

# MoStGAN-V: Video Generation with Temporal Motion Styles

Xiaoqian Shen    Xiang Li    Mohamed Elhoseiny  
 King Abdullah University of Science and Technology (KAUST)  
 {xiaoqian.shen, xiang.li.1, mohamed.elhoseiny}@kaust.edu.sa

## Abstract

Video generation remains a challenging task due to spatiotemporal complexity and the requirement of synthesizing diverse motions with temporal consistency. Previous works attempt to generate videos in arbitrary lengths either in an autoregressive manner or regarding time as a continuous signal. However, they struggle to synthesize detailed and diverse motions with temporal coherence and tend to generate repetitive scenes after a few time steps. In this work, we argue that a single time-agnostic latent vector of style-based generator is insufficient to model various and temporally-consistent motions. Hence, we introduce additional time-dependent motion styles to model diverse motion patterns. In addition, a **Motion Style Attention** modulation mechanism, dubbed as **MoStAtt**, is proposed to augment frames with vivid dynamics for each specific scale (i.e., layer), which assigns attention score for each motion style w.r.t deconvolution filter weights in the target synthesis layer and softly attends different motion styles for weight modulation. Experimental results show our model achieves state-of-the-art performance on four unconditional  $256^2$  video synthesis benchmarks trained with only 3 frames per clip and produces better qualitative results with respect to dynamic motions. Code and videos have been made available at <https://github.com/xiaoqian-shen/MoStGAN-V>.

## 1. Introduction

Learning algorithms for image generation are rapidly reaching maturity and are expected to soon approach human levels of performance. However, the video generation task does not share similar success and remains a research opportunity to be further explored. Since videos are computationally intensive to model, a key question is how to generate high-quality videos with limited computation resources. Another challenge is that due to the very nature of video data, video synthesis is not simply tantamount to generating high-quality images that change over time but also requires generating these frames with temporal con-

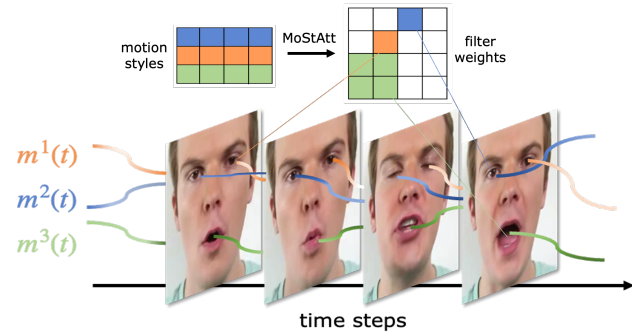


Figure 1. Intuitively show how motion styles model different motion patterns along time steps. Each curve corresponds to one motion style as a function of time step  $t$ . Each motion style will adaptively modulate filter weights of the generator with our proposed MoStAtt mechanism and further stylize input features with dynamic motions.

sistency, i.e., natural transitions between frames as well as realistic motions. This problem gets even more severe when the number of elements (e.g., objects and background) that move spatiotemporally grows.

Towards low computation consumption, several RNN or LSTM-based autoregressive approaches [5, 22, 28, 29] reduce computation complexity and make long video generation possible but the resulting videos tend to accumulate errors over time and contain inconsistent frames. Recent works use implicit neural representations by regarding time as continuous signals and mapping time to either spatial coordinates [39] or StyleGAN [12] feature inputs [26]. Although these methods leverage sparse training to enhance training efficiency and produce arbitrarily long videos, they fall short in capturing diverse patterns of changing motions. For instance, a generated video of a person talking may contain only low-frequency global motions like moving head to different angles but lack the ability to synthesize high-frequency mouth and eye actions, e.g., open/close mouth and blink eyes.

The limitation of previous works leads us to rethink how to explicitly model the diversity of motions and leverage them for motion synthesis in the video generation pro-

cess. Motion segmentation methods [3, 6, 15] assume similar trajectories representing for similar motions and segment videos into static background and diverse moving objects, which explicitly model them on perception level. Drawing inspiration from this perceptual grouping concept commonly employed in motion segmentation literature, we formulate the notion of diverse motion modeling for video generation on top the style-based framework [12, 13, 26] which leverage style parameters to control different levels of synthesis by modulating the weights of scale-specific layers in the generator. This means we aim to design group of motion styles, each of which corresponds to one motion pattern and is able to stylize input features with dynamic motions by adaptively modulating the weights of deconvolutional filters.

In this work, we aim at generating videos with vivid dynamics and rich motions over time while reserving temporal consistency. We argue that single time-agnostic latent vector of the StyleGAN-based model is insufficient to model different motion patterns in complicated motion-variant video datasets and consider auxiliary style control to increase motion awareness in the generation process. To this end, we introduce the concept of motion styles for video generation and develop a motion network to generate time-dependent motion styles to model diverse motion patterns along continuous time. In addition, we consider how to make multiple motion styles dynamically contribute to motion synthesis since motions might temporally exist in several frames while disappear in other time steps. Therefore, a motion-style attention modulation mechanism, dubbed as MoStAtt, is designed to perform cross-attention between deconvolutional filters and motions styles and linearly combine filter-specific and time-dependent motion styles for dynamic filter weight modulation. Variant motion styles that correspond to different motion patterns will adaptively attend for frames stylization and thus the modulated weights can better represent time-dependent and diverse motions. Note that our approach differs from previous works, which either predict latent motion trajectories for an image generator [28] or generate continuous motion codes and concatenate them with constant vectors to serve as initial feature inputs for a generator network [26]. In contrast, our proposed motion styles and MoStAtt modules are designed to modulate kernel weights, rather than functioning as feature inputs.

Our contributions are highlighted below:

- Introducing time-dependent motion styles for video generation task in addition to the original time-agnostic content style of style-based models to facilitate weight modulation and thus raise temporal awareness and enhance motion synthesis.
- Motion style attention modulation mechanism, dubbed

as MoStAtt, which softly attends different motion styles for weight modulation in each synthesis layer and facilitates the modulated weights to augment individual static frames with dynamic motions across holistic time continuous videos.

- Our simple yet effective approach achieves state-of-the-art performance in unconditional video generation and enhances qualitative results of motion synthesis.

## 2. Related Work

### 2.1. Video Synthesis

Different from image generation, which only models pixels at the spatial level, generating video from scratch is a more challenging task since it takes an additional temporal dimension into consideration. VGAN [33] adapts GAN to generate foreground scenes using 3D deconvolution and combines 2D background to create videos with a mask. MoCoGAN [29] and TGAN [21] model spatial and temporal dimensions with an image generator and a RNN model separately. MoCoGAN-HD [28] proposes to predict a sequence of latent motion trajectory by training a motion generator and fed it into pre-trained StyleGAN to synthesize a sequence of images. StyleGAN-V [26] [12] regards time as continuous signals while DIGAN [39] builds on top of INR-GAN [25] and treats videos as continuous signals at both spatial and temporal dimensions. LongVideoGAN [4] leverages a hierarchical generator in a multi-scale training strategy to synthesis first-person viewpoint videos. Our method builds on top of style-based model and considers time as a continuous signal similar to [26, 39] and extends style-based models with additional temporal dependent motion styles and prioritizes motion awareness with newly proposed style-based attention weight modulation technique.

In addition, several works compress high dimensional video data into a discretized latent space. For instance, VideoGPT [37] utilizes VQ-VAE [31] to encode video data to latent sequences for the prior model to predict target sequences. TATS [8] extends VQ-GAN [7] to 3D and conditions one interpolation transformer on a sparse autoregressive transformer to predict tokens. Recently, diffusion models have also been applied in the video generation task [10, 11, 38]. But they require huge computation resources, with millions of time steps to achieve high quality results and also suffer from slow inference speed. For example, a single frame of video generated with Disco Diffusion [14] takes on the order of 5 minutes and 17 seconds for animation adaptations of Stable Diffusion [19], and 1 minute for CogVideo [11]. However, our method upholds a computation effective manner that trains model with only 3 frames for each video clip and is able to generate videos at 3.12 millisecond per frame during inference time.

## 2.2. HyperNetworks

HyperNetworks refer to the models that use an auxiliary light network to generate parameters for the main networks [9, 17]. Hypernetworks have been proved useful in many fields like few-shot learning [2], continual learning [32], and language modeling [27]. In the case of generative models, [16] utilizes HyperNetworks to transform text-conditional signal to modulate weights of the generator for text-controllable image generation. In this work, we leverage a HyperNet-augmented modulation approach to facilitate motion synthesis in video generation task, which means weight modulation matrix of each deconvolutional layer are produced with a HyperNetwork. HyperNetworks suffer from extreme memory consumption due to additional parametrization. To alleviate this issue, we adopt a low-rank modulation technique to obtain the modulation weight matrix as the product of factorized modulating tensors produced by a HyperNetwork to improve training efficiency.

## 2.3. Latent Space Decomposition

Style-based models [12, 13, 26] transform latent vectors to different modulation based convolutional layers for fine-grained synthesis control. Alterations of the latent code correspond to particular manipulations in generated images. Recent works [24, 35] attempt to figure out the valid directions in the high-dimensional latent space for interpretable style control. [34] decompose motion codes into orthogonal basis deviating from source images for image animation. [18] decomposes images into appearance and pose representations and re-compose both representations to obtain latent code  $w$  for facial video generation. Our method separates a new branch to generate time-dependent motion styles in parallel with original content vector of style-based models and separately modulate the weights of deconvolutional layers with both time-agnostic content style as global context and time-sensitive motion styles to augment static frames with dynamic motions.

## 3. Method

### 3.1. Method Overview

**Generator.** Our generator derives from StyleGAN-V [26], which treats time as a continuous signal and can generate arbitrary length of frames. It models motion noise  $z_{t_0}^m, \dots, z_{t_n}^m \in \mathcal{N}(0, I)$  with Conv1d layers, where  $m$  represents for motion noise and  $t_0, \dots, t_n$  denote discrete time steps in chronological order. Then it uses acyclic positional encoding to predicts amplitudes, periods, and phases of corresponding waves for motions of different frequencies to obtain motion code  $v_t$  for current time step  $t$  (see left-top part of Figure 2). Then  $v_t$  is concatenated with trainable constant vector as generator inputs.

Same as style-based methods [12, 13, 26], our generator samples a noise vector  $z^c \in \mathcal{N}(0, I)$  from a Gaussian distribution and passes it through a Mapping Network  $F_c$  to get latent vector  $w \in \mathbb{R}^{d_c}$  in intermediate latent space  $\mathcal{W}$  [1]. For each synthesis layer,  $w$  will be transformed by an Affine Network  $A$  to generate content style  $s \in \mathbb{R}^{c_{in}}$  in  $S$  space [35]. The content style then modulates the filter weight  $\mathbf{W}_G \in \mathbb{R}^{c_{out} \times c_{in} \times k_h \times k_w}$  of corresponding convolutional layer in the generator by element-wise product, i.e.,  $\mathbf{W}'_G = \mathbf{W}_G \odot \mathbf{W}_c$ , where  $\mathbf{W}_c$  is obtained by broadcasting  $s$  to  $\mathbb{R}^{c_{in} \times k_h \times k_w}$ , and  $c_{out}, c_{in}, k_h, k_w$  represent for output channel, input channel, and kernel size of current filter respectively.

In this study, instead of using motion codes only as generator inputs, we argue that time-dependent motion styles can be generated from motion codes and further used to modulate the weights of synthesis layers for diverse motion synthesis. Figure 2 gives an overview of the generator network. We elaborate on the generation process of time-sensitive motion styles in Section 3.2 and propose a new attention-based modulation mechanism MoStAtt in Section 3.3.

**Discriminator.** For computation efficiency, we follow previous work [26] that uses a 2D discriminator network instead of 3D convolutional networks to independently extract features for each individual frame. After reaching a lower resolution block of the discriminator, it concatenates all the frame features within a video by time dimension as the global representation for the whole video. At the last block, the discriminator computes time distance information encoded with positional encodings and outputs real/fake logits as a dot product between time difference embeddings and the global feature vector.

However, we think such an operation still could not enable the discriminator to capture the motion differences in feature level, thus, we also feed frame differences as auxiliary discriminator inputs to facilitate motion artifact capture, which will be elaborated in Section. 3.5.

### 3.2. Time-sensitive Motion Styles

Although StyleGAN-V [26] represents time as a continuous signal and concatenates the time-dependent motion code  $v_t$  on top of the constant vector as generator inputs; however, it struggles to model motions of different patterns, which means the dynamic motion and the relatively static background will have a similar frequency to move, (i.e., the head might resemble a similar moving trajectory as the mouth in a talking head generation case). The main reason is that the content style mentioned above is invariant to time and lacks temporal information, thus, it is insufficient to model the whole video, which contains complex motions of variant motion patterns. To address this issue, we *augment the generation process with additional time-dependent*

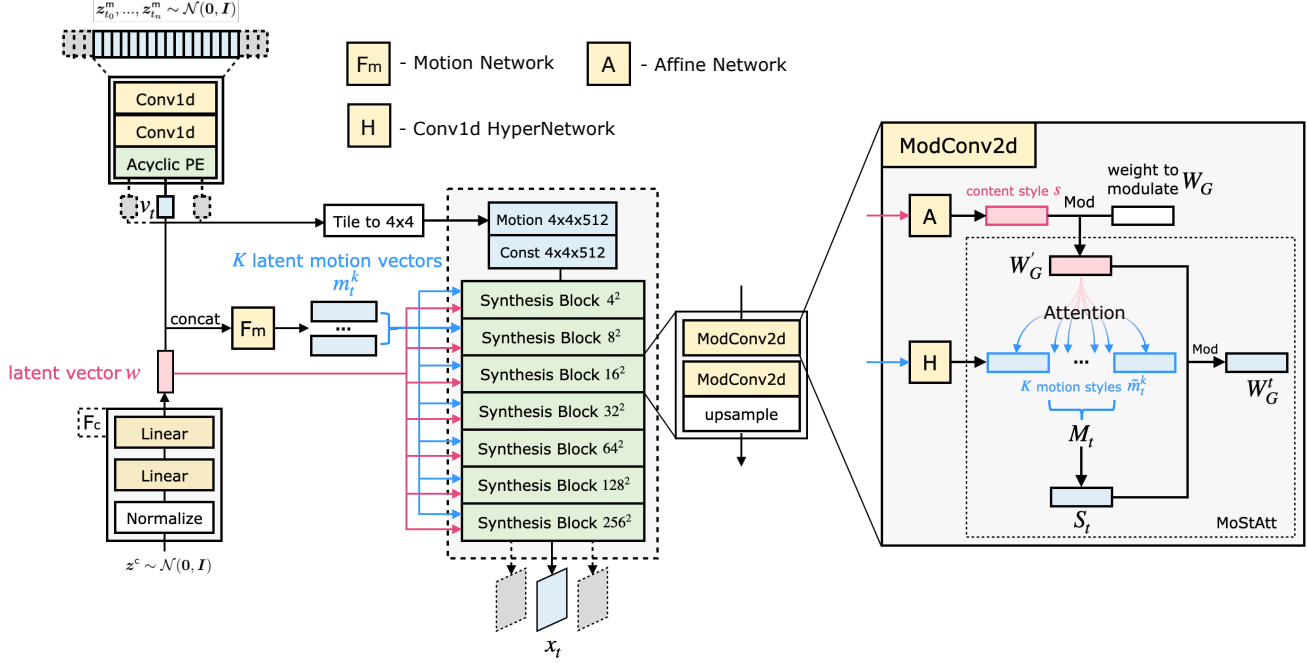


Figure 2. Overview of MoStGAN-V generator. The Motion Network  $F_m$  takes the concatenation of motion code  $v_t$  and latent vector  $w$  as input to generate time-dependent  $K$  latent motion vectors  $\{m_t^k\}_{k=1,2,\dots,K}$  for each time step  $t$ . The latent vector  $w$  and  $K$  latent motion vectors will be transformed by Affine Network  $A$  and HyperNet  $H$  into content style  $s$  and motion styles  $\{\tilde{m}_t^k\}_{k=1,2,\dots,K}$  respectively to control the weights of each Synthesis Block. Within  $\text{ModConv2d}$ , the content style first modulate the weight  $W_G$  of deconvolutional filter same as vanilla style-based model [12]. Then MoStAtt mechanism first performs cross attention between weight matrix  $W'_G$  and modulation matrix  $M_t$  (consists of  $K$  motion styles) and softly attends different motion styles as motion modulation matrix  $S_t$ . And then the weight matrix  $W'_G$  will be modulated by  $S_t$ , resulting in the final modulated weights  $W_G^t$  of deconvolution layer for time step  $t$ . This mechanism will be elaborated in Section 3.3.

*motion styles apart from the original content style by leveraging two types of styles to separately modulate weights of each synthesis layer of the generator.*

In original StyleGAN-V [26] method, each motion code  $v_t$  is generated from a sequence of trajectory noise and processed by several `conv1d` layers with a large filter (e.g., kernel size is 11), it, therefore, has already received long-term temporal information across different time steps. In this work, we want to utilize such knowledge to generate time-dependent motion styles. More specifically, the latent vector  $w$  will be concatenated with motion code  $v_t \in \mathbb{R}^{d_v}$  and then pass through a Motion Network  $F_m$  to generate  $K$  time-dependent motion vectors  $\{m_t^k\}_{k=1,2,\dots,K} \subset \mathbb{R}^{d_m}$ . A HyperNetwork is then designed to transform these motion vectors into a weight modulation matrix to modulate the weights in each synthesis layer. However, a HyperNetwork that directly produces a full rank matrix of dimensionality  $c_{in} \times c_{out} \times k_h \times k_w$  for each deconvolutional filter will consume extremely large memory. To address memory-intensive hyperscaling, we adopt a low-rank tensor decomposition technique to improve training efficiency [27]. More specifically, our HyperNetwork  $H$  trans-

forms motion vectors  $\{m_t^k\}_{k=1,2,\dots,K}$  into motion styles  $\{\tilde{m}_t^k\}_{k=1,2,\dots,K} \subset \mathbb{R}^{R \times (c_{in} + k_h + k_w)}$  for each synthesis layer to modulate the deconvolution weights, where  $R$  denotes the rank number in tensor decomposition process. The full rank motion modulation matrix for each motion style  $\tilde{m}_t^k$  can be calculated as:

$$\mathbf{M}_t^k = \sum_{r=1}^R v_1^r \otimes v_2^r \otimes v_3^r \quad (1)$$

where  $\mathbf{M}_t^k \in \mathbb{R}^{c_{in} \times k_h \times k_w}$  and  $\otimes$  denotes outer product.  $v_1^r$ ,  $v_2^r$ , and  $v_3^r$  denote modulating vectors slicing from  $\tilde{m}_t^k$ , with sizes of  $c_{in}$ ,  $k_h$ ,  $k_w$  respectively. Eventually, the whole modulation matrix for  $K$  motion styles is  $\mathbf{M}_t \in \mathbb{R}^{K \times (c_{in} \times k_h \times k_w)}$ .

### 3.3. MoStAtt: Motion Style Attention Modulation

Inspired by [16] which utilizes text-conditional signal for weight modulation and leverage attention mechanism for text-controllable image generation, we explore attention mechanisms for motion style weight modulation in the case of unconditional video generation and aim to make motion

styles softly attend modulation operation for diverse motion pattern modeling. Representing videos as a combination of individual frames, we integrate  $t$  into batch dimension and perform attention between the modulation matrix  $\mathbf{M}_t$  and deconvolutional filter weights in each synthesis layer formulated as below:

$$\mathbf{S}_t = \text{Softmax}\left(\frac{\mathbf{W}'_G(\mathbf{M}_t)^T}{\sqrt{c_{in} \times k_h \times k_w}}\right)\mathbf{M}_t \quad (2)$$

where  $\mathbf{S}_t \in \mathbb{R}^{c_{out} \times (c_{in} \times k_h \times k_w)}$  denotes the final motion modulation matrix. Intuitively, this operation determines which motions should attend in current frame, given the whole context. Therefore, the motion-attended weights will decorate the generated frames with diverse motions and progressively interact with contextualized information layer by layer via MoStAtt mechanism. The modulated weights for current frame will be finalized as  $\mathbf{W}^t_G = \mathbf{W}'_G \odot \mathbf{S}_t$ , thus the modulated weights of each ModConv2d layer will be aware of temporal information from  $\mathbf{S}_t$ . Please check the right part of Figure 2 for illustration.

### 3.4. Motion Diversity

To encourage motion styles to be disentangled from each other for modeling various motions patterns, we further leverage orthogonality constraints introduced in [23] to prevent motion modulation weights from overlapping. More specifically, we obtain the attention matrix  $\mathbf{A}_t = \frac{\mathbf{W}'_G(\mathbf{M}_t)^T}{\sqrt{c_{in} \times k_h \times k_w}}$  (from Eq. 2) before softmax layer and calculate a regularization loss for each ModConv2d as follows:

$$L_{div} = \frac{1}{T} \sum_{t=1}^T \|\mathbf{A}_t^T \mathbf{A}_t\|_F \quad (3)$$

where  $\|\cdot\|_F$  is the Frobenius norm and we average over all ModConv2d layers to form a motion regularization loss.

### 3.5. Motion Consistency

Another issue that hinders the performance of existing video generation methods is that they fail to preserve long-term consistency, e.g., motions will become receptive after several frames. To alleviate periodic artifacts, we strengthen the discriminator with frame discrepancy to capture motion changes among different frames. More specifically, we compute differences between every frame pair in a video and concatenate them with either generated or real videos as the discriminator’s input. Therefore, the input of the discriminator will be:

$$(x_1, \dots, x_N, \delta x_1, \dots, \delta x_{N-1}), \quad \delta x_i := |x_{i+1} - x_i| \quad (4)$$

where  $(x_1, \dots, x_N)$  represent for  $N$  generated frames for a video clip. Intuitively, since consecutive frames are cor-

related due to their time continuity, this operation encourages the discriminator to focus on motion artifacts among frames.

Note that StyleGAN-V [26] also attempted in a similar way that it computes differences between activations of next/previous frames in a video and concatenates this difference to the original activation maps, but it did not work because the discriminator is far more powerful and outpaces the generator much. In contrast, we observed that it works for our approach since our layer-wise MoStAtt augmented generator can produce realistic enough motions and fight against the discriminator, making it more capable of capturing motion artifacts.

## 4. Experiments

### 4.1. Experimental Setups

**Datasets.** Our MoStGAN-V model is evaluated on 4 unconditional benchmarks: FaceForensics 256<sup>2</sup> [20], Sky-Timelapse 256<sup>2</sup> [36], RainbowJelly 256<sup>2</sup> [26] and CelebV-HQ 256<sup>2</sup> [40]. We provide the details of the dataset in supplementary.

**Evaluation metrics.** Following previous works, we use Frechet Video Distance (FVD) [30] to evaluate the video generation performance. However, FVD is heavily dependent on data preprocessing procedures and sensitive to different sampling strategies (Appendix C in [26]). Following [26], we prepare 2048 frames from real video and sub-sample the generated video into 16-frame and 128-frame segments, respectively, denoted as FVD<sub>16</sub> and FVD<sub>128</sub>, to evaluate the generation quality. For all the experiments we select the results with the best FVD<sub>16</sub> performance. In addition, we also conduct a human evaluation with respect to motion diversity and time consistency.

### Comparison Approaches.

- **MoCoGAN-HD** [28] decomposes content and motion synthesis and introduces a motion generator to discover the desired trajectory in the latent space of a pre-trained image generator.

- **VideoGPT** [37] first trains a VQ-VAE [31] model for discrete latent representations of videos and then a GPT-like architecture to autoregressively model the discrete latents.

- **DIGAN** [39] utilizes an INR-based video generator to improve the motion dynamics by manipulating the space and time coordinates separately.

- **StyleGAN-V** [26] regards time as continuous signals and concatenates the continuous motion codes on top of a constant vector in StyleGAN2 [12] architecture.

**Implementation Details.** All models are trained on 32GB NVIDIA V100 GPUs. Our model uses one node of 4 GPUs and takes less than 2 days to converge to the lowest FVD<sub>16</sub>. The proposed Motion Network  $F_m$  is implemented by several MLP layers which map the concatenation of motion

Method	FaceForensics 256 <sup>2</sup>		SkyTimelapse 256 <sup>2</sup>		RainbowJelly 256 <sup>2</sup>		CelebV-HQ 256 <sup>2</sup>	
	FVD <sub>16</sub>	FVD <sub>128</sub>	FVD <sub>16</sub>	FVD <sub>128</sub>	FVD <sub>16</sub>	FVD <sub>128</sub>	FVD <sub>16</sub>	FVD <sub>128</sub>
MoCoGAN-HD [28]	111.8	653.0	164.1	878.1	579.1	628.2	212.4	753.1
VideoGPT [37]	185.9	N/A	222.7	N/A	136.0	N/A	177.8	N/A
DIGAN [39]	62.5	1824.7	83.1	196.7	436.6	369.0	72.9	163.2
StyleGAN-V [26]	47.4	89.3	79.5	197.0	195.4	262.5	68.0	158.6
MoStGAN-V (ours)	<b>39.7</b>	<b>72.6</b>	<b>65.3</b>	<b>162.4</b>	<b>70.1</b>	<b>74.3</b>	<b>56.1</b>	<b>132.1</b>

Table 1. Comparison of quantitative performance among unconditional video generation models. For each method, we report the result with the best FVD<sub>16</sub> performance.

code  $v_t$  and latent vector  $w$  into  $K$  latent motion vectors  $\{m_t^k\}_{k=1,2,\dots,K}$  with a dimension of 128, followed by a conv1d-based HyperNetwork with kernel size 1 to transform them into motion styles  $\{\tilde{m}_t^k\}_{k=1,2,\dots,K}$ . We observe that increasing the rank in Eq. (1) would not improve performance but increase parameters thus set rank  $R$  to 1. By default, we set  $K$  to 8 in our experiments.

## 4.2. Unconditional Video Generation

**Quantitative Comparison.** As shown in Table 1, our MoStGAN-V achieves new state-of-the-art performance on all datasets. Generally, FVD<sub>128</sub> score is higher than FVD<sub>16</sub> since 128-frames-long videos require more coherent features to fit the real distribution. It is interesting to find that DIGAN [39] reports a lower FVD<sub>128</sub> score than FVD<sub>16</sub> on the RainbowJelly dataset, while ours has a closer FVD<sub>16</sub> and FVD<sub>128</sub> score. One possible explanation is this dataset contains jellyfish moving back and forth, resulting in similar frames after several time steps. Note that the RainbowJelly dataset contains complex hierarchical motions, which makes this benchmark more challenging than others. Our MoStGAN-V method successfully achieves the best performance on this motion-diverse dataset with the proposed MoStAtt mechanism.

**Qualitative Results.** Figure 3 shows qualitative results of unconditional video generation among five comparing methods. The major qualitative difference in results is that our model preponderates other approaches by generating realistic detailed motions assisted with motion style modulation. For instance, in CelebV-HQ 256<sup>2</sup> [40] datasets, MoCoGAN-HD [28] merely generates consecutive images without any consistent motions. VideoGPT [37] and DIGAN [39] pay attention to motion synthesis but fail to preserve quality in longer time steps. StyleGAN-V [26] can only generate stable longer videos with moving faces but lacks of detailed motion synthesis. In contrast, our model succeeds in fine-grained and consistent facial expression and motion generation, (e.g. open/close for mouth and blinking eyes). Therefore, after introducing motion style and implementing MoStAtt mechanism on top of

StyleGAN-V [26], people can actually talk!

**Human Evaluation.** In addition, we use Mechanical Turk to assess the quality of 100 generated videos per method for each dataset on both motion diversity and temporal consistency aspect. Given a pair of videos generated by StyleGAN-V [26] and our MoStGAN-V models trained on the same dataset, people are asked to decide which video is better w.r.t motion diversity and time consistency separately. Each video is evaluated by 5 unique workers. We provide a neutral option if the volunteers find it’s hard to decide which model is better. From Figure 4, our MoStGAN-V model shows significantly better video generation quality on all evaluated datasets w.r.t both motion diversity and consistency.

## 4.3. Ablation study

**Motion styles and MoStAtt.** We first investigate the effect of motion styles and how MoStAtt mechanism contributes to the performance. We compare with the baseline method StyleGAN-V [26] which does not have time-variant motion styles, i.e.  $K = 0$ . We also conduct experiments using different numbers of motion styles, with  $K$  selecting from [1, 6, 8, 10]. Note that when  $K = 1$ , which means only one motion style will attend in the modulation matrix instead of adaptively attentive modulation what MoStAtt is designed for, hence we regard it as w/o MoStAtt.

Table 2 shows the performance of our method with only one motion style ( $K = 1$ ) w/o using MoStAtt drops compare to the baseline, but is largely boosted when multiple motion styles are used along with the MoStAtt mechanism. A small number of motion styles (e.g.  $K = 6$ ) can not fully capture various motion patterns; meanwhile, too many motion styles ( $K = 10$ ) could not improve performance but lead to ineffective computation. In addition, redundant motion styles would not contribute to motion synthesis because they will potentially receive small attention scores. Our method achieves the best performance when  $K = 8$ .

**Modulation Order.** In our MoStGAN-V, two different styles are used for weight modulation, i.e., content style and motion style. An important design choice is the order of

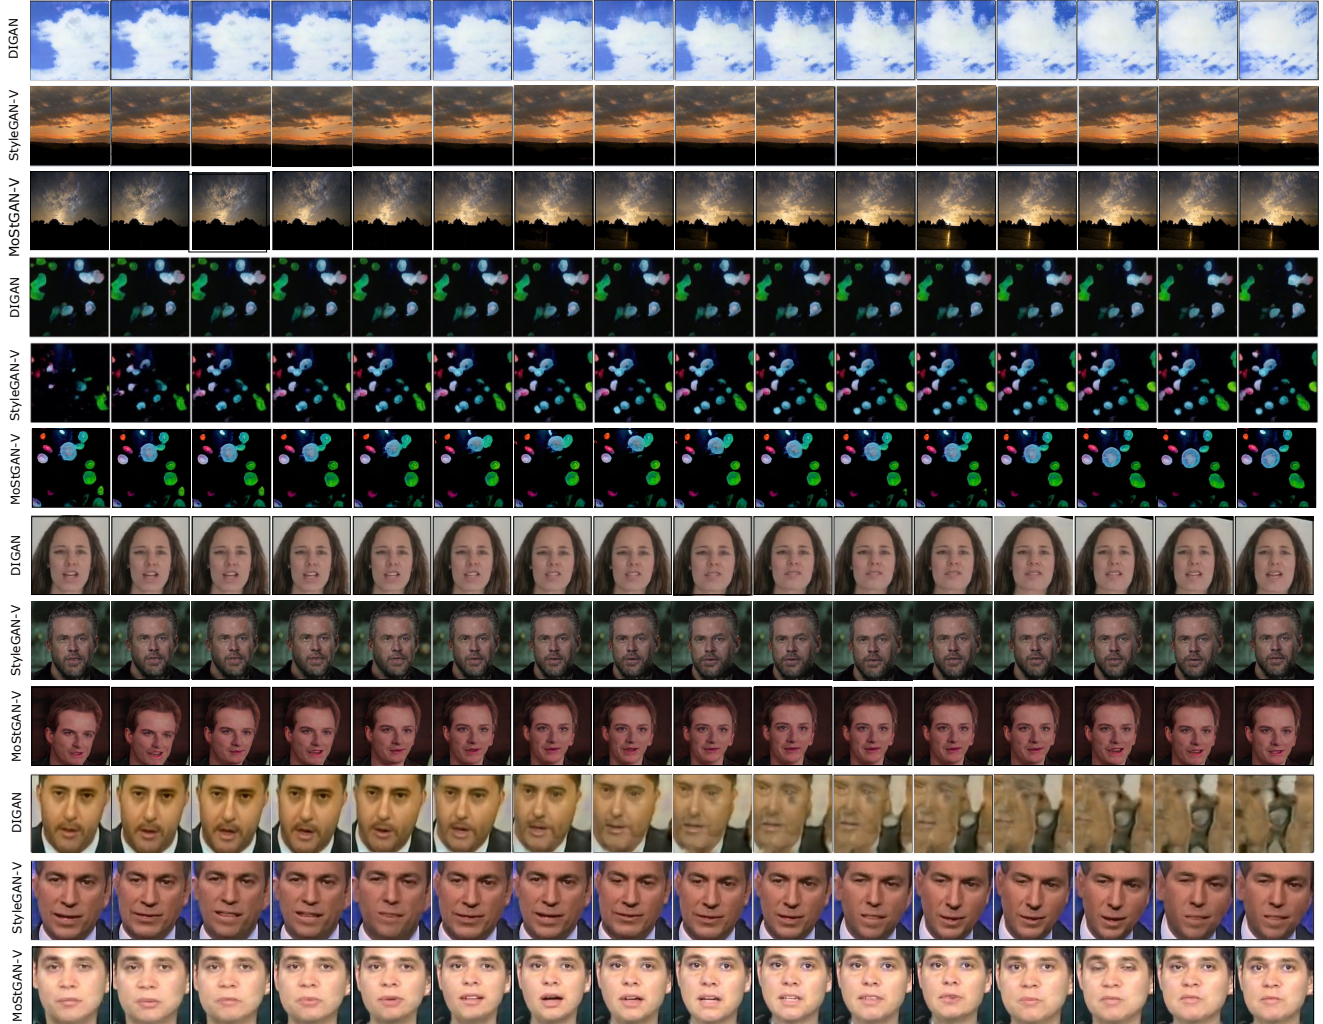


Figure 3. Random samples from the comparison baselines and our model on SkyTimelapse  $256^2$ , RainbowJelly  $256^2$ , CelebV-HQ  $256^2$  and FaceForensics  $256^2$  respectively. Start from  $t = 0$  and report every 4-th frame from a 64-frame video clip. To better compare the qualitative results, please refer to synthesised videos in our project webpage.

Number of Motion Styles	FaceForensics $256^2$		RainbowJelly $256^2$		CelebV-HQ $256^2$	
	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓
K=0 (StyleGAN-V [26])	47.4	89.3	195.4	262.5	68.0	158.6
K=1 (w/o MoStAtt)	50.9	104.6	248.7	122.1	68.4	172.0
K=6	46.3	83.9	96.1	89.9	60.8	143.5
K=8 (default)	<b>39.7</b>	<b>72.6</b>	<b>70.1</b>	<b>74.3</b>	56.1	132.1
K=10	40.4	79.8	77.6	77.9	<b>54.8</b>	<b>125.8</b>

Table 2. Effect of MoStAtt and different numbers of motion styles. Note that  $K = 0$  is StyleGAN-V [26] baseline for comparison and  $K = 1$  indicates only single motion style will attend in modulation matrix and thus regarded as w/o MoStAtt.

style-based weight modulation. We attempt two strategies:

**i)** This is the default setting of our method. We first modulate weights of each synthesis layer with the content style  $s$  and then use  $K$  motion styles  $\{\tilde{m}_t^k\}_{k=1,2,\dots,K}$  with MoStAtt to modulate the weights as Section 3.3 introduced.

**ii)** We first modulate the weights of each synthesis layer with  $K$  motion styles  $\{\tilde{m}_t^k\}_{k=1,2,\dots,K}$  with MoStAtt mechanism (use  $\mathbf{W}_G$  instead of  $\mathbf{W}'_G$  in Eq. (2)) and then, similar to original StyleGAN2 [12], use content style  $s$  for weight modulation.

Intuitively, the first strategy regards the content-style-modulated weights as the whole context of a video. Then the following MoStAtt assigns significance to motion styles that correspond to possible motions belonging to the current frame given the whole context. Table 3 shows the first strategy has significantly better performance than the one that switches the modulation order. The results verify the advantage of our design choice that first modulates the weights using a time-agnostic style at a global perspective and then time-sensitive motion styles at a local perspective.

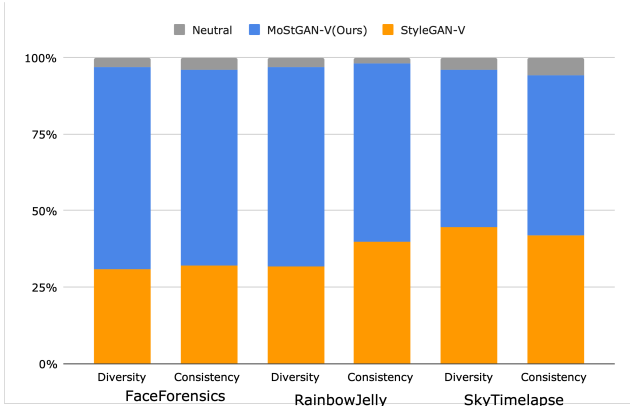


Figure 4. Human evaluation results on FaceForensics  $256^2$ , RainbowJelly  $256^2$ , and SkyTimelapse  $256^2$  datasets w.r.t motion diversity and temporal consistency.

Modulation	FaceForensics $256^2$		RainbowJelly $256^2$		CelebV-HQ $256^2$	
	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓
strategy i (default)	39.76	72.64	70.10	74.39	56.17	132.14
strategy ii	65.07	150.74	124.18	121.83	76.85	209.90

Table 3. Different weight modulation strategies.

**Motion Difference Capture.** We further investigate the effect of using motion difference as discriminator inputs to encourage motion consistency among frames. As Table 4 shows, adding motion difference as additional discriminator inputs does not improve the performance of StyleGAN-V model, while significantly boost the performance of our MoStGAN-V model. This indicates our MoStAtt-augmented generator can produce realistic enough motions and fight against the discriminator, making it more capable of capturing motion artifacts.

Methods	FaceForensics $256^2$		CelebV-HQ $256^2$	
	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓	FVD <sub>16</sub> ↓	FVD <sub>128</sub> ↓
StyleGAN-V [26]	47.4	89.3	68.0	158.6
w/ motion-diff	58.1	140.6	88.4	204.2
MoStGAN-V (ours)	49.1	117.8	61.0	138.6
w/ motion-diff (default)	<b>39.7</b>	<b>72.6</b>	<b>56.1</b>	<b>132.1</b>

Table 4. Motion difference capture. We conduct experiments w and w/o taking motion difference as discriminator inputs on StyleGAN-V [26] and our model.

#### 4.4. Motion Diversity

The motion styles are introduced to model different motion patterns in the generated video. Therefore, we quantitatively analyze a) the diversity of the motion styles; b) how the attention scores change over time, to interpret the behavior of motion styles in  $\mathcal{S}$  space. Figure 5 shows cosine similarities between  $K$  motion styles in different scales, i.e., the first ModConv2d layer of synthesis blocks at resolutions

of  $8^2$ ,  $64^2$  and  $256^2$  respectively. Figure 5 shows different motions styles have small similarity scores, which indicates that they are encouraged to be disentangled from each other and model different motion patterns. Figure 6 shows attention scores assigned by deconvolutional filters for different motion styles among time steps. From this figure, we can tell that motion styles are dynamically attending weight modulation matrix to model different motion patterns.




Figure 5. Cosine similarities between different motion styles at the first ModConv2d layer of synthesis blocks at resolutions of  $8^2$ ,  $64^2$  and  $256^2$  respectively on CelebV-HQ  $256^2$  [40].




Figure 6. Attention scores for different motion styles in  $\mathbb{R}^{cout}$  over time at the first ModConv2d layer of synthesis blocks at resolutions of  $8^2$ ,  $64^2$  and  $256^2$  respectively on CelebV-HQ  $256^2$  [40].

## 5. Conclusion

In this paper, we rethink a single time-agnostic latent vector of traditional style-based method is insufficient to model temporal motion patterns and thus introduce time-dependent motion styles to facilitate video generation with vivid dynamics. A motion style attention modulation mechanism is further designed to dynamically assign different motion styles to the current frame, which will be used to modulate the deconvolution filter weights in the target synthesis layer. Our method achieves state-of-the-art performance on four video generation benchmarks and the MoStAtt-enhanced generator can produce better qualitative results by generating diverse and temporal coherent motions.

## Acknowledgment

This work was supported by KAUST BAS/1/1685-01-01, SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence.



## References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. [3](#)
- [2] Luca Bertinetto, João F Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. *Advances in neural information processing systems*, 29, 2016. [3](#)
- [3] Pia Bideau, Rakesh R Menon, and Erik Learned-Miller. Moa-net: self-supervised motion segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. [2](#)
- [4] Tim Brooks, Janne Hellsten, Miika Aittala, Ting-Chun Wang, Timo Aila, Jaakko Lehtinen, Ming-Yu Liu, Alexei A Efros, and Tero Karras. Generating long videos of dynamic scenes. *arXiv preprint arXiv:2206.03429*, 2022. [2](#)
- [5] Aidan Clark, Jeff Donahue, and Karen Simonyan. Adversarial video generation on complex datasets. *arXiv preprint arXiv:1907.06571*, 2019. [1](#)
- [6] Ali Elqursh and Ahmed Elgammal. Online motion segmentation using dynamic label propagation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2008–2015, 2013. [2](#)
- [7] 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*, pages 12873–12883, 2021. [2](#)
- [8] 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. *arXiv preprint arXiv:2204.03638*, 2022. [2](#)
- [9] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. [3](#)
- [10] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv:2204.03458*, 2022. [2](#)
- [11] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022. [2](#)
- [12] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33:12104–12114, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [13] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. [2](#), [3](#)
- [14] Crowson Katherine. Disco diffusion. <http://discodiffusion.com>, 2021. [2](#)
- [15] Margret Keuper, Siyu Tang, Bjoern Andres, Thomas Brox, and Bernt Schiele. Motion segmentation & multiple object tracking by correlation co-clustering. *IEEE transactions on pattern analysis and machine intelligence*, 42(1):140–153, 2018. [2](#)
- [16] Jovana Lazarevic Ivan Skorokhodov Mohamed Elhoseiny KILICHBEK HAYDAROV, Aashiq Muhamed. Hypergan: Text-to-image synthesis with hypernet-modulated conditional generative adversarial networks. In *Preprint*, 2022. [3](#), [4](#)
- [17] Etai Littwin, Tomer Galanti, and Lior Wolf. On the optimization dynamics of wide hypernetworks. *arXiv preprint arXiv:2003.12193*, 2020. [3](#)
- [18] Haonan Qiu, Yuming Jiang, Hang Zhou, Wayne Wu, and Ziwei Liu. Stylefacev: Face video generation via decomposing and recomposing pretrained stylegan3. *arXiv preprint arXiv:2208.07862*, 2022. [3](#)
- [19] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [2](#)
- [20] Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics: A large-scale video dataset for forgery detection in human faces. *arXiv preprint arXiv:1803.09179*, 2018. [5](#)
- [21] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *Proceedings of the IEEE international conference on computer vision*, pages 2830–2839, 2017. [2](#)
- [22] Masaki Saito, Shunta Saito, Masanori Koyama, and Sotuke Kobayashi. Train sparsely, generate densely: Memory-efficient unsupervised training of high-resolution temporal gan. *International Journal of Computer Vision*, 128(10):2586–2606, 2020. [1](#)
- [23] Mathieu Salzmann, Carl Henrik Ek, Raquel Urtasun, and Trevor Darrell. Factorized orthogonal latent spaces. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 701–708. JMLR Workshop and Conference Proceedings, 2010. [5](#)
- [24] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1532–1540, 2021. [3](#)
- [25] Ivan Skorokhodov, Savva Ignatyev, and Mohamed Elhoseiny. Adversarial generation of continuous images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10753–10764, 2021. [2](#)
- [26] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3626–3636, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [27] Joseph Suarez. Language modeling with recurrent highway hypernetworks. *Advances in neural information processing systems*, 30, 2017. [3](#), [4](#)
- [28] 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. *arXiv preprint arXiv:2104.15069*, 2021. 1, 2, 5, 6
- [29] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 1, 2
- [30] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 5
- [31] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 5
- [32] Johannes Von Oswald, Christian Henning, João Sacramento, and Benjamin F Grewe. Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695*, 2019. 3
- [33] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. *Advances in neural information processing systems*, 29, 2016. 2
- [34] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. *arXiv preprint arXiv:2203.09043*, 2022. 3
- [35] Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12863–12872, 2021. 3
- [36] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2364–2373, 2018. 5
- [37] Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. Videogpt: Video generation using vq-vae and transformers. *arXiv preprint arXiv:2104.10157*, 2021. 2, 5, 6
- [38] Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. Diffusion probabilistic modeling for video generation. *arXiv preprint arXiv:2203.09481*, 2022. 2
- [39] 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. *arXiv preprint arXiv:2202.10571*, 2022. 1, 2, 5, 6
- [40] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. Celebv-hq: A large-scale video facial attributes dataset. *arXiv preprint arXiv:2207.12393*, 2022. 5, 6, 8