

Towards End-to-End Generative Modeling of Long Videos with Memory-Efficient Bidirectional Transformers

Jaehoon Yoo¹, Semin Kim¹, Doyup Lee², Chiheon Kim², Seunghoon Hong¹

¹KAIST, ²Kakao Brain

{wogns98, kimsm1121, seunghoon.hong}@kaist.ac.kr, {damien.re, frost.conv}@kakaobrain.com

Abstract

Autoregressive transformers have shown remarkable success in video generation. However, the transformers are prohibited from directly learning the long-term dependency in videos due to the quadratic complexity of self-attention, and inherently suffering from slow inference time and error propagation due to the autoregressive process. In this paper, we propose Memory-efficient Bidirectional Transformer (MeBT) for end-to-end learning of long-term dependency in videos and fast inference. Based on recent advances in bidirectional transformers, our method learns to decode the entire spatio-temporal volume of a video in parallel from partially observed patches. The proposed transformer achieves a linear time complexity in both encoding and decoding, by projecting observable context tokens into a fixed number of latent tokens and conditioning them to decode the masked tokens through the cross-attention. Empowered by linear complexity and bidirectional modeling, our method demonstrates significant improvement over the autoregressive transformers for generating moderately long videos in both quality and speed. Videos and code are available at <https://sites.google.com/view/mebt-cvpr2023>.

1. Introduction

Modeling the generative process of videos is an important yet challenging problem. Compared to images, generating convincing videos requires not only producing high-quality frames but also maintaining their semantic structures and dynamics coherently over long timescale [11, 16, 27, 36, 38, 54].

Recently, autoregressive transformers on discrete representation of videos have shown promising generative modeling performances [11, 32, 51, 53]. Such methods generally involve two stages, where the video frames are first turned into discrete tokens through vector quantization, and then their sequential dynamics are modeled by an autoregressive transformer. Powered by the flexibility of discrete distri-

butions and the expressiveness of transformer architecture, these methods demonstrate impressive results in learning and synthesizing high-fidelity videos.

However, autoregressive transformers for videos suffer from critical scaling issues in both training and inference. During training, due to the quadratic cost of self-attention, the transformers are forced to learn the joint distribution of frames entirely from short videos (e.g., 16 frames [11, 53]) and cannot directly learn the statistical dependencies of frames over long timescales. During inference, the models are challenged by two issues of autoregressive generative process – its serial process significantly slows down the inference speed, and perhaps more importantly, autoregressive prediction is prone to error propagation where the prediction error of the frames accumulates over time.

To (partially) address the issues, prior works proposed improved transformers for generative modeling of videos, which are categorized as the following: **(a)** Employing sparse attention to improve scaling during training [12, 16, 51], **(b)** Hierarchical approaches that employ separate models in different frame rates to generate long videos with a smaller computation budget [11, 16], and **(c)** Removing autoregression by formulating the generative process as masked token prediction and training a bidirectional transformer [12, 13]. While each approach is effective in addressing specific limitations in autoregressive transformers, none of them provides comprehensive solutions to aforementioned problems – **(a, b)** still inherits the problems in autoregressive inference and cannot leverage the long-term dependency by design due to the local attention window, and **(c)** is not appropriate to learn long-range dependency due to the quadratic computation cost. We believe that an alternative that jointly resolves all the issues would provide a promising approach towards powerful and efficient video generative modeling with transformers.

In this work, we propose an efficient transformer for video synthesis that can fully leverage the long-range dependency of video frames during training, while being able to achieve fast generation and robustness to error propagation. We achieve the former with a linear complexity ar-

chitecture that still imposes *dense* dependencies across all timesteps, and the latter by removing autoregressive serialization through masked generation with a bidirectional transformer. While conceptually simple, we show that efficient dense architecture and masked generation are highly complementary, and when combined together, lead to substantial improvements in modeling longer videos compared to previous works both in training and inference. The contributions of this paper are as follows:

- We propose Memory-efficient Bidirectional Transformer (MeBT) for generative modeling of video. Unlike prior methods, MeBT can directly learn long-range dependency from training videos while enjoying fast inference and robustness in error propagation.
- To train MeBT for moderately long videos, we propose a simple yet effective curriculum learning that guides the model to learn short- to long-term dependencies gradually over training.
- We evaluate MeBT on three challenging real-world video datasets. MeBT achieves a performance competitive to state-of-the-arts for short videos of 16 frames, and outperforms all for long videos of 128 frames while being considerably more efficient in memory and computation during training and inference.

2. Background

This section introduces generative transformers for videos that utilize discrete token representation, which can be categorized into autoregressive and bidirectional models.

Let $\mathbf{x} \in \mathbb{R}^{T \times H \times W \times 3}$ be a video. To model its generative distribution $p(\mathbf{x})$, prior works on transformers employ discrete latent representation of frames $\mathbf{y} \in \mathbb{R}^{t \times h \times w \times d}$ and model the prior distribution on the latent space $p(\mathbf{y})$.

Vector Quantization To map a video \mathbf{x} into discrete tokens \mathbf{y} , previous works [11, 13, 32, 48, 51] utilize vector quantization with an encoder E that maps \mathbf{x} onto a learnable codebook $F = \{e_i\}_{i=1}^U$ [45]. Specifically, given a video \mathbf{x} , the encoder produces continuous embeddings $\mathbf{h} = E(\mathbf{x}) \in \mathbb{R}^{t \times h \times w \times d}$ and searches for the nearest code $e_u \in F$.

The encoder E is trained through an autoencoding framework by introducing a decoder D that takes discrete tokens \mathbf{y} and produces reconstruction $\hat{\mathbf{x}} = D(\mathbf{y})$. The encoder E , codebook F , and the decoder D are optimized with the following training objective:

$$\mathcal{L}_q = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 + \|\text{sg}(\mathbf{h}) - \mathbf{y}\|_2^2 + \beta \|\text{sg}(\mathbf{y}) - \mathbf{h}\|_2^2, \quad (1)$$

where sg denotes stop-gradient operator. In practice, to improve the quality of discrete representations, additional perceptual loss and adversarial loss are often introduced [9].

For the choice of the encoder E , we utilize 3D convolutional networks that compress a video in both spatial and temporal dimensions following prior works [11, 53].

Autoregressive Transformers Given the discrete latent representations $\mathbf{y} \in \mathbb{R}^{t \times h \times w \times d}$, generative modeling of videos boils down to modeling the prior $p(\mathbf{y})$. Prior work based on autoregressive transformers employ a sequential factorization $p(\mathbf{y}) = \prod_{i \leq N} p(y_i | y_{<i})$ where $N = thw$, and use a transformer to model the conditional distribution of each token $p(y_i | y_{<i})$. The transformer is trained to minimize the following negative log-likelihood of training data:

$$\mathcal{L}_a = \sum_{i \leq N} -\log p(y_i | y_{<i}). \quad (2)$$

During inference, the transformer generates a video by sequentially sampling each token y_i from the conditional $p(y_i | y_{<i})$ based on context $y_{<i}$. The sampled tokens \mathbf{y} are then mapped back to a video using the decoder D .

While simple and powerful, autoregressive transformers for videos suffer from critical scaling issues. First, each conditional $p(y_n | y_{<n})$ involves $\mathcal{O}(n^2)$ computational cost due to the quadratic complexity of self-attention. This forces the model to only utilize short-term context in both training and inference, making it inappropriate in modeling spatio-temporal long-term coherence. Furthermore, during inference, the sequential decoding requires N model predictions that recursively depend on the previous one. This leads to a slow inference, and more notably, a potential error propagation over space and time since the prediction error at a certain token accumulates over the remaining decoding steps. This is particularly problematic for videos since N is often very large as tokenization spans both spatial and temporal dimensions.

Bidirectional Transformers To improve the decoding efficiency of autoregressive transformers, bidirectional generative transformers have been proposed [5, 13, 56]. Contrary to autoregressive models that predict a single consecutive token at each step, a bidirectional transformer learns to predict *multiple* masked tokens at once based on the previously generated context. Specifically, given the random masking indices $\mathbf{m} \subseteq \{1, \dots, N\}$, it models the joint distribution over masked tokens $\mathbf{y}_M = \{y_i | i \in \mathbf{m}\}$ conditioned on the visible context $\mathbf{y}_C = \{y_i | i \notin \mathbf{m}\}$, and is trained with the below objective:

$$\mathcal{L}_b = -\log p(\mathbf{y}_M | \mathbf{y}_C, \mathbf{z}_M) \approx -\sum_{i \in \mathbf{m}} -\log p(y_i | \mathbf{y}_C, \mathbf{z}_M), \quad (3)$$

where mask embeddings \mathbf{z}_M encode positions of the mask tokens with learnable vectors. Each conditional in Eq. (3) is modeled by a transformer, but contrary to autoregressive

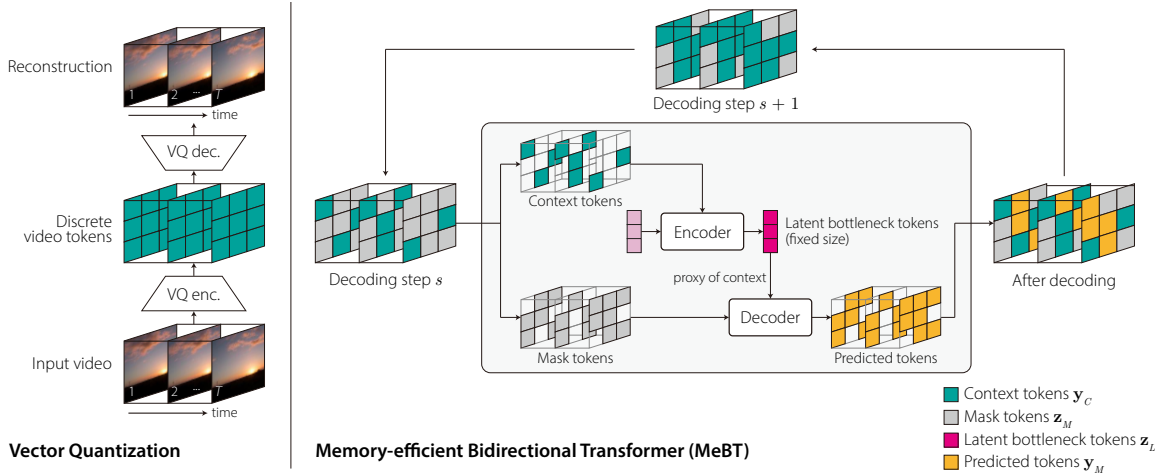


Figure 1. Overview of our method. Our model learns to predict masked tokens from the context tokens with linear complexity encoder and decoder. The encoder and decoder utilize latent bottlenecks to achieve linear complexity while performing full dependency modeling.

models that apply causal masking, bidirectional transformers operate on entire tokens $\mathbf{y} = \mathbf{y}_C \cup \mathbf{y}_M$. During inference, the model is initialized with empty context (*i.e.*, $\mathbf{y}_C = \emptyset$), and performs iterative decoding by predicting the probability over all masked tokens \mathbf{y}_M by Eq. (3) and sampling their subset as the context for the next iteration.

Adopting bidirectional transformers to model videos has two advantages compared to autoregressive transformers. First, by decoding multiple tokens at once, bidirectional transformers enjoy a better parallelization and smaller decoding steps than autoregressive transformers which results in a faster sampling. Second, bidirectional transformers are more robust to error propagation than autoregressive counterparts since the decoding of the masked token is independent of the temporal order, allowing consistent prediction quality over time. However, complexity of the bidirectional transformers is still quadratic to the number of tokens, limiting the length of training videos to only short sequences thus hindering learning long-term dependencies.

3. Memory-efficient Bidirectional Transformer

Our goal is to design a generative transformer for videos that can perform fast inference with a robustness to error propagation, while being able to fully leverage the long-range statistical dependency of video frames in training.

For inference speed and robustness, we adopt the bidirectional approach and parameterize the joint distribution of masked tokens $p(\mathbf{y}_M | \mathbf{y}_C, \mathbf{z}_M) \approx \prod_{i \in \mathbf{m}} p(y_i | \mathbf{y}_C, \mathbf{z}_M)$ (Eq. (3)) with a transformer. For learning long-range dependency, we take a simple approach of employing an efficient transformer architecture [40] of sub-quadratic complexity and directly training it with longer videos. While sparse attention (*e.g.*, local, axial, strided) is dominant in autoregressive transformers to reduce complexity [11, 12, 16, 51], we find it potentially problematic for bidirectional transformers

since the context \mathbf{y}_C can be provided for an arbitrary subset of token positions that often cannot be covered by the fixed sparse attention patterns¹. Thus, unlike prior work, we design an efficient bidirectional transformer based on low-rank latent bottleneck [17, 18, 21, 25, 30, 49] that always enables a *dense* attention dependency over all token positions while guaranteeing *linear* complexity.

To this end, we propose Memory-efficient Bidirectional Transformer (MeBT) for generative modeling of videos. The overall framework of MeBT is illustrated in Fig. 1. To predict the masked tokens \mathbf{y}_M from the context tokens \mathbf{y}_C , MeBT employs an encoder-decoder architecture based on a fixed number of latent bottleneck tokens $\mathbf{z}_L \in \mathbb{R}^{N_L \times d}$ with $N_L \ll N$. The encoder projects the context tokens \mathbf{y}_C to the latent tokens \mathbf{z}_L , and the decoder utilizes the latent tokens \mathbf{z}_L to predict the masked tokens \mathbf{y}_M . Overall, the encoder and decoder exploit the latent bottleneck to achieve a linear complexity $\mathcal{O}(N)$ while performing a full dependency modeling across the token positions. In the following sections, we describe the details of the encoder, decoder, and the training and inference procedures of MeBT.

3.1. Encoder Architecture

The encoder aims to project all previously generated context tokens \mathbf{y}_C to fixed-size latent bottleneck \mathbf{z}_L with a memory and time cost linear to context size $\mathcal{O}(N_C)$. Following prior work on transformers with latent bottlenecks [17, 18, 21, 25, 30], we construct the encoder as an alternating stack of two types of attention layers that progressively update the latent tokens based on the provided context tokens. Specifically, as illustrated in Fig. 2, the first layer updates the latent tokens by cross-attending to context tokens, and the second one updates the latent tokens by per-

¹Note that, in autoregressive transformers, the context $\mathbf{y}_{<i}$ (Eq. (2)) is always the entire past, allowing sparse attention to robustly see the context.




Figure 2. Encoder architecture. Our encoder compresses the context tokens into fixed number of latent tokens.

forming self-attention within them. By stacking the layers, the encoder progressively gathers appropriate contextual information into the latent tokens as well as flexibly updating their state through the exchanges of latent information.

Our encoder with latent bottleneck can benefit bidirectional modeling of videos in several ways. First, since we employ the fixed-size latent tokens independent to the video length, the complexity of the encoder scales linearly to the number of context tokens, allowing it to directly learn and model long-range dependency by taking longer videos with a reduced memory footprint. More importantly, unlike the sparse attention [12, 15] or hierarchical methods [11, 16] that capture partial dependency among the predefined local subset of context tokens, our encoder always captures full observations of any given set of context tokens by exploiting the latent bottleneck. This property suits well with the bidirectional modeling since the context tokens can be given as an arbitrary subset of a video often beyond the hand-designed patterns of sparse attention.

3.2. Decoder Architecture

Given the latent bottleneck \mathbf{z}_L as a proxy of context², the decoder aims to predict the masked tokens \mathbf{y}_M by updating the mask embeddings \mathbf{z}_M . Since the standard masked modeling [5, 8, 14] of applying a bidirectional transformer on mask and context $\{\mathbf{z}_M, \mathbf{z}_L\}$ leads to a complexity of $\mathcal{O}((N_M + N_L)^2)$, we opt into an efficient decoder that reduces the cost to linear to mask size $\mathcal{O}(N_M)$ while preserving as much modeling capacity as possible.

We construct our decoder as a stack of two types of attention layers that update the latent tokens \mathbf{z}_L and mask tokens \mathbf{z}_M in an alternating manner. More specifically, as illustrated in Fig. 3, the first layer updates latent tokens \mathbf{z}_L by attending to *both* latent and mask $\{\mathbf{z}_L, \mathbf{z}_M\}$, and the second layer updates mask tokens \mathbf{z}_M by attending to the latent tokens \mathbf{z}_L . After being processed by the decoder, the mask tokens \mathbf{z}_M are then used to predict the masked tokens \mathbf{y}_M .

Our decoder achieves linear complexity while retaining a high expressive power, thus being suitable for directly learn-

²Since the latent tokens \mathbf{z}_L are a proxy of context, we use *latent* and *context* interchangeably to refer to them when describing the decoder.

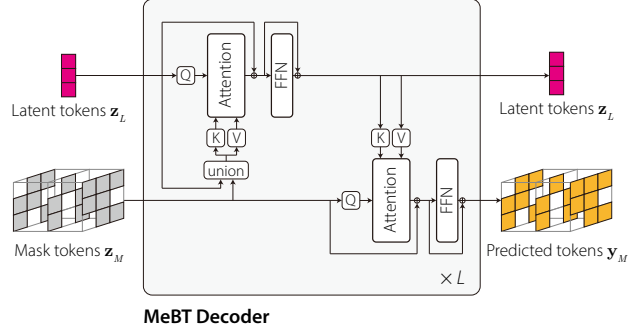


Figure 3. Decoder architecture. Our decoder predicts the masked tokens based on the latent tokens encoding the entire context.

ing and generating long-term dependency. Since the two attention types have $\mathcal{O}(N_L \times (N_L + N_M))$ and $\mathcal{O}(N_L N_M)$ complexities respectively, our decoder overall achieves a linear complexity to mask size N_M . Still, we note that this efficient decoder retains much of the modeling capacity of standard masked decoders [5, 8, 14] that employ full self-attention on context and mask $\{\mathbf{z}_L, \mathbf{z}_M\}$. This is seen by decomposing their *all-to-all* dependency $\{\mathbf{z}_L, \mathbf{z}_M\} \rightarrow \{\mathbf{z}_L, \mathbf{z}_M\}$ into *all-to-context* $\{\mathbf{z}_L, \mathbf{z}_M\} \rightarrow \mathbf{z}_L$, *context-to-mask* $\mathbf{z}_L \rightarrow \mathbf{z}_M$, and *mask-to-mask* $\mathbf{z}_M \rightarrow \mathbf{z}_M$ dependencies. In our efficient decoder, the all-to-context and context-to-mask dependencies are precisely modeled by the two types of attention layers, and the mask-to-mask dependency is approximately modeled through the *composition* of the layers by exploiting the latent tokens \mathbf{z}_L as an intermediate bottleneck. More specifically, to model $\mathbf{z}_M \rightarrow \mathbf{z}_M$, our decoder imposes a low-rank approximation through the fixed-size latent representation $\mathbf{z}_M \rightarrow \mathbf{z}_L \rightarrow \mathbf{z}_M$. While such approximation appears in a range of prior work on efficient transformers [21, 25, 30, 49], to our knowledge, we are the first to employ it for bidirectional video generation.

3.3. Training and Inference

Training To train bidirectional transformers for masked token prediction objective in Eq. (3), prior works [5, 12] employed random sampling of the masking indices \mathbf{m} . Specifically, it choose masking ratio $r \in (0, 1]$ according to a predefined schedule function, and sample $N_M = \lceil rN \rceil$ unique masking indices randomly by $\mathbf{m} \sim \text{Uniform}(1, N)$. Although such strategy works well with images or short videos, we observe that training with the strategy directly on moderately long videos often leads to suboptimal solution. This is presumably because sampling temporally distant tokens at early stage of training makes the task too difficult, preventing model to learn important local structures.

To resolve the issue, we propose a simple curriculum-based training that guides the model to learn from short- to long-term sequences gradually over training. To this end, we define an *interval* scheduling function that dictates the lengths (intervals) v of training videos at a particular train-




Figure 4. Visualization of the interval scheduling function. ing iteration n by:

$$p_n(v) \propto \exp\left(-\frac{(v-1-\frac{n}{\alpha})^2}{2\beta^2}\right) \quad (4)$$

where α and β are hyperparameters. Eq. (4) characterizes a Gaussian distribution of training video lengths v whose mean gradually increases by the training iterations n with a speed determined by α (Fig. 4). At every training iteration n , we sample the interval $v \in (1, t]$ by first sampling $\hat{v} \sim p_n(v)$ and truncating it within a valid range by $v = \min(\max(1, \lceil \hat{v} \rceil), t)$. Then we sample v consecutive frames randomly in a video, and sample the masking indices \mathbf{m} and the context tokens randomly within the sampled interval as discussed above. We found that such curriculum learning is effective in stabilizing training and improving the prediction quality for long videos (Section 5.4).

Inference During inference, MeBT generates videos by iteratively decoding subsets of mask tokens. To this end, we define a mask sampling function γ on $[0, 1]$ that decreases from $\gamma(0) = 1$ to $\gamma(1) = 0$. At each decoding step $s \in \{0, \dots, S-1\}$, our method produces predictions on the mask tokens by Eq. (3), and updates the context tokens with top N_s mask tokens according to their sampled probability where $N_s = N - \lceil \gamma(\frac{s}{S})N \rceil$. The process is iteratively applied S times until all tokens are included in the context. In addition to the naïve decoding, we also adopt the revision phase proposed in [24]. In this stage, we randomly divide the indices evenly into R partitions $\Pi = \{\mathbf{m}_1, \dots, \mathbf{m}_R\}$, and revise the tokens in \mathbf{m}_i based on the contexts. We found that the additional revision is helpful to improve the consistency and fidelity of generated videos.

4. Related Work

Video Generative Models Video generative models have shown rapid growth with the advance of generative frameworks. Various approaches based on Generative Adversarial Networks (GAN) [6, 19, 33, 34, 41, 43], Variational Autoencoders (VAE) [2, 7, 10, 23, 36], and Implicit Neural Representations (INR) [38, 54] are proposed to extend the image generative models to powerful video sequence models. Despite the impressive success, these methods are mostly displayed in simple and short videos, while progress in modeling complex scenes and motions is still behind ones in the image domains. Recently, autoregressive transformers [11, 16, 32, 51, 53] demonstrated promising results in

modeling complex videos, closing the gap between the image and video generative models. However, they are mostly limited in short-term videos due to quadratic computation, and inherently suffer from slow inference time and error propagation due to autoregressive process. This motivates us to seek efficient yet robust transformers for long-term videos. Long-term video generation has been studied under various context, often conditioned on powerful side information such as real frames [2, 4, 10, 26, 36, 50], dense pixel labels [1, 27, 28], action sequences [22, 46], or text description [16, 51]. Apart from these, we study the potential of transformers for unconditional long-term video generation.

Efficient Transformers To improve the scalability of transformers for long sequences, a large volume of work has been developed to approximate the quadratic computation. One promising direction is sparse attention that approximates dense pair-wise attention by restricting the size of the neighborhood that each token can attend to [3, 12, 15, 16, 51, 55]. However, the fixed sparsity patterns are often incomprehensive to cover various token-wise dependencies in broad applications, especially when combined with masked bidirectional transformers with high masking ratio. The other line of work proposed latent bottleneck [17, 18, 21, 25, 30] that projects the input to fewer latent embeddings. Latent bottleneck can be easily combined with bidirectional transformers as it models dense dependencies between tokens, and this work is the first attempt to demonstrate it in videos.

5. Experiments

5.1. Experimental Setup

Datasets We evaluate our method on three popular video benchmarks; SkyTimelapse [52], Taichi-HD [37], and UCF-101 [39]. For each dataset, we use 16- and 128-frame-long videos for evaluation of short-term and long-term generation, respectively. Following prior works, all videos are pre-processed with 128×128 spatial resolution. More details are described in the Appendix (Section A.2).

Evaluation metrics We evaluate the quality of the generated videos using Fréchet Video Distance (FVD) [44] and Kernel Video Distance (KVD) [44] based on an I3D network trained on Kinetics-600 [20]. For 16-frame videos, we follow the evaluation protocol of [34] and report the average score over 10 runs where each run samples 2048 videos. For 128-frame videos, we sample 512 videos 5 times and report the average score. For UCF-101 dataset, we also measure the Inception Score (IS) [35] with a pre-trained C3D model [42] following the prior works [11, 31, 47]. To evaluate the generation quality over time on longer videos (128 frames), we also measure FVD for every 8 frames using a sliding window of 16 frame length. Furthermore, we evalu-

Table 1. Quantitative results on 16-frame video generation. * denotes the models trained on both training and validation splits.

(a) SkyTimelapse					(c) UCF-101		
Method	Memory (GB)	Inf.Time	FVD ₁₆ (↓)	KVD ₁₆ (↓)	Method	IS (↑)	FVD ₁₆ (↓)
MoCoGAN-HD	22.4	0.04s	183.6 \pm 5.2	13.9 \pm 0.7	TGAN	11.85 \pm .07	-
DIGAN	4.8	0.04s	114.6 \pm 4.9	6.8 \pm 0.5	LDVD-GAN	22.91 \pm .19	-
TATS-base	21.9	14.3s	132.6 \pm 2.6	5.7 \pm 0.3	VideoGPT	24.69 \pm .30	-
MeBT (Ours)	9.5	1.10s	90.4 \pm 1.8	2.6 \pm 0.1	MoCoGAN-HD*	32.36	838
(b) Taichi-HD					DIGAN	29.71 \pm .53	655 \pm 22
Method	Memory (GB)	Inf.Time	FVD ₁₆ (↓)	KVD ₁₆ (↓)	DIGAN*	32.70 \pm .35	577 \pm 21
MoCoGAN-HD	22.4	0.04s	144.7 \pm 6.0	25.4 \pm 1.9	CCVS*+StyleGAN	24.47 \pm .13	386 \pm 15
DIGAN	4.8	0.04s	128.1 \pm 4.9	20.6 \pm 1.1	StyleGAN-V*	23.94 \pm .73	-
TATS-base	21.9	14.3s	94.6 \pm 2.7	9.8 \pm 1.0	Video Diffusion*	57.00 \pm .62	-
MeBT (Ours)	9.5	1.36s	122.7 \pm 4.9	16.9 \pm 1.11	TATS-base	57.63 \pm .24	420 \pm 18
					Real data	90.52	-
					MeBT (Ours)	65.93 \pm .79	438 \pm 18

ate the efficiency of the models by measuring the peak training memory and inference time with batch size 4.

Implementation For discrete tokenization, we use the official VQGAN checkpoints from [11] that have a compression ratio of 4 and 8 in temporal and spatial dimensions, respectively. For a fair comparison, we match the number of parameters and layers in our method to be roughly the same as other transformer-based baseline [11]. We set the size of the latent bottleneck to $N_L = 256$ for every experiment. We train 16-frame MeBT without curriculum while 128-frame MeBT models are trained with the proposed learning schedule in Section 3.3. For the interval scheduling function, we use $\beta = 2$ for all datasets while $\alpha = 30K$ for SkyTimelapse and $\alpha = 100K$ for Taichi-HD and UCF-101. More details on training are described in Appendix (Section A.1).

5.2. Results on Short-term Video Generation

Baselines We compare our method with state-of-the-art video generative models whose pre-trained models are publicly available. In all datasets, we employ MoCoGAN-HD [41] and DIGAN [54] as methods based on Generative Adversarial Networks (GAN), and TATS [11] as the state-of-the-art autoregressive transformer that differs only in architecture while sharing the same quantization codebook with ours. We also present comparisons to other unconditional models in UCF-101 using the reported numbers.

Results Table 1 summarizes the results on short-term video generation (16 frames). Overall, we observe that our method exhibits competitive performance to the state-of-the-art approaches. Our method outperforms all non-transformer baselines by a large margin in all datasets, and competitive performance to other autoregressive transformers (TATS and CCVS); it achieves the best and the second-best performance in SkyTimelapse and Taichi-HD, respectively, and the best Inception score in UCF-101. Compared to TATS, however, our method exhibits considerable im-

provement in efficiency in both training and inference; it reduces the peak training memory usage by 33-45% while improving the inference speed by 10-20 times. Note that the gap increases rapidly in longer videos as the asymptotic complexity is quadratic in TATS but linear in MeBT. Overall, our results show that MeBT provides compelling short-term generation performance while reducing computation cost significantly. It allows our model to be trained with longer videos to model interesting long-term dynamics, which is described in the next section.

5.3. Results on Long-term Video Generation

Baselines We evaluate our method with five strong baselines: MoCoGAN-HD [41], CCVS [31]³, and TATS [11] as strong autoregressive approaches that recursively generate long videos conditioned on the previous generation, and DIGAN [54] as an implicit network that directly generates each frame independently conditioned on latent variable. We also compare with the hierarchical variants of TATS [11] that generates videos in coarse-to-fine manner using two separate transformers; one generates keyframes in coarse temporal scale, and the other interpolates missing frames between two consecutive keyframes. We trained the baselines on 128-frame videos if their training peak memory stays under the VRAM capacity, and used 16-frames otherwise. See the Appendix for more details on the baselines (Section A.3).

Results Table 2 and Figure 5 summarize the quantitative and qualitative comparisons on long-term video synthesis, respectively. The quantitative results show that our method outperforms all baselines by a large margin over all datasets. More specifically, the qualitative results display that the prediction error accumulates over time in naïve autoregressive transformers (Fig. 5(b)), which is partly ad-

³For CCVS, we report the results only in UCF-101 dataset where the official pre-trained model is available.

Table 2. Quantitative results on 128-frame video generation. The subscripts on methods denote the length of training videos.

Method	SkyTimelapse			Taichi-HD			UCF-101		
	FVD ₁₂₈ (↓)	KVD ₁₂₈ (↓)	Inf.Time	FVD ₁₂₈ (↓)	KVD ₁₂₈ (↓)	Inf.Time	FVD ₁₂₈ (↓)	KVD ₁₂₈ (↓)	Inf.Time
MoCoGAN-HD ₁₂₈	498 \pm 15	23 \pm 1.39	0.37s	991 \pm 23	290 \pm 12	0.37s	1622 \pm 53	118 \pm 6.4	0.37s
DIGAN ₁₂₈	331 \pm 13	9 \pm 0.38	0.12s	764 \pm 19	230 \pm 25	0.12s	1627 \pm 51	122 \pm 10	0.12s
CCVS ₁₆	N/A	N/A	N/A	N/A	N/A	N/A	1411 \pm 45	121 \pm 11	219s
TATS-base ₁₆	435 \pm 12	19 \pm 1.50	110s	458 \pm 21	211 \pm 18	110s	1107 \pm 31	91.8 \pm 8.5	110s
TATS-hierarchical ₁₂₈	455 \pm 12	21 \pm 0.99	127s	803 \pm 40	371 \pm 31	127s	1138 \pm 53	83.4 \pm 11.6	127s
MeBT ₁₂₈ (Ours)	239 \pm 3.6	5.1 \pm 0.16	6.53s	399 \pm 19	139 \pm 13	6.62s	968 \pm 75	75.5 \pm 12.1	7.53s




Figure 5. Qualitative results on 128-frame video generation with different models on the UCF-101 dataset: (a) MeBT (Ours), (b) TATS-base, (c) TATS-hierarchical, (d) DIGAN. We present every 10th frames from the generated videos. More results are in Appendix (Fig. 9-11).




Figure 6. Qualitative results of MeBT on 128-frame videos. We present every 16th frames from the generated videos.

dressed in the hierarchical method yet suffering from the inconsistency at keyframe boundaries and inaccurate interpolation (Fig. 5(c)). The DIGAN based on implicit representation exhibits more robustness to error propagation due to non-autoregressive generation process, yet fails to capture meaningful structures in complex videos (Fig. 5(d)). Compared to these methods, MeBT generates convincing

videos by generating not only high-fidelity frames consistently over time but also coherent semantic structures and dynamics, such as a person leaving and re-entering the scene (Fig. 5(a)), by exploiting long-term dependencies. Figure 6 illustrates more qualitative results of our method. It shows that MeBT can capture interesting long-term dynamics, such as global dynamics as the sun goes down over the water (Fig. 6(a)), complex dynamics of a human (Fig. 6(b)) and combination of human and camera motion (Fig. 6(c)).

To assess the consistency of generation quality, we measure FVD over time relative to the initial predictions and summarize the results in Figure 7. The autoregressive models (MoCoGAN-HD, CCVS, and TATS-base) suffer from dramatic error propagation. Conversely, the hierarchical model (TATS-hierarchical) is plagued by inconsistencies around keyframe boundaries that manifest as zigzag patterns. In contrast, DIGAN and MeBT achieve near-zero quality degradation in all datasets (Δ FVD \approx 0) due to their non-autoregressive design and end-to-end training with long sequences. Nevertheless, unlike MeBT, DIGAN fails to generate meaningful structures as discussed in previous paragraphs (Table 2).

5.4. Ablation Studies

This section presents component-wise analysis of MeBT. All experiments are conducted with SkyTimelapse videos. Except for the ablation studies about long-range dependency modeling, we used 16-frame model for simplicity.




Figure 7. Evaluation of generation quality over time. FVD is measured in 16-frame interval relative to the initial prediction.

Table 3. Ablation study on decoder design and bottleneck sizes.

Decoder	N_L	Memory	FVD ₁₆	KVD ₁₆
Full	128	18.2 GB	92.2 \pm 2.2	2.6 \pm .12
Ours	32	6.0 GB	110.8 \pm 2.4	4.0 \pm .15
Ours	128	7.4 GB	95.2 \pm 1.5	2.8 \pm .11
Ours	256	9.5 GB	90.4 \pm 1.8	2.8 \pm .17
Ours	512	14.7 GB	90.8 \pm 2.1	2.6 \pm .14

Table 4. Ablation study on attention and scheduling functions. The training peak memory is measured with batch size 1.

Model	Schedule	FVD ₁₂₈	KVD ₁₂₈	Memory	Time
Window	Gaussian	458 \pm 21.2	33.2 \pm 3.31	34.9 GB	13.7s
Axial	Gauss.	377 \pm 12.2	13.6 \pm 0.81	21.2 GB	12.0s
Ours	None	276 \pm 10.9	7.5 \pm .99	13.3 GB	6.53s
Ours	Uniform	253 \pm 4.6	6.7 \pm .49	13.3 GB	6.53s
Ours	Gaussian	239 \pm 3.6	5.1 \pm 0.16	13.3 GB	6.53s

Impact of latent bottleneck size To study the efficiency and performance trade-off introduced by the latent bottleneck, we train MeBT with varying number of latent codes and report the results in Table 3. Since the entire context tokens are compressed into much smaller latents by the encoder (Section 3.1), employing small latents generally sacrifices the performance. Yet, the performance quickly saturates at a reasonable scale $N_L = 256$, which is still considerably smaller than the context size $N_C = 1024$. It shows that MeBT can achieve linear complexity without compromising much performance.

Impact of the decoder architecture To evaluate the impact of our decoder (Section 3.2), we compare our method with the bidirectional decoder that leverages the full self-attention among the latent and masked tokens. Comparing the first and third rows in Table 3 shows that approximating all pair-wise interaction of context and mask tokens through linear-time MeBT decoder merely degrades the performance while considerably reducing the peak memory.

Impact of attention structure We evaluate the impact of attention structure on 128-frame videos, by replacing the latent bottleneck attention in MeBT with sparse attention methods used in video modeling: axial attention [15] alternates attention over horizontal, vertical, and temporal axes, and window attention [12] alternates attention over spatial and temporal axes. See Section A.4 for more details on the baselines. As shown in Table 4, ours shows clear improvements over the baselines in both efficiency and perfor-

mance. In terms of efficiency, MeBT outperforms baselines in both memory and inference time thanks to its linear complexity, while baselines have quadratic complexity to the temporal length from self-attention on the temporal axis. Interestingly, MeBT also exhibits improvement in sample quality since the heuristic sparsity patterns in the baselines are insufficient to capture complex long-range dependency while ours can *learn* them adaptively through the latents.

Impact of interval scheduling function To evaluate the effectiveness of the interval scheduling function in Section 3.3, we train MeBT on 128-frame videos with three different strategies: *None* refers random sampling over entire video volume ($p_n(v) = \delta(n = t)$), *Uniform* refers sampling the video interval uniformly random ($p_n(v) = \frac{1}{t}$), and the *Gaussian* refers the proposed curriculum sampling in Eq. (4). As summarized in Table 4, sampling the context and mask tokens randomly from long videos often leads to suboptimal results (*None*), while limiting the interval of the samples makes the training easier (*Uniform*). By learning from short to long videos (*Gaussian*), the proposed strategy stabilizes the training and improves the performance.

6. Conclusion

We proposed a Memory-efficient Bidirectional Transformer (MeBT), a transformer-based generative model for moderately long-term video synthesis. By formulating the video synthesis as an iterative mask prediction task and employing bidirectional transformers with efficient encoder-decoder architecture, we showed that we could push the transformers to leverage much longer videos in training while enjoying fast inference speed and robustness in error propagation. Our experiments demonstrated that by training with longer videos, MeBT could learn long-term dependencies and generate coherent videos over longer time horizon.

Acknowledgements This work was supported in part by the Institute of Information & communications Technology Planning & Evaluation (IITP) (No. and 2019-0-00075) and the National Research Foundation of Korea (NRF) (No. 2021R1C1C1012540 and 2021R1A4A3032834) funded by the Korea government (MSIT), Korea Meteorological Administration under Grant (KMA2021-00121), and Kakao Brain corporations. We thank to Jinwoo Kim for his helpful comments and discussion.

References

- [1] Adil Kaan Akan, Sadra Safadoust, Erkut Erdem, Aykut Erdem, and Fatma Güney. Stochastic video prediction with structure and motion. *CoRR*, abs/2203.10528, 2022. 5
- [2] Mohammad Babaeizadeh, Mohammad Taghi Saffar, Suraj Nair, Sergey Levine, Chelsea Finn, and Dumitru Erhan. Fitvid: Overfitting in pixel-level video prediction. *CoRR*, abs/2106.13195, 2021. 5
- [3] Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer. *CoRR*, abs/2004.05150, 2020. 5
- [4] Lluís Castrejón, Nicolas Ballas, and Aaron C. Courville. Improved conditional vrns for video prediction. In *ICCV*, 2019. 5
- [5] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T. Freeman. Maskgit: Masked generative image transformer. In *CVPR*, 2022. 2, 4
- [6] Aidan Clark, Jeff Donahue, and Karen Simonyan. Efficient video generation on complex datasets. *CoRR*, abs/1907.06571, 2019. 5
- [7] Emily Denton and Rob Fergus. Stochastic video generation with a learned prior. In *ICML*, 2018. 5
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 4
- [9] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021. 2
- [10] Jean-Yves Franceschi, Edouard Delasalles, Mickaël Chen, Sylvain Lamprier, and Patrick Gallinari. Stochastic latent residual video prediction. In *ICML*, 2020. 5
- [11] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic VQGAN and time-sensitive transformer. In *ECCV*, 2022. 1, 2, 3, 4, 5, 6, 11, 12
- [12] Agrim Gupta, Stephen Tian, Yunzhi Zhang, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Maskvit: Masked visual pre-training for video prediction. *CoRR*, abs/2206.11894, 2022. 1, 3, 4, 5, 8, 12
- [13] Ligong Han, Jian Ren, Hsin-Ying Lee, Francesco Barbieri, Kyle Olszewski, Shervin Minaee, Dimitris N. Metaxas, and Sergey Tulyakov. Show me what and tell me how: Video synthesis via multimodal conditioning. In *CVPR*, 2022. 1, 2
- [14] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 4
- [15] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *CoRR*, abs/1912.12180, 2019. 4, 5, 8, 12
- [16] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *CoRR*, abs/2205.15868, 2022. 1, 3, 4, 5
- [17] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, Olivier J. Hénaff, Matthew M. Botvinick, Andrew Zisserman, Oriol Vinyals, and João Carreira. Perceiver IO: A general architecture for structured inputs & outputs. In *ICLR*, 2022. 3, 5
- [18] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and João Carreira. Perceiver: General perception with iterative attention. In *ICML*, 2021. 3, 5
- [19] Emmanuel Kahembwe and Subramanian Ramamoorthy. Lower dimensional kernels for video discriminators. *Neural Networks*, 2020. 5
- [20] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017. 5
- [21] Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. Setvae: Learning hierarchical composition for generative modeling of set-structured data. In *CVPR*, 2021. 3, 4, 5
- [22] Yunji Kim, Seonghyeon Nam, In Cho, and Seon Joo Kim. Unsupervised keypoint learning for guiding class-conditional video prediction. In *NeurIPS*, 2019. 5
- [23] Alex X. Lee, Richard Zhang, Frederik Ebert, Pieter Abbeel, Chelsea Finn, and Sergey Levine. Stochastic adversarial video prediction. *CoRR*, abs/1804.01523, 2018. 5
- [24] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Draft-and-revise: Effective image generation with contextual rq-transformer. In *NeurIPS*, 2022. 5
- [25] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R. Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, 2019. 3, 4, 5
- [26] Sangmin Lee, Hak Gu Kim, Dae Hwi Choi, Hyung-II Kim, and Yong Man Ro. Video prediction recalling long-term motion context via memory alignment learning. In *CVPR*, 2021. 5
- [27] Wonkwang Lee, Whie Jung, Han Zhang, Ting Chen, Jing Yu Koh, Thomas E. Huang, Hyungsuk Yoon, Honglak Lee, and Seunghoon Hong. Revisiting hierarchical approach for persistent long-term video prediction. In *ICLR*, 2021. 1, 5
- [28] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P. Xing. Dual motion GAN for future-flow embedded video prediction. In *ICCV*, 2017. 5
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 11
- [30] Xuezhe Ma, Xiang Kong, Sinong Wang, Chunting Zhou, Jonathan May, Hao Ma, and Luke Zettlemoyer. Luna: Linear unified nested attention. In *NeurIPS*, 2021. 3, 4, 5
- [31] Guillaume Le Moing, Jean Ponce, and Cordelia Schmid. CCVS: context-aware controllable video synthesis. In *NeurIPS*, 2021. 5, 6, 11
- [32] Ruslan Rakhimov, Denis Volkhonskiy, Alexey Artemov, Denis Zorin, and Evgeny Burnaev. Latent video transformer. In *VISIGRAPP*, 2021. 1, 2, 5

- [33] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *ICCV*, 2017. 5
- [34] Masaki Saito, Shunta Saito, Masanori Koyama, and Sotuke Kobayashi. Train sparsely, generate densely: Memory-efficient unsupervised training of high-resolution temporal GAN. *Int. J. Comput. Vis.*, 2020. 5
- [35] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NeurIPS*, 2016. 5
- [36] Vaibhav Saxena, Jimmy Ba, and Danijar Hafner. Clockwork variational autoencoders. In *NeurIPS*, 2021. 1, 5
- [37] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *NeurIPS*, 2019. 5
- [38] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *CVPR*, 2022. 1, 5
- [39] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012. 5
- [40] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv*, 2020. 3
- [41] Yu Tian, Jian Ren, Menglei Chai, Kyle Olszewski, Xi Peng, Dimitris N. Metaxas, and Sergey Tulyakov. A good image generator is what you need for high-resolution video synthesis. In *ICLR*, 2021. 5, 6, 11, 12
- [42] Du Tran, Lubomir D. Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015. 5
- [43] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *CVPR*, 2018. 5
- [44] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *CoRR*, abs/1812.01717, 2018. 5
- [45] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. 2
- [46] Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to generate long-term future via hierarchical prediction. In Doina Precup and Yee Whye Teh, editors, *ICML*, 2017. 5
- [47] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. In *NeurIPS*, 2016. 5
- [48] Jacob Walker, Ali Razavi, and Aäron van den Oord. Predicting video with VQVAE. *CoRR*, abs/2103.01950, 2021. 2
- [49] Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv*, 2020. 3, 4
- [50] Bohan Wu, Suraj Nair, Roberto Martín-Martín, Li Fei-Fei, and Chelsea Finn. Greedy hierarchical variational autoencoders for large-scale video prediction. In *CVPR*, 2021. 5
- [51] Chenfei Wu, Jian Liang, Lei Ji, Fan Yang, Yuejian Fang, Daxin Jiang, and Nan Duan. Nüwa: Visual synthesis pre-training for neural visual world creation. In *ECCV*, 2022. 1, 2, 3, 5, 12
- [52] Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, and Jiebo Luo. Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In *CVPR*, 2018. 5, 11
- [53] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 1, 2, 5
- [54] Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, and Jinwoo Shin. Generating videos with dynamics-aware implicit generative adversarial networks. In *ICLR*, 2022. 1, 5, 6, 11, 12
- [55] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big bird: Transformers for longer sequences. In *NeurIPS*, 2020. 5
- [56] Zhu Zhang, Jianxin Ma, Chang Zhou, Rui Men, Zhikang Li, Ming Ding, Jie Tang, Jingren Zhou, and Hongxia Yang. UFC-BERT: unifying multi-modal controls for conditional image synthesis. In *NeurIPS*, 2021. 2