

# Photorealistic Video Generation with Diffusion Models

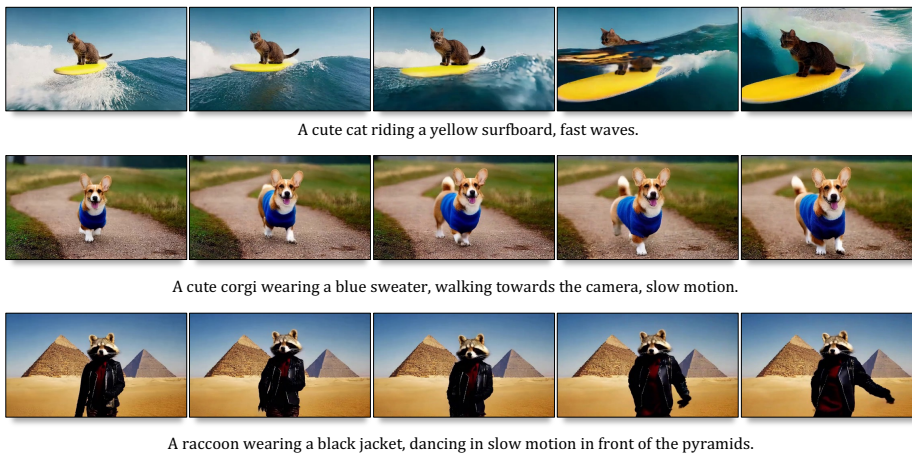
Agrim Gupta<sup>1,2,\*</sup>, Lijun Yu<sup>2</sup>, Kihyuk Sohn<sup>2</sup>, Xiuye Gu<sup>2</sup>, Meera Hahn<sup>2</sup>,  
Fei-Fei Li<sup>1</sup>, Irfan Essa<sup>2,3</sup>, Lu Jiang<sup>2</sup>, and José Lezama<sup>2</sup>

<sup>1</sup> Stanford University

<sup>2</sup> Google Research

<sup>3</sup> Georgia Institute of Technology

<https://walt-video-diffusion.github.io/>



**Fig. 1: *W.A.L.T.* samples for text-to-video generation.** *W.A.L.T.* generates high-resolution, temporally consistent photorealistic videos from text prompts. The samples shown are  $512 \times 896$  resolution over 3.6 seconds duration at 8 frames per second.

**Abstract.** We present *W.A.L.T.*, a diffusion transformer for photorealistic video generation from text prompts. Our approach has two key design decisions. First, we use a causal encoder to jointly compress images and videos within a unified latent space, enabling training and generation across modalities. Second, for memory and training efficiency, we use a window attention architecture tailored for joint spatial and spatiotemporal generative modeling. Taken together these design decisions enable us to achieve state-of-the-art performance on established video (UCF-101 and Kinetics-600) and image (ImageNet) generation benchmarks without using classifier free guidance. Finally, we also train a cascade of three models for the task of text-to-video generation consisting of a base latent video diffusion model, and two video super-resolution diffusion models to generate videos of  $512 \times 896$  resolution at 8 frames per second.

---

\*Work partially done during an internship at Google.

## 1 Introduction

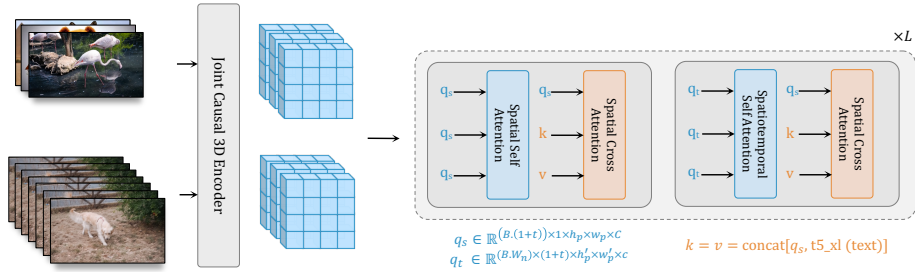
Modeling the complex and dynamic visual world with a high level of detail is the ultimate goal of visual generative models. To accomplish this goal, we need neural architectures that can efficiently process high dimensional inputs, capture long range dependencies and have powerful scaling and parallelization properties.

The transformer [74] is a neural network architecture designed with such desirable properties. This has encouraged the research community to increasingly favor it over domain-specific architectures in diverse fields such as language [28, 56, 58, 59], audio [1], speech [57], vision [19, 32], and robotics [6, 8, 91]. Such a trend towards unification allows researchers to share and build upon advances in traditionally disparate domains. Thus, leading to a virtuous cycle of innovation and improvement in model design favoring transformers.

Generative modelling of images and videos has been however a notable exception to this trend. Diffusion models [68, 70], which have emerged as a leading paradigm for generative modelling of images [17, 36] and videos [38], predominantly rely on the U-Net architecture [36, 63] as backbone, consisting of a series of convolutional [47] and self-attention [74] layers. This preference stems from the fact that the memory demands of full attention mechanisms in transformers scales quadratically with input sequence length. Such scaling leads to prohibitively high costs when processing high-dimensional signals like video with transformers.

The computational requirements of diffusion models can be reduced by operating in a lower-dimensional latent space derived from an autoencoder [21, 73, 76], giving rise to latent diffusion models (LDMs) [62]. A critical design choice in this context is the type of latent space employed: spatial compression (per frame latents) versus spatiotemporal compression. Spatial compression is often preferred because it enables leveraging pre-trained image autoencoders and LDMs, which are trained on large paired image-text datasets. However, this choice increases network complexity and limits the use of transformers as backbones, especially in generating high-resolution videos due to memory constraints. On the other hand, while spatiotemporal compression can mitigate these issues, it precludes the use of paired image-text datasets, which are much larger and diverse than their video counterparts. Hence, to leverage all sources of visual data and unlock scaling via transformers, it is critical to have a unified compressed representation of images and videos.

In this work, we present **Window Attention Latent Transformer** (*W.A.L.T*): a novel transformer-based method for latent video diffusion models (LVDMs). Our method consists of two stages. First, an autoencoder maps both videos and images into a unified, lower-dimensional latent space. This design choice enables training a single generative model *jointly* on image and video datasets and significantly reduces the computational burden for generating high resolution videos. Subsequently, we propose a new design of transformer blocks for latent video diffusion modeling composed of self-attention layers that alternate between non-overlapping, window-restricted spatial and spatiotemporal attention. This design offers two primary benefits: Firstly, the use of local window attention



**Fig. 2: W.A.L.T.** We encode images and videos into a shared latent space. The transformer backbone processes these latents with blocks having two layers of window-restricted attention: spatial layers capture spatial relations in both images and video, while spatiotemporal layers model temporal dynamics in videos and *passthrough* images via identity attention mask. Text conditioning is done via spatial cross-attention.

significantly lowers computational demands. Secondly, it facilitates joint training, where the spatial layers independently process images and video frames, while the spatiotemporal layers are dedicated to modeling the temporal relationships in videos.

While conceptually simple, our method provides the first empirical evidence of transformers’ superior generation quality and parameter efficiency in latent video diffusion on public benchmarks. Specifically, we report state-of-the-art results on class-conditional video generation (UCF-101 [71]), frame prediction (Kinetics-600 [10]) and class conditional image generation (ImageNet [16]) without using classifier free guidance. Finally, to showcase the scalability and efficiency of our method we also demonstrate results on the challenging task of photorealistic text-to-video generation. We train a cascade of three models consisting of a base latent video diffusion model, and two video super-resolution diffusion models to generate videos of  $512 \times 896$  resolution at 8 frames per second and report state-of-the-art zero-shot FVD score on the UCF-101 benchmark.

Our contributions can be summarized as follows,

- We propose W.A.L.T, a framework for transformer-based latent diffusion in a *unified* image and video latent space.
- We validate W.A.L.T by performing extensive ablations and achieving state-of-the-art video generation on frame prediction (K600), class-conditional video (UCF-101) and image (ImageNet) generation.
- We demonstrate the scaling properties of W.A.L.T and achieve state-of-the-art results in zero-shot UCF-101 text-to-video benchmark.

## 2 Related Work

*Video Diffusion Models.* Diffusion models have shown impressive results in image [36, 40, 53, 62, 68, 69] and video generation [5, 25, 31, 35, 38, 67]. Video diffusion models can be categorized into pixel-space [35, 38, 67] and latent-space [5, 25, 33, 84]

approaches, the later bringing important efficiency advantages when modeling videos. Ho *et al.* [38] demonstrated that the quality of text conditioned video generation can be significantly improved by jointly training on image and video data. Similarly, to leverage image datasets, latent video diffusion models inflate a pre-trained image model, typically a U-Net [63], into a video model by adding temporal layers, and initializing them as the identity function [5, 35, 67]. Although computationally efficient, this approach couples the design of video and image models, and precludes spatiotemporal compression. In this work, we operate on a unified latent space for images and videos, allowing us to leverage large scale image and video datasets while enjoying computational efficiency gains from spatiotemporal compression of videos.

*Transformers for Generative Modeling.* Multiple classes of generative models have utilized Transformers [74] as backbone, such as, Generative adversarial networks [43, 48, 87], autoregressive [11, 12, 21, 22, 29, 60, 75, 78, 79, 81, 82] and diffusion [2, 23, 42, 51, 54, 89] models. Inspired by the success of autoregressive pretraining of large language models [56, 58, 59], Ramesh *et al.* [60] trained a text-to-image generation model by predicting the next visual token obtained from an image tokenizer. Subsequently, this approach was applied to multiple applications including class-conditional image generation [21, 80], text-to-image [18, 60, 77, 81] or image-to-image translation [22, 78]. Similarly, for video generation, transformer-based models were proposed to predict next tokens using 3D extensions of VQGAN [24, 39, 79, 82] or using per frame image latents [29]. Autoregressive sampling of videos is typically impractical given the very long sequences involved. To alleviate this issue, non-autoregressive sampling [11, 12], *i.e.* parallel token prediction, has been adopted as a more efficient solution for transformer-based video generation [29, 75, 82]. Recently, the community has started adopting transformers as the denoising backbone for diffusion models in place of U-Net [13, 40, 51, 54, 89]. To the best of our knowledge, our work is the first successful empirical demonstration (§ 5.1) of a transformer-based backbone for jointly training image and video latent diffusion models.

### 3 Background

**Diffusion formulation.** Diffusion models [36, 68, 70] are a class of generative models which learn to generate data by iteratively denoising samples drawn from a noise distribution. Gaussian diffusion models assume a forward noising process which gradually applies noise ( $\epsilon$ ) to real data ( $\mathbf{x}_0 \sim p_{\text{data}}$ ). Concretely,

$$\mathbf{x}_t = \sqrt{\gamma(t)} \mathbf{x}_0 + \sqrt{1 - \gamma(t)} \epsilon, \quad (1)$$

where  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $t \in [0, 1]$ , and  $\gamma(t)$  is a monotonically decreasing function (noise schedule) from 1 to 0. Diffusion models are trained to learn the reverse process that inverts the forward corruptions:

$$\mathbb{E}_{\mathbf{x} \sim p_{\text{data}}, t \sim \mathcal{U}(0,1), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[ \|\mathbf{y} - f_{\theta}(\mathbf{x}_t; \mathbf{c}, t)\|^2 \right], \quad (2)$$

where  $f_\theta$  is the denoiser model parameterized by a neural network,  $\mathbf{c}$  is conditioning information e.g., class labels or text prompts, and the target  $\mathbf{y}$  can be random noise  $\epsilon$ , denoised input  $\mathbf{x}_0$  or  $\mathbf{v} = \sqrt{1 - \gamma(t)} \epsilon - \sqrt{\gamma(t)} \mathbf{x}_0$ . Following [35, 65], we use *v-prediction* in all our experiments.

**Latent diffusion models (LDMs).** Processing high-resolution images and videos using raw pixels requires considerable computational resources. To address this, LDMs operate on the low dimensional latent space of a VQ-VAE [21, 73]. VQ-VAE consists of an encoder  $E(\mathbf{x})$  that encodes an input video  $\mathbf{x} \in \mathbb{R}^{T \times H \times W \times 3}$  into a latent representation  $\mathbf{z} \in \mathbb{R}^{t \times h \times w \times c}$ . The encoder downsamples the video by a factor of  $f_s = H/h = W/w$  and  $f_t = T/t$ , where  $T = t = 1$  corresponds to using an image autoencoder. An important distinction from the original VQ-VAE is the absence of a codebook of quantized embeddings as diffusion models can operate on continuous latent spaces. A decoder  $D$  is trained to predict a reconstruction of the video,  $\hat{\mathbf{x}}$ , from  $\mathbf{z}$ . Following VQ-GAN [21], reconstruction quality can be further improved by adding adversarial [27] and perceptual losses [44, 88].

## 4 W.A.L.T

### 4.1 Learning Visual Tokens

A key design decision in video generative modeling is the choice of latent space representation. Ideally, we want a shared and unified compressed visual representation that can be used for generative modeling of both images and videos [75, 83]. The unified representation is important because joint image-video learning is preferable due to a scarcity of labeled video data [35], such as text-video pairs. Concretely, given a video sequence  $\mathbf{x} \in \mathbb{R}^{(1+T) \times H \times W \times C}$ , we aim to learn a low-dimensional representation  $\mathbf{z} \in \mathbb{R}^{(1+t) \times h \times w \times c}$  that performs spatial-temporal compression by a factor of  $f_s = H/h = W/w$  in space and a factor of  $f_t = T/t$  in time. To enable a unified representation for both videos and static images, the first frame is always encoded independently from the rest of the video. This allows static images  $\mathbf{x} \in \mathbb{R}^{1 \times H \times W \times C}$  to be treated as videos with a single frame, *i.e.*  $\mathbf{z} \in \mathbb{R}^{1 \times h \times w \times c}$ .

We instantiate this design with the causal 3D CNN encoder-decoder architecture of the MAGVIT-v2 tokenizer [83]. Typically the encoder-decoder consists of regular 3D convolution layers which cannot process the first frame independently [24, 82]. This limitation stems from the fact that a regular convolutional kernel of size  $(k_t, k_h, k_w)$  will operate on  $\lfloor \frac{k_t-1}{2} \rfloor$  frames before and  $\lfloor \frac{k_t}{2} \rfloor$  frames after the input frames. *Causal* 3D convolution layers solve this issue as the convolutional kernel operates on only the past  $k_t - 1$  frames. This ensures that the output for each frame is influenced solely by the preceding frames, enabling the model to tokenize the first frame independently.

After this stage, the input to our model is a batch of latent tensors  $\mathbf{z} \in \mathbb{R}^{(1+t) \times h \times w \times c}$  representing a single video or a stack of  $1 + t$  independent images (Fig. 2). Different from [83], our latent representation is real-valued and quantization-free. In the section below we describe how our model jointly processes a mixed batch of images and videos.

## 4.2 Learning to Generate Images and Videos

**Patchify.** Following the original ViT [19], we “patchify” each latent frame independently by converting it into a sequence of non-overlapping  $h_p \times w_p$  patches where  $h_p = h/p$ ,  $w_p = w/p$  and  $p$  is the patch size. We use learnable positional embeddings [74], which are the sum of space and time positional embeddings. Position embeddings are added to the linear projections [19] of the patches. Note that for images, we simply add the temporal position embedding corresponding to the first latent frame.

**Window attention.** Transformer models composed entirely of global self-attention modules incur significant compute and memory costs, especially for video tasks. For efficiency and for processing images and videos jointly we compute self-attention in windows [29, 74], based on two types of non-overlapping configurations: spatial (S) and spatiotemporal (ST), cf. Fig. 2. *Spatial Window (SW)* attention is restricted to all the tokens within a latent frame of size  $1 \times h_p \times w_p$  (the first dimension is time). SW models the spatial relations in images and videos. *Spatiotemporal Window (STW)* attention is restricted within a 3D window of size  $(1 + t) \times h'_p \times h'_w$ , modeling the temporal relationships among video latent frames. For images, we simply use an *identity* attention mask ensuring that the *value* embeddings corresponding to the image frame latents are passed through the layer as is. Finally, in addition to absolute position embeddings we also use relative position embeddings [50].

Our design, while conceptually straightforward, achieves computational efficiency and enables joint training on image and video datasets. In contrast to methods based on frame-level autoencoders [5, 25, 29], our approach does not suffer from flickering artifacts, which often result from encoding and decoding video frames independently. However, similar to Blattmann *et al.* [5], we can also potentially leverage pre-trained image LDMs with transformer backbones by simply interleaving STW layers.

## 4.3 Conditional Generation

To enable controllable video generation, in addition to conditioning on timestep  $t$ , diffusion models are often conditioned on additional conditional information  $\mathbf{c}$  such as class labels, natural language, past frames or low resolution videos. In our transformer backbone, we incorporate three types of conditioning mechanisms as described in what follows:

**Cross-attention.** In addition to self-attention layers in our window transformer blocks, we add a cross-attention layer for text conditioned generation. When training models on just videos, the cross-attention layer employs the same window-restricted attention as the self-attention layer, meaning S/ST blocks will have SW/STW cross-attention layers (Fig. 2). However, for joint training, we only use SW cross-attention layers. For cross-attention we concatenate the input signal (query) with the conditioning signal (key, value) as our early experiments showed this improves performance.

Method	K600 FVD↓	UCF FVD↓	params.	steps
TrIVD-GAN-FP [52]	25.7 $\pm$ 0.7	–	–	1
Video Diffusion [38]	16.2 $\pm$ 0.3	–	1.1B	256
RIN [42]	10.8	–	411M	1000
TATS [24]	–	332 $\pm$ 18	321M	1024
Phenaki [75]	36.4 $\pm$ 0.2	–	227M	48
MAGVIT [82]	9.9 $\pm$ 0.3	76 $\pm$ 2	306M	12
MAGVITv2 [83]	4.3 $\pm$ 0.1	58 $\pm$ 2	307M	24
<i>W.A.L.T-L Ours</i>	<b>3.3</b> $\pm$ 0.0	46 $\pm$ 2	313M	50
<i>W.A.L.T-XL Ours</i>	–	<b>36</b> $\pm$ 2	460M	50

**Table 1: Video generation** evaluation on frame prediction on Kinetics-600 and class-conditional generation on UCF-101.

**AdaLN-LoRA.** Adaptive normalization layers are an important component in a broad range of generative and visual synthesis models [17, 20, 45, 53–55]. A simple way to incorporate adaptive layer normalization is to include for each layer  $i$ , an MLP layer to regress a vector of conditioning parameters  $A^i = \text{MLP}(\mathbf{c} + \mathbf{t})$ , where  $A^i = \text{concat}(\gamma_1, \gamma_2, \beta_1, \beta_2, \alpha_1, \alpha_2)$ ,  $A^i \in \mathbb{R}^{6 \times d_{\text{model}}}$ , and  $\mathbf{c} \in \mathbb{R}^{d_{\text{model}}}$ ,  $\mathbf{t} \in \mathbb{R}^{d_{\text{model}}}$  are the condition and timestep embeddings. In the transformer block,  $\gamma$  and  $\beta$  scale and shift the inputs of the multi-head attention and MLP layers, respectively, while  $\alpha$  scales the output of both the multi-head attention and MLP layers. The parameter count of these additional MLP layers scales linearly with the number of layers and quadratically with the model’s dimensional size ( $\text{num\_blocks} \times d_{\text{model}} \times 6 \times d_{\text{model}}$ ). For instance, in a ViT-g model with 1B parameters, the MLP layers contribute an additional 475M parameters. Inspired by [41], we propose a simple solution dubbed *AdaLN-LoRA*, to reduce the model parameters. For each layer, we regress conditioning parameters as

$$A^1 = \text{MLP}(\mathbf{c} + \mathbf{t}), \quad A^i = A^1 + W_b^i W_a^i (\mathbf{c} + \mathbf{t}) \quad \forall i \neq 1, \quad (3)$$

where  $W_b^i \in \mathbb{R}^{d_{\text{model}} \times r}$ ,  $W_a^i \in \mathbb{R}^{r \times (6 \times d_{\text{model}})}$ . This reduces the number of trainable model parameters significantly when  $r \ll d_{\text{model}}$ . For example, a ViT-g model with  $r = 2$  reduces the MLP parameters from 475M to 12M.

**Self-conditioning.** In addition to being conditioned on external inputs, iterative generative algorithms can also be conditioned on their own previously generated samples during inference [3, 14, 66]. Specifically, Chen *et al.* [14] modify the training process for diffusion models, such that with some probability  $p_{\text{sc}}$  the model first generates a sample  $\tilde{\mathbf{z}}_0 = f_\theta(\mathbf{z}_t; \mathbf{0}, \mathbf{c}, t)$  and then refines this estimate using another forward pass conditioned on this initial sample:  $f_\theta(\mathbf{z}_t; \text{stopgrad}(\tilde{\mathbf{z}}_0), \mathbf{c}, t)$ . With probability  $1 - p_{\text{sc}}$ , only a single forward pass is done. We concatenate the model estimate with the input along the channel dimension and found this simple technique to work well when used in conjunction with *v-prediction*.

#### 4.4 Autoregressive Generation

For generating long videos via autoregressive prediction we also train our model *jointly* on the task of *frame prediction*. This is achieved by conditioning the model on past frames with a probability of  $p_{fp}$  during training. Specifically, the model is conditioned using  $c_{fp} = \text{concat}(m_{fp} \circ z_t, m_{fp})$ , where  $m_{fp}$  is a binary mask. The binary mask indicates the number of past frames used for conditioning. We condition on either 1 latent frame (image to video generation) or 2 latent frames (video prediction). This conditioning is integrated into the model through concatenation along the channel dimension of the noisy latent input. During inference, we use standard classifier-free guidance with  $c_{fp}$  as the conditioning signal.

#### 4.5 Video Super Resolution

Generating high-resolution videos with a single model is computationally prohibitive. Following [37], we use a cascaded approach with three models operating at increasing resolutions. Our base model generates videos at  $128 \times 128$  resolution which are subsequently upsampled twice via two super resolution stages. We first spatially upscale the low resolution input  $z^{lr}$  (video or image) using a depth-to-space convolution operation. Note that, unlike training where ground truth low-resolution inputs are available, inference relies on latents produced by preceding stages (*cf.* teaching-forcing). To reduce this discrepancy and improve the robustness of the super-resolution stages in handling artifacts generated by lower resolution stages, we use noise conditioning augmentation [37]. Concretely, noise is added in accordance with  $\gamma(t)$ , by sampling a noise level as  $t_{sr} \sim \mathcal{U}(0, t_{\max\_noise})$  and is provided as input to our *AdaLN-LoRA* layers.

**Aspect-ratio finetuning.** To simplify training and leverage broad data sources with different aspect ratios, we train our base stage using a square aspect ratio. We fine-tune the base stage on a subset of data to generate videos with a 9 : 16 aspect ratio by interpolating position embeddings.

## 5 Experiments

In this section, we evaluate our method on multiple tasks: class-conditional image and video generation, frame prediction and text conditioned video generation and perform extensive ablation studies of different design choices. For qualitative results, see Fig. 1, Fig. 3, Fig. 4 and videos on our [project website](#). See appendix for additional details.

### 5.1 Visual Generation

**Video generation.** We consider two standard video benchmarks, UCF-101 [71] for class-conditional generation and Kinetics-600 [10] for video prediction with 5 conditioning frames. We use FVD [72] as our primary evaluation metric.



Method	Cost (Iter×BS)	FID↓	IS↑	params.	steps
BigGAN-deep [7]	-	6.95	171.4	160M	1
LDM-4 [62]	178k×1200	10.56	103.5	400M	250
DiT-XL/2 [54]	7000k×256	9.62	121.5	675M	250
ADM [17]	-	7.49	127.5	608M	2000
MDT [23]	6500k×256	6.23	143.0	676M	250
MaskDiT [89]	1200k×1024	5.69	178.0	736M	40
RIN [42]	600k×1024	3.42	182.0	410M	1000
simple diffusion [40]	500K×2048	2.77	211.8	2B	512
VDM++ [46]	-	2.40	225.3	2B	512
<i>W.A.L.T-L Ours</i>	437k×1024	<b>2.56</b>	<b>215.1</b>	460M	50

**Table 2: Class-conditional image generation on ImageNet 256×256.** We adopt the evaluation protocol and implementation of ADM [17] and report results without classifier free guidance.

Across both datasets, *W.A.L.T* significantly outperforms all prior works (Tab. 1). Compared to prior video diffusion models, we achieve state-of-the-art performance with less model parameters, and require 50 DDIM [69] inference steps.

**Image generation.** To verify the modeling capabilities of *W.A.L.T* on the image domain, we train a version of *W.A.L.T* for the standard ImageNet class-conditional setting. For evaluation, we follow ADM [17] and report the FID [34] and Inception [64] scores calculated on 50K samples generated in 50 DDIM steps. We compare (Table 2) *W.A.L.T* with state-of-the-art image generation methods for  $256 \times 256$  resolution. Our model outperforms prior works without requiring specialized schedules, convolution inductive bias, improved diffusion losses, and classifier free guidance. Although VDM++ [46] has slightly better FID score, the model has significantly more parameters (2B).

## 5.2 Ablation Studies

We ablate *W.A.L.T* to understand the contribution of various design decisions with the default settings: model L, patch size 1,  $1 \times 16 \times 16$  spatial window,  $5 \times 8 \times 8$  spatiotemporal window,  $p_{sc} = 0.9$ ,  $c = 8$  and  $r = 2$ .

**Patch size.** In various computer vision tasks utilizing ViT [19]-based models, a smaller patch size  $p$  has been shown to consistently enhance performance [9, 19, 30, 86]. Similarly, our findings also indicate that a reduced patch size improves performance (Table 3a).

**Window attention.** We compare three different STW window configurations with full self-attention (Table 3b). We find that local self-attention can achieve competitive (or better) performance while being significantly faster (up to  $2\times$ ) and requiring less accelerator memory.

**Self-conditioning.** In Table 3c we study the influence of varying the self-conditioning rate  $p_{sc}$  on generation quality. We notice a clear trend: increasing the self conditioning rate from 0.0 (no self-conditioning) to 0.9 improves the FVD score substantially (44%).

patch size $p$	FVD↓	IS↑	st window	FVD↓	IS↑	sps
1	60.7	87.2	$5 \times 4 \times 4$	56.9	87.3	2.24
2	134.4	82.2	$5 \times 8 \times 8$	59.6	87.4	2.00
4	461.8	63.9	$5 \times 16 \times 16$	55.3	87.4	1.75
			full self attn.	59.9	87.8	1.20

(a) **Patch size.** Lower patch size is significantly better.(b) **Spatiotemporal window size.** Full self-attention is not essential for good performance. sps is steps per sec.

$p_{sc}$	FVD↓	IS↑	$r$	FVD↓	IS↑	params
0.0	109.9	82.6	2	60.7	87.2	313 M
0.3	76.0	86.5	4	56.6	87.3	314 M
0.6	60.0	86.8	16	55.5	88.0	316 M
0.9	61.4	87.1	64	54.4	87.9	324 M
			256	52.5	88.5	357 M

(c) **Self-conditioning.** Higher  $p_{sc}$  is better. (d) **AdaLN-LoRA.** Bigger  $r$  is better.

	FVD↓	IS↑	$c$	rFVD↓	FVD↓	IS↑
w/o qk norm [15]	59.0	86.8	4	37.7	86.4	84.9
w/o latent norm	67.9	87.1	8	17.1	75.4	86.3
w/o zero snr [49]	91.0	84.2	16	8.2	67.0	86.0
full method	60.7	87.2	32	3.5	83.4	82.9

(e) **Other improvements.** See text for details.(f) **Latent dimension  $c$ .** Higher  $c$  is better.Table 3: **Ablation experiments** on UCF-101 [71]. We compare FVD and inception scores to ablate important design decisions with the default setting: L model,  $1 \times 16 \times 16$  spatial window,  $5 \times 8 \times 8$  spatiotemporal (st) window,  $p_{sc} = 0.9$ ,  $c = 8$  and  $r = 2$ .

**AdaLN-LoRA.** An important design decision in diffusion models is the conditioning mechanism. We investigate the effect of increasing the bottleneck dimension  $r$  in our proposed *AdaLN-LoRA* layers (Table 3d). This hyperparameter provides a flexible way to trade off between number of model parameters and generation performance. As shown in Table 3d, increasing  $r$  improves performance but also increases model parameters. This highlights an important model design question: given a fixed parameter budget, how should we allocate parameters - either by using *separate* AdaLN layers, or by increasing base model parameters while using *shared* AdaLN-LoRA layers? We explore this in Table 4 by comparing two model configurations: *W.A.L.T-L* with separate AdaLN layers and *W.A.L.T-XL* with AdaLN-LoRA and  $r = 2$ . While both configurations yield similar FVD and Inception scores, *W.A.L.T-XL* achieves a lower final loss value, suggesting the advantage of allocating more parameters to the base model and choosing an appropriate  $r$  value within accelerator memory limits.

Model	AdaLN	FVD↓	IS↑	params.	final loss
L	separate	34.6	90.2	458M	0.274
XL	LoRA-2	36.7	89.4	460M	0.268

**Table 4: Parameter matched** comparison between AdaLN-LoRA and per layer adaln layers. See text for details.

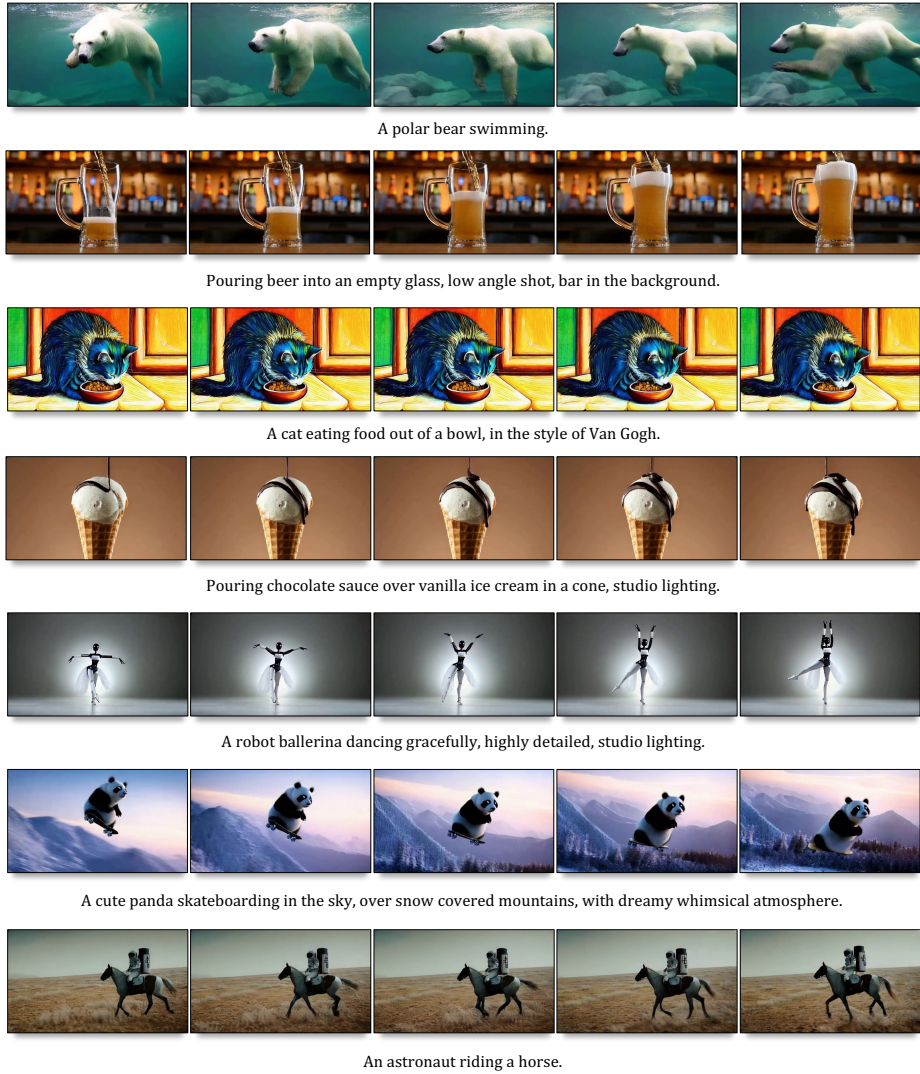
**Noise schedule.** Common latent diffusion noise schedules [62] typically do not ensure a zero signal-to-noise ratio (SNR) at the final timestep, i.e., at  $t = 1, \gamma(t) > 0$ . This leads to a mismatch between training and inference phases. During inference, models are expected to start from purely Gaussian noise, whereas during training, at  $t = 1$ , a small amount of signal information remains accessible to the model. This is especially harmful for video generation as videos have high temporal redundancy. Even minimal information leakage at  $t = 1$  can reveal substantial information to the model. Addressing this mismatch by enforcing a zero terminal SNR [49] significantly improves performance (Table 3e). Note that this approach was originally proposed to fix over-exposure problems in image generation, but we find it effective for video generation as well.

**Autoencoder.** Finally, we investigate one critical but often overlooked hyperparameter in the first stage of our model: the channel dimension  $c$  of the autoencoder latent  $z$ . As shown in Table 3f, increasing  $c$  significantly improves the reconstruction quality (lower rFVD) while keeping the same spatial  $f_s$  and temporal compression  $f_t$  ratios. Empirically, we found that both lower and higher values of  $c$  lead to poor FVD scores in generation, with a sweet spot of  $c = 8$  working well across most datasets and tasks we evaluated. We also normalize the latents before processing them via transformer which further improves performance.

In our transformer models, we use query-key normalization [15] as it helps stabilize training for larger models. Finally, we note that some of our default settings are not optimal, as indicated by ablation studies. These defaults were chosen early on for their robustness across datasets, though further tuning may improve performance.

### 5.3 Text-to-video

We train *W.A.L.T* for text-to-video jointly on text-image and text-video pairs (Sec. 4.2). We used a dataset of  $\sim 970$ M text-image pairs and  $\sim 89$ M text-video pairs from the public internet and internal sources. We train our base model at resolution  $17 \times 128 \times 128$  (3B parameters), and two  $2\times$  cascaded super-resolution models for  $17 \times 128 \times 224 \rightarrow 17 \times 256 \times 448$  (L, 1.3B,  $p = 2$ ) and  $17 \times 256 \times 448 \rightarrow 17 \times 512 \times 896$  (L, 419M,  $p = 2$ ) respectively. We fine-tune the base stage for the  $9 : 16$  aspect ratio to generate videos at resolution  $128 \times 224$ . We use classifier free guidance for all our text-to-video results.



**Fig. 3: Qualitative evaluation.** Example videos generated by *W.A.L.T* from natural language prompts at  $512 \times 896$  resolution over 3.6 seconds duration at 8 frames per second. The *W.A.L.T* model is able to generate temporally consistent photorealistic videos that align with the textual prompt. There are more results on our [project website](#).

Method	IS ( $\uparrow$ )	FVD ( $\downarrow$ )
CogVideo (Chinese) [39]	23.6	751.3
CogVideo (English) [39]	25.3	701.6
MagicVideo [90]	-	699.0
EMU Video* [26]	42.7	606.2
Video LDM [5]	33.5	550.6
Make-A-Video [67]	33.0	367.2
PYoCo [25]	<b>47.8</b>	355.2
PixelDance* [85]	42.1	242.8
SVD* [4]	-	<b>242.0</b>
<i>W.A.L.T (Ours)</i> 419M (video only)	26.8	598.8
<i>W.A.L.T (Ours)</i> 419M (video + image)	31.7	344.5
<i>W.A.L.T (Ours)</i> 3B (video + image)	35.1	258.1

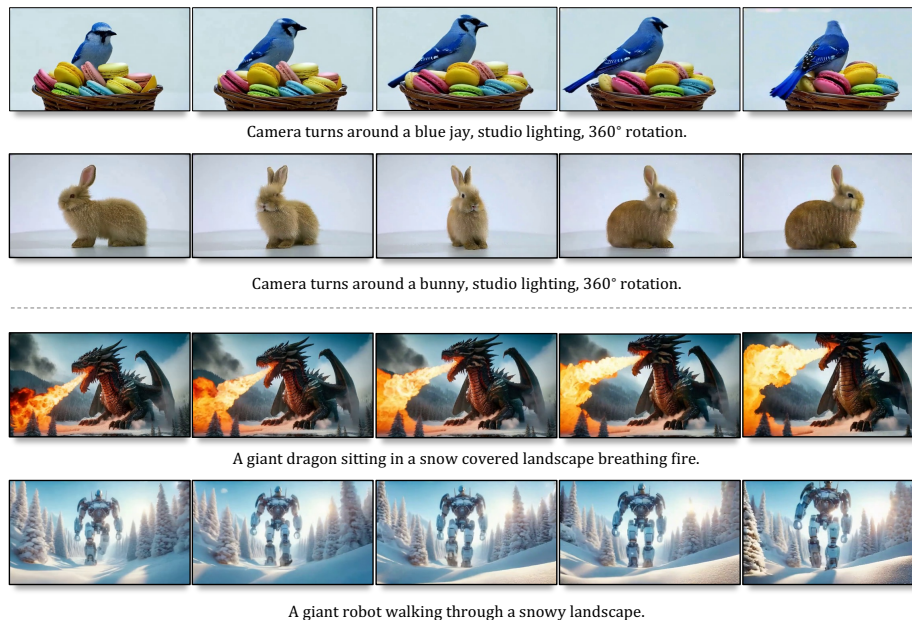
**Table 5: UCF-101 text-to-video generation.** Joint training on image and video datasets in conjunction with scaling the model parameters is essential for high quality video generation. \* Denotes concurrent work.

**Quantitative Evaluation** Evaluating text-conditioned video generation scientifically remains a significant challenge, in part due to the absence of standardized training datasets and benchmarks. So far we have focused our experiments and analyses on the standard academic benchmarks, which use the same training data to ensure controlled and fair comparisons. Nevertheless, to compare with prior work on text-to-video, we also report results on the UCF-101 dataset in the zero-shot evaluation protocol in Table 5 [25, 39, 67]. Also see supplement.

**Joint training.** A primary strength of our framework is its ability to train simultaneously on both image and video datasets. In Table 5 we ablate the impact of this joint training. Specifically, we trained two versions of *W.A.L.T*-L (each with 419M params.) models using the default settings specified in § 5.2. We find that joint training leads to a notable improvement across both metrics. Our results align with the findings of Ho *et al.* [38], who demonstrated the benefits of joint training for pixel-based video diffusion models with U-Net backbones.

**Scaling.** Transformers are known for their ability to scale effectively in many tasks [6, 15, 58]. In Table 5 we show the benefits of scaling our transformer model for video diffusion. Scaling our base model size leads to significant improvements on both the metrics. It is important to note, however, that our base model is considerably smaller than leading text-to-video systems. For instance, Ho *et al.* [35] trained base model of 5.7B parameters. Hence, we believe scaling our models further is an important direction of future work.

**Comparison with prior work.** In Table 5, we present a system-level comparison of various text-to-video methods. Our results are promising, surpassing all previous work in the FVD metric, and being competitive with concurrent work. In terms of the IS, *W.A.L.T* outperforms all previous work but PYoCo [25]. A possible explanation for this discrepancy might be PYoCo’s use of stronger text embeddings. Specifically, they utilize both CLIP [56] and T5-XXL [61] encoders, whereas we employ a T5-XL [61] text encoder only.



**Fig. 4: Samples** at  $512 \times 896$  resolution. *Top*: Text-to-video samples of consistent 3D camera motion (5.1 secs). *Bottom*: Image-to-video samples (3.8 secs). Images are generated using off-shelf models. Visit our [project website](#) to view them in video format.

**Qualitative Results** As mentioned in § 4.4, we jointly train our model on the task of frame prediction conditioned on 1 or 2 latent frames. Hence, our model can be used for animating images (**image-to-video**) and generating longer videos with consistent camera motion (Fig. 4). See videos on our [project website](#).

## 6 Conclusion

In this work, we introduce *W.A.L.T.*, a simple, scalable, and efficient transformer-based framework for latent video diffusion models. A key design decision in our framework is to operate on a shared, compressed visual representation that enables generative modeling of both images and videos. We demonstrate state-of-the-art results for image and video generation using a transformer backbone with windowed attention. We also train a cascade of three *W.A.L.T.* models jointly on image and video datasets, to synthesize high-resolution, temporally consistent photorealistic videos from natural language descriptions. While generative modeling has seen tremendous recent advances for images, progress on video generation has lagged behind. We hope that scaling our framework will enable the creation of powerful video generative models which can understand and model the physical world.

## Acknowledgements

We thank Bryan Seybold, Dan Kondratyuk, David Ross, Hartwig Adam, Huisheng Wang, Jason Baldridge, Mauricio Delbracio and Orly Liba for helpful discussions and feedback.

## References

1. Agostinelli, A., Denk, T.I., Borsos, Z., Engel, J., Verzetti, M., Caillon, A., Huang, Q., Jansen, A., Roberts, A., Tagliasacchi, M., et al.: MusicLM: Generating music from text. arXiv:2301.11325 (2023) [2](#)
2. Bao, F., Li, C., Cao, Y., Zhu, J.: All are worth words: a vit backbone for score-based diffusion models. In: NeurIPS 2022 Workshop on Score-Based Methods (2022) [4](#)
3. Bengio, S., Vinyals, O., Jaitly, N., Shazeer, N.: Scheduled sampling for sequence prediction with recurrent neural networks. *Advances in neural information processing systems* **28** (2015) [7](#)
4. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023) [13](#)
5. Blattmann, A., Rombach, R., Ling, H., Dockhorn, T., Kim, S.W., Fidler, S., Kreis, K.: Align your latents: High-resolution video synthesis with latent diffusion models. In: CVPR (2023) [3](#), [4](#), [6](#), [13](#)
6. Bousmalis, K., Vezzani, G., Rao, D., Devin, C., Lee, A.X., Bauza, M., Davchev, T., Zhou, Y., Gupta, A., Raju, A., et al.: Robocat: A self-improving foundation agent for robotic manipulation. arXiv preprint arXiv:2306.11706 (2023) [2](#), [13](#)
7. Brock, A., Donahue, J., Simonyan, K.: Large scale GAN training for high fidelity natural image synthesis. In: ICLR (2018) [9](#)
8. Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al.: Rt-1: Robotics transformer for real-world control at scale. arXiv preprint arXiv:2212.06817 (2022) [2](#)
9. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021) [9](#)
10. Carreira, J., Noland, E., Banki-Horvath, A., Hillier, C., Zisserman, A.: A short note about Kinetics-600. arXiv:1808.01340 (2018) [3](#), [8](#)
11. Chang, H., Zhang, H., Barber, J., Maschinot, A., Lezama, J., Jiang, L., Yang, M.H., Murphy, K., Freeman, W.T., Rubinstein, M., et al.: Muse: Text-to-image generation via masked generative transformers. In: ICML (2023) [4](#)
12. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: MaskGIT: Masked generative image transformer. In: CVPR (2022) [4](#)
13. Chen, J., Yu, J., Ge, C., Yao, L., Xie, E., Wu, Y., Wang, Z., Kwok, J., Luo, P., Lu, H., et al.: Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. arXiv preprint arXiv:2310.00426 (2023) [4](#)
14. Chen, T., Zhang, R., Hinton, G.: Analog bits: Generating discrete data using diffusion models with self-conditioning. arXiv preprint arXiv:2208.04202 (2022) [7](#)
15. Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A.P., Caron, M., Geirhos, R., Alabdulmohsin, I., et al.: Scaling vision transformers to 22 billion parameters. In: International Conference on Machine Learning. pp. 7480–7512. PMLR (2023) [10](#), [11](#), [13](#)

16. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: CVPR (2009) [3](#)
17. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. In: NeurIPS (2021) [2](#), [7](#), [9](#)
18. Ding, M., Yang, Z., Hong, W., Zheng, W., Zhou, C., Yin, D., Lin, J., Zou, X., Shao, Z., Yang, H., et al.: CogView: Mastering text-to-image generation via transformers. In: NeurIPS (2021) [4](#)
19. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2020) [2](#), [6](#), [9](#)
20. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. arXiv preprint arXiv:1610.07629 (2016) [7](#)
21. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: CVPR (2021) [2](#), [4](#), [5](#)
22. Gafni, O., Polyak, A., Ashual, O., Sheynin, S., Parikh, D., Taigman, Y.: Make-a-scene: Scene-based text-to-image generation with human priors. In: European Conference on Computer Vision. pp. 89–106. Springer (2022) [4](#)
23. Gao, S., Zhou, P., Cheng, M.M., Yan, S.: Masked diffusion transformer is a strong image synthesizer. arXiv:2303.14389 (2023) [4](#), [9](#)
24. Ge, S., Hayes, T., Yang, H., Yin, X., Pang, G., Jacobs, D., Huang, J.B., Parikh, D.: Long video generation with time-agnostic VQGAN and time-sensitive transformer. In: ECCV (2022) [4](#), [5](#), [7](#)
25. Ge, S., Nah, S., Liu, G., Poon, T., Tao, A., Catanzaro, B., Jacobs, D., Huang, J.B., Liu, M.Y., Balaji, Y.: Preserve your own correlation: A noise prior for video diffusion models. arXiv preprint arXiv:2305.10474 (2023) [3](#), [6](#), [13](#)
26. Girdhar, R., Singh, M., Brown, A., Duval, Q., Azadi, S., Rambhatla, S.S., Shah, A., Yin, X., Parikh, D., Misra, I.: Emu video: Factorizing text-to-video generation by explicit image conditioning. arXiv preprint arXiv:2311.10709 (2023) [13](#)
27. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014) [5](#)
28. Google: PaLM 2 technical report. arXiv:2305.10403 (2023) [2](#)
29. Gupta, A., Tian, S., Zhang, Y., Wu, J., Martín-Martín, R., Fei-Fei, L.: MaskViT: Masked visual pre-training for video prediction. In: ICLR (2022) [4](#), [6](#)
30. Gupta, A., Wu, J., Deng, J., Fei-Fei, L.: Siamese masked autoencoders. arXiv preprint arXiv:2305.14344 (2023) [9](#)
31. Harvey, W., Naderiparizi, S., Masrani, V., Weilbach, C., Wood, F.: Flexible diffusion modeling of long videos. *Advances in Neural Information Processing Systems* **35**, 27953–27965 (2022) [3](#)
32. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: CVPR (2022) [2](#)
33. He, Y., Yang, T., Zhang, Y., Shan, Y., Chen, Q.: Latent video diffusion models for high-fidelity long video generation. arXiv preprint arXiv:2211.13221 (2023) [3](#)
34. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017) [9](#)
35. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv:2210.02303 (2022) [3](#), [4](#), [5](#), [13](#)
36. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020) [2](#), [3](#), [4](#)



37. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. *JMLR* **23**(1), 2249–2281 (2022) [8](#)
38. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. In: *ICLR Workshops (2022)* [2](#), [3](#), [4](#), [7](#), [13](#)
39. Hong, W., Ding, M., Zheng, W., Liu, X., Tang, J.: CogVideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv:2205.15868 (2022)* [4](#), [13](#)
40. Hoogeboom, E., Heek, J., Salimans, T.: simple diffusion: End-to-end diffusion for high resolution images. In: *ICML (2023)* [3](#), [4](#), [9](#)
41. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: *ICLR (2021)* [7](#)
42. Jabri, A., Fleet, D.J., Chen, T.: Scalable adaptive computation for iterative generation. In: *ICML (2023)* [4](#), [7](#), [9](#)
43. Jiang, Y., Chang, S., Wang, Z.: Transgan: Two pure transformers can make one strong gan, and that can scale up. *Advances in Neural Information Processing Systems* **34**, 14745–14758 (2021) [4](#)
44. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *ECCV (2016)* [5](#)
45. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *CVPR (2019)* [7](#)
46. Kingma, D.P., Gao, R.: Understanding the diffusion objective as a weighted integral of elbos. *arXiv:2303.00848 (2023)* [9](#)
47. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998) [2](#)
48. Lee, K., Chang, H., Jiang, L., Zhang, H., Tu, Z., Liu, C.: Vitgan: Training gans with vision transformers. *arXiv preprint arXiv:2107.04589 (2021)* [4](#)
49. Lin, S., Liu, B., Li, J., Yang, X.: Common diffusion noise schedules and sample steps are flawed. *arXiv preprint arXiv:2305.08891 (2023)* [10](#), [11](#)
50. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 10012–10022 (2021) [6](#)
51. Lu, H., Yang, G., Fei, N., Huo, Y., Lu, Z., Luo, P., Ding, M.: Vdt: General-purpose video diffusion transformers via mask modeling. *arXiv preprint arXiv:2305.13311 (2023)* [4](#)
52. Luc, P., Clark, A., Dieleman, S., Casas, D.d.L., Doron, Y., Cassirer, A., Simonyan, K.: Transformation-based adversarial video prediction on large-scale data. *arXiv:2003.04035 (2020)* [7](#)
53. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021) [3](#), [7](#)
54. Peebles, W., Xie, S.: Scalable diffusion models with transformers. *arXiv:2212.09748 (2022)* [4](#), [7](#), [9](#)
55. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 32 (2018) [7](#)
56. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *ICML (2021)* [2](#), [4](#), [13](#)
57. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. In: *International Conference on Machine Learning*. pp. 28492–28518. PMLR (2023) [2](#)

58. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018) [2](#), [4](#), [13](#)
59. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019) [2](#), [4](#)
60. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML (2021) [4](#)
61. Roberts, A., Raffel, C., Lee, K., Matena, M., Shazeer, N., Liu, P.J., Narang, S., Li, W., Zhou, Y.: Exploring the limits of transfer learning with a unified text-to-text transformer (2019) [13](#)
62. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) [2](#), [3](#), [9](#), [11](#)
63. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015) [2](#), [4](#)
64. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training GANs. In: NeurIPS (2016) [9](#)
65. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. arXiv preprint arXiv:2202.00512 (2022) [5](#)
66. Savinov, N., Chung, J., Binkowski, M., Elsen, E., Oord, A.v.d.: Step-unrolled denoising autoencoders for text generation. arXiv preprint arXiv:2112.06749 (2021) [7](#)
67. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., et al.: Make-a-video: Text-to-video generation without text-video data. arXiv:2209.14792 (2022) [3](#), [4](#), [13](#)
68. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: ICML (2015) [2](#), [3](#), [4](#)
69. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) [3](#), [9](#)
70. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: NeurIPS (2019) [2](#), [4](#)
71. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv:1212.0402 (2012) [3](#), [8](#), [10](#)
72. Unterthiner, T., van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. arXiv:1812.01717 (2018) [8](#)
73. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. In: NeurIPS (2017) [2](#), [5](#)
74. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017) [2](#), [4](#), [6](#)
75. Villegas, R., Babaeizadeh, M., Kindermans, P.J., Moraldo, H., Zhang, H., Saffar, M.T., Castro, S., Kunze, J., Erhan, D.: Phenaki: Variable length video generation from open domain textual description. arXiv:2210.02399 (2022) [4](#), [5](#), [7](#)
76. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. pp. 1096–1103 (2008) [2](#)
77. Wu, C., Huang, L., Zhang, Q., Li, B., Ji, L., Yang, F., Sapiro, G., Duan, N.: Godiva: Generating open-domain videos from natural descriptions. arXiv preprint arXiv:2104.14806 (2021) [4](#)
78. Wu, C., Liang, J., Ji, L., Yang, F., Fang, Y., Jiang, D., Duan, N.: Nüwa: Visual synthesis pre-training for neural visual world creation. In: European conference on computer vision. pp. 720–736. Springer (2022) [4](#)

79. Yan, W., Zhang, Y., Abbeel, P., Srinivas, A.: VideoGPT: Video generation using VQ-VAE and transformers. arXiv:2104.10157 (2021) [4](#)
80. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldrige, J., Wu, Y.: Vector-quantized image modeling with improved VQGAN. In: ICLR (2022) [4](#)
81. Yu, J., Xu, Y., Koh, J.Y., Luong, T., Baid, G., Wang, Z., Vasudevan, V., Ku, A., Yang, Y., Ayan, B.K., et al.: Scaling autoregressive models for content-rich text-to-image generation. arXiv:2206.10789 (2022) [4](#)
82. Yu, L., Cheng, Y., Sohn, K., Lezama, J., Zhang, H., Chang, H., Hauptmann, A.G., Yang, M.H., Hao, Y., Essa, I., et al.: MAGVIT: Masked generative video transformer. In: CVPR (2023) [4](#), [5](#), [7](#)
83. Yu, L., Lezama, J., Gundavarapu, N.B., Versari, L., Sohn, K., Minnen, D., Cheng, Y., Gupta, A., Gu, X., Hauptmann, A.G., et al.: Language model beats diffusion-tokenizer is key to visual generation. arXiv preprint arXiv:2310.05737 (2023) [5](#), [7](#)
84. Yu, S., Sohn, K., Kim, S., Shin, J.: Video probabilistic diffusion models in projected latent space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18456–18466 (2023) [3](#)
85. Zeng, Y., Wei, G., Zheng, J., Zou, J., Wei, Y., Zhang, Y., Li, H.: Make pixels dance: High-dynamic video generation. arXiv preprint arXiv:2311.10982 (2023) [13](#)
86. Zhai, X., Kolesnikov, A., Houlsby, N., Beyer, L.: Scaling vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12104–12113 (2022) [9](#)
87. Zhang, B., Gu, S., Zhang, B., Bao, J., Chen, D., Wen, F., Wang, Y., Guo, B.: Styleswin: Transformer-based gan for high-resolution image generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11304–11314 (2022) [4](#)
88. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) [5](#)
89. Zheng, H., Nie, W., Vahdat, A., Anandkumar, A.: Fast training of diffusion models with masked transformers. arXiv:2306.09305 (2023) [4](#), [9](#)
90. Zhou, D., Wang, W., Yan, H., Lv, W., Zhu, Y., Feng, J.: Magicvideo: Efficient video generation with latent diffusion models. arXiv preprint arXiv:2211.11018 (2022) [13](#)
91. Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F., Wu, J., Wohlhart, P., Welker, S., Wahid, A., et al.: RT-2: Vision-language-action models transfer web knowledge to robotic control. In: CoRL (2023) [2](#)