Ode to Joy" does not fully match the training data, since Beethoven composed in a later era than J.S. Bach. However we see the trained model is more than capable of inpainting based on a pre-defined mask.

We test masks which hold the whole soprano fixed, masks which cover only parts of the soprano but do not allow randomization away from those notes, and masks which cover parts of the soprano but allow rewriting of the non-masked parts of the soprano, as well as rewriting all other voices. The use of masks to focus on subsets of variables while preserving undferlying intermediate predictions is unique to SUNMASK, as SUNDAE does not have an explicitly controllable input mask and Coconet does not have the ability to mask without also blanking the underlying variable. The high level task specification of imputing parts given the rest is shared across all three methods, with both Coconet and SUNDAE demonstrating effective task performance for infilling and completion in their respective publications. All random seeds in this task control experiment were held identical for each task, meaning the underlying changes in the output content are directly driven by changes in the specified mask.

5 Conclusion

We introduce SUNMASK, a method for masked unrolled denoising modeling of structured data. SUNMASK separates the role of masking and correction by conditioning predictions on the mask,

allowing for fine-grained control at inference. When applied to character level and word level text as well as symbolic polyphonic music, SUNMASK is competitive with strong baselines, outperforming reference baselines on music modeling. Leveraging the separation of mask and noise allows for subtle control at inference, paving the way for a variety of domain specific applications and generative pipelines for human-in-the-loop creation.

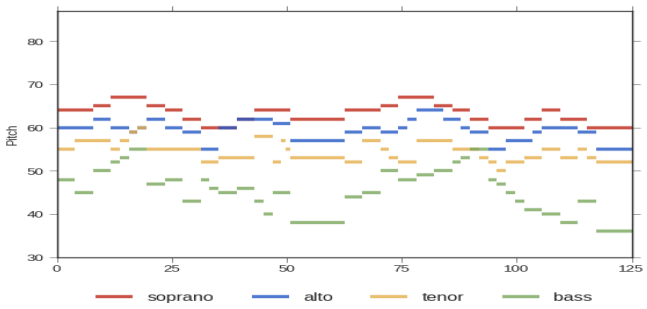

Figure 3: SUNMASK harmonization (bass, tenor, alto) of existing melody (soprano) based on mask.

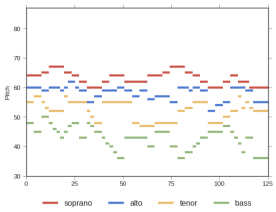

Figure 4: A

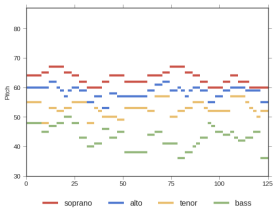

Figure 5: B

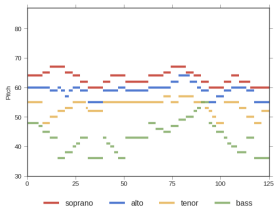

Figure 6: C

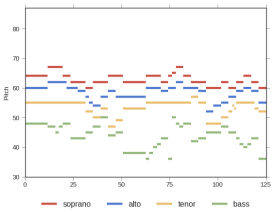

Figure 7: D

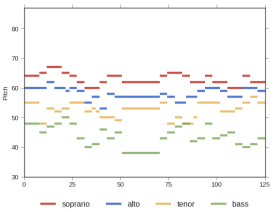

Figure 8: E

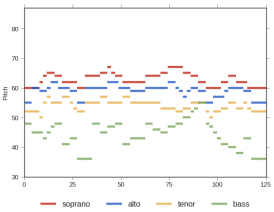

Figure 9: F

A: SUNMASK harmonization (bass, tenor, alto) of an existing melody (soprano), based on pre-defined mask which only highlights the first beat of each measure (0 through 4 of every 16 on the x-axis), but preserves all notes of the melody by overriding predictions in that voice each step.
B: SUNMASK harmonization (bass, tenor, alto) of an existing melody (soprano), based on pre-defined mask which only highlights the first half of the melody (0 through 64 on the x-axis), but preserves all notes of the melody by overriding predictions in that voice at each diffusion step.
C: SUNMASK harmonization (bass, tenor, alto) of an existing melody (soprano), based on pre-defined mask which only highlights the last half of the melody (64 through 128 on the x-axis), but preserves all notes of the melody by overriding predictions in that voice at each diffusion step.
D: SUNMASK harmonization and infilling of an existing melody (soprano), based on pre-defined mask on the first beat of each measure (0 to 4 of every 16 on the x-axis).
E: SUNMASK harmonization and infilling of an existing melody (soprano), based on pre-defined mask which only preserves the first half of the melody (0 to 64 on the x-axis).
F: SUNMASK harmonization and infilling of an existing melody (soprano), based on pre-defined mask which only preserves the last half of the melody (64 through 128 on the x-axis).

References

[1] Guillaume Alain and Yoshua Bengio. What regularized auto-encoders learn from the data-generating distribution. *The Journal of Machine Learning Research*, 15(1):3563–3593, 2014.

[2] Moray Allan and Christopher Williams. Harmonising chorales by probabilistic inference. *Advances in neural information processing systems*, 17, 2004.

[3] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.

[4] Philip Bachman and Doina Precup. Data generation as sequential decision making. *Advances in Neural Information Processing Systems*, 28, 2015.

[5] Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. Modeling temporal dependencies in high-dimensional sequences: application to polyphonic music generation and transcription. In *Proceedings of the 29th International Coference on International Conference on Machine Learning*, pages 1881–1888, 2012.

[6] Massimo Caccia, Lucas Caccia, William Fedus, Hugo Larochelle, Joelle Pineau, and Laurent Charlin. Language gans falling short. *arXiv preprint arXiv:1811.02549*, 2018.

[7] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked generative image transformer. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.

[8] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. In *International Conference on Learning Representations*, 2020.

[9] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.

[10] Katherine Crowson, Stella Biderman, Daniel Kornis, Dashiell Stander, Eric Hallahan, Louis Castricato, and Edward Raff. Vqgan-clip: Open domain image generation and editing with natural language guidance, 2022.

[11] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.

[12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[13] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.

[14] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

[15] Chris Donahue, Mina Lee, and Percy Liang. Enabling language models to fill in the blanks. *arXiv preprint arXiv:2005.05339*, 2020.

[16] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12873–12883, June 2021.

[17] Alexander Fang, Alisa Liu, Prem Seetharaman, and Bryan Pardo. Bach or mock? a grading function for chorales in the style of js bach. *arXiv preprint arXiv:2006.13329*, 2020.

[18] William Fedus, Ian Goodfellow, and Andrew M. Dai. Maskgan: Better text generation via filling in the_____, 2018.

[19] Yaroslav Ganin, Tejas Kulkarni, Igor Babuschkin, SM Ali Eslami, and Oriol Vinyals. Synthesizing programs for images using reinforced adversarial learning. In *International Conference on Machine Learning*, pages 1666–1675. PMLR, 2018.

[20] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pages 881–889. PMLR, 2015.

[21] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[22] Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.

[23] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. In *International Conference on Machine Learning*, pages 1462–1471. PMLR, 2015.

[24] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. *arXiv preprint arXiv:2111.14822*, 2021.

[25] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010.

[26] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv:2111.06377*, 2021.

[27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.

[28] Emiel Hoogeboom, Alexey A Gritsenko, Jasmijn Bastings, Ben Poole, Rianne van den Berg, and Tim Salimans. Autoregressive diffusion models. *arXiv preprint arXiv:2110.02037*, 2021.

[29] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34, 2021.

[30] Cheng-Zhi Anna Huang, Tim Cooijmans, Adam Roberts, Aaron Courville, and Douglas Eck. Counterpoint by convolution. *arXiv preprint arXiv:1903.07227*, 2019.

[31] Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, Ian Simon, Curtis Hawthorne, Noam Shazeer, Andrew M Dai, Matthew D Hoffman, Monica Dinculescu, and Douglas Eck. Music transformer: Generating music with long-term structure. In *International Conference on Learning Representations*, 2018.

[32] Chin-Wei Huang, Jae Hyun Lim, and Aaron C Courville. A variational perspective on diffusion-based generative models and score matching. *Advances in Neural Information Processing Systems*, 34, 2021.

[33] Rongjie Huang, Max WY Lam, Jun Wang, Dan Su, Dong Yu, Yi Ren, and Zhou Zhao. Fastdiff: A fast conditional diffusion model for high-quality speech synthesis. *arXiv preprint arXiv:2204.09934*, 2022.

[34] Daphne Ippolito, Anna Huang, Curtis Hawthorne, and Douglas Eck. Infilling piano performances. In *NIPS Workshop on Machine Learning for Creativity and Design*, 2018.

[35] Nal Kalchbrenner, Ivo Danihelka, and Alex Graves. Grid long short-term memory. *arXiv preprint arXiv:1507.01526*, 2015.

[36] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aaron Oord, Sander Dieleman, and Koray Kavukcuoglu. Efficient neural audio synthesis. In *International Conference on Machine Learning*, pages 2410–2419. PMLR, 2018.

[37] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In *Advances in Neural Information Processing Systems*, 2021.

[38] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[39] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29, 2016.

[40] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2020.

[41] Max WY Lam, Jun Wang, Rongjie Huang, Dan Su, and Dong Yu. Bilateral denoising diffusion models. *arXiv preprint arXiv:2108.11514*, 2021.

[42] Alex M Lamb, Devon Hjelm, Yaroslav Ganin, Joseph Paul Cohen, Aaron C Courville, and Yoshua Bengio. Gibbsnet: Iterative adversarial inference for deep graphical models. *Advances in Neural Information Processing Systems*, 30, 2017.

[43] Jason Lee, Elman Mansimov, and Kyunghyun Cho. Deterministic non-autoregressive neural sequence modeling by iterative refinement. In *EMNLP*, 2018.

[44] Feynman Liang. Bachbot: Automatic composition in the style of bach chorales. *University of Cambridge*, 8:19–48, 2016.

[45] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

[46] Sidi Lu, Yaoming Zhu, Weinan Zhang, Jun Wang, and Yong Yu. Neural text generation: Past, present and beyond. *arXiv preprint arXiv:1803.07133*, 2018.

[47] Matthew V Mahoney. Large text compression benchmark. `http://mattmahoney.net/dc/text.html`, 2009.

[48] Clara Meister, Tiago Pimentel, Gian Wiher, and Ryan Cotterell. Typical decoding for natural language generation. *arXiv preprint arXiv:2202.00666*, 2022.

[49] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021.

[50] Gautam Mittal, Jesse Engel, Curtis Hawthorne, and Ian Simon. Symbolic music generation with diffusion models. *arXiv preprint arXiv:2103.16091*, 2021.

[51] Alex Nichol. Vq-draw: A sequential discrete vae. *arXiv preprint arXiv:2003.01599*, 2020.

[52] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.

[53] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30, 2017.

[54] Ashis Pati, Alexander Lerch, and Gaëtan Hadjeres. Learning to traverse latent spaces for musical score inpainting. *arXiv preprint arXiv:1907.01164*, 2019.

[55] Christine Payne. Musenet. *openai.com/blog/musenet*, 2019.

[56] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022.

[57] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.

[58] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.

[59] Chitwan Saharia, William Chan, Huiwen Chang, Chris A Lee, Jonathan Ho, Tim Salimans, David J Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. *arXiv preprint arXiv:2111.05826*, 2021.

[60] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021.

[61] Nikolay Savinov, Junyoung Chung, Mikolaj Binkowski, Erich Elsen, and Aaron van den Oord. Step-unrolled denoising autoencoders for text generation. *arXiv preprint arXiv:2112.06749*, 2021.

[62] Raphael Shu, Jason Lee, Hideki Nakayama, and Kyunghyun Cho. Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8846–8853, 2020.

[63] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.

[64] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems*, 34, 2021.

[65] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019.

[66] Lucas Theis and Matthias Bethge. Generative image modeling using spatial lstms. *Advances in neural information processing systems*, 28, 2015.

[67] Benigno Uria, Marc-Alexandre Côté, Karol Gregor, Iain Murray, and Hugo Larochelle. Neural autoregressive distribution estimation. *The Journal of Machine Learning Research*, 17(1):7184–7220, 2016.

[68] Benigno Uria, Iain Murray, and Hugo Larochelle. A deep and tractable density estimator. In *International Conference on Machine Learning*, pages 467–475. PMLR, 2014.

[69] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR, 2016.

[70] Sean Vasquez and Mike Lewis. Melnet: A generative model for audio in the frequency domain. *arXiv preprint arXiv:1906.01083*, 2019.

[71] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[72] Francesco Visin, Kyle Kastner, Kyunghyun Cho, Matteo Matteucci, Aaron Courville, and Yoshua Bengio. Renet: A recurrent neural network based alternative to convolutional networks. *arXiv preprint arXiv:1505.00393*, 2015.

[73] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.

[74] Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. *arXiv preprint arXiv:2204.13902*, 2022.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[TODO]**[Yes]
   (b) Did you describe the limitations of your work? **[TODO]**[Yes] Throughout the paper we list downsides to including the mask in the input, as well as the complexity of our inference pipeline
   (c) Did you discuss any potential negative societal impacts of your work? **[TODO]**[Yes] We plan to address this in the supplementary material and appendix due to space limitations
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[TODO]**[Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? **[TODO]**[Yes] Our theorems state the relevant conditions
   (b) Did you include complete proofs of all theoretical results? **[TODO]**[Yes] We also cite relevant prior work on said theorems
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[TODO]**[No] We will include a (likely non-runnable) version of the experimental code in the supplemental material, and work toward a more generally usable release during the review cycle
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[TODO]**[No] To fully disclose this, requires full experimental code
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[TODO]**[Yes] Variance in the scoring table
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[TODO]**[No] We discuss the GPU type, but do not do a full accounting of resources used. All experiments were trained on commodity hardware (P100/V100/A100) using single GPUs, for less than 24 hours an experiment run
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? **[TODO]**[Yes]
   (b) Did you mention the license of the assets? **[TODO]**[No] All material is extensively used in prior work
   (c) Did you include any new assets either in the supplemental material or as a URL? **[TODO]**[No] No new assets
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[TODO]**[No] No new assets, and assets used are public domain
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[TODO]**[No]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[TODO]**[N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[TODO]**[N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[TODO]**[N/A]