

MaskINT: Video Editing via Interpolative Non-autoregressive Masked Transformers

Haoyu Ma^{1,2*}, Shahin Mahdizadehghadam², Bichen Wu², Zhipeng Fan², Yuchao Gu³,
Wenliang Zhao², Lior Shapira², Xiaohui Xie¹

¹University of California, Irvine, ²GenAI, Meta, ³National University of Singapore

Project Page: <https://maskint.github.io>

Abstract

Recent advances in generative AI have significantly enhanced image and video editing, particularly in the context of text prompt control. State-of-the-art approaches predominantly rely on diffusion models to accomplish these tasks. However, the computational demands of diffusion-based methods are substantial, often necessitating large-scale paired datasets for training, and therefore challenging the deployment in real applications. To address these issues, this paper breaks down the text-based video editing task into two stages. First, we leverage a pre-trained text-to-image diffusion model to simultaneously edit few keyframes in a zero-shot way. Second, we introduce an efficient model called MaskINT, which is built on non-autoregressive masked generative transformers and specializes in frame interpolation between the edited keyframes, using the structural guidance from intermediate frames. Experimental results suggest that our MaskINT achieves comparable performance with diffusion-based methodologies, while significantly improve the inference time. This research offers a practical solution for text-based video editing and showcases the potential of non-autoregressive masked generative transformers in this domain.

1. Introduction

Text-based video editing, which aims to modify a video’s content or style in accordance with a provided text description while preserving the motion and semantic layout, plays an important role in a wide range of applications, including advertisement, live streaming, and the movie industry, etc. This challenging task requires that edited video frames not only match the given text prompt but also ensure the consistency across all video frames.

Recently, numerous studies have showcased the impressive capabilities of diffusion models [22] in image genera-

tion and editing with text prompts [6, 13, 39, 44, 45, 51, 61]. When extending these foundation text-to-image (T2I) models to video generation and editing, a challenge is to maintain the temporal consistency. Existing solutions can be mainly divided into two ways: One is to finetune T2I models with additional temporal modules on paired text-video datasets [5, 9, 17, 23, 32]. However, the lack of extensive paired text-video datasets hinders these works from attaining the same level of performance as seen in image editing. The other involves leveraging a T2I models in a zero-shot manner [8, 19, 29, 41, 62]. These works find that extending the spatial self-attention across frames can produce consistent content across frames [56]. However, extended attention can lead to inconsistent editing of details due to its implicit temporal constraints. Moreover, while these techniques can generate high-fidelity videos, using cumbersome diffusion models to generate all frames is time-consuming. The integration of the extended spatial-temporal attention further exacerbates the running time, making them less practical for real-world applications. This raises the question: *Is it necessary to utilize tedious diffusion models to generate all video frames?*

In this paper, we propose *MaskINT*, an efficient prompt-based video editing framework designed to generate high-quality videos swiftly (Fig. 1). MaskINT disentangles the task into two separate stages. In the first stage, we utilize pre-trained T2I models with extended attention to jointly edit only two keyframes (i.e., the initial and last frames) from the video clip, guided by the provided text prompt. In the second stage, we introduce a novel structure-aware frame interpolation module based on non-autoregressive generative Transformers [10, 11, 60], which generates all intermediate frames in parallel with structural cues and iteratively refine predictions in a few steps. Our framework offers several advantages: First, through the disentanglement of frame editing and interpolation, MaskINT eliminates the requirement for paired text-video datasets during training, thereby enabling us to train the MaskINT using large-scale video-only datasets. Second, the usage of

*Work done during an internship at GenAI, Meta

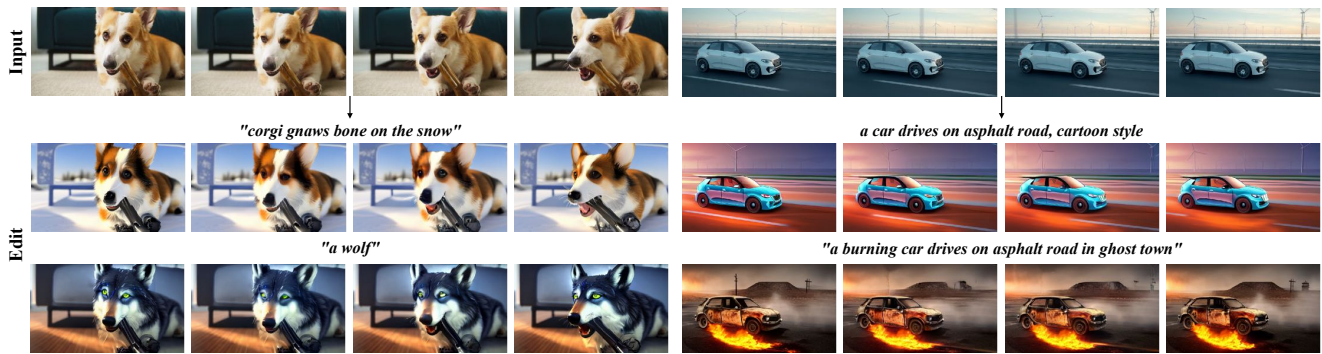


Figure 1. Examples of video editing with MaskINT.

non-autoregressive generation significantly accelerates the processing time. Experimental results indicate that MaskINT achieves comparable performance with pure diffusion methods in terms of temporal consistency and alignment with text prompts, while provides 5-7 times faster inference time. Our work not only provides a practical solution for text-based video editing in terms of the trade-off between quality and efficiency, but also highlights the potential of masked generative transformers in this domain.

Our major contributions are summarized as follows:

- We introduce MaskINT, a two-stage pipeline for text-based video editing, which involves keyframes joint editing and structure-aware frame interpolation, eliminating the need for paired text-video datasets.
- The proposed structure-aware frame interpolation is the pioneer work that explicitly introduces structure control into non-autoregressive generative transformers.
- Experimental results demonstrate that MaskINT achieves comparable performance with diffusion methods, while significantly reducing the inference time.

2. Related Work

Generative Transformers. Following GPT [7], previous works [16, 18, 24, 31, 59] first tokenize images/videos into discrete tokens, and train *Autoregressive Generative Transformers* to perform image/video generation, where tokens are generated sequentially based on previous outputs. However, these methods become exceedingly time-consuming when the length of the token sequence increases.

Recently, *Non-autoregressive Generative Transformers* have emerged as efficient solutions [10, 11, 20, 54, 60]. These works train bidirectional transformers [14] with masked token modeling (MTM). During inference, all discrete tokens are generated in parallel and refined iteratively in a few steps, known as non-autoregressive decoding. Specifically, in image generation, MaskGiT [10] first shows the capability and efficiency of non-autoregressive transformers. It can be seamlessly extended to tasks like inpaint-

ing and outpainting by applying various initial mask constraints. Muse [11] achieves state-of-the-art performance in text-to-image generation by training on large-scale text-image datasets and brings significantly efficiency improvement. StyleDrop [48] further finetunes Muse with human feedback to perform text-to-image generation guided with a reference style image. MaskSketch [3] introduces implicit structural guidance into MaskGiT by calculating the similarity of attention maps in the sampling step. In video generation, MaskViT [20] employ 2D tokenizer and trains a bidirectional window transformer to perform frame prediction. Phenaki [54] trains a masked transformer to generate short video clips condition on text prompt and extends it to arbitrary long video with different prompts in an autoregressive way. MAGViT [60] utilizes 3D tokenizer to quantize videos and trains a single model to perform multiple video generation tasks such as inpainting, outpainting, frame interpolation, etc. However, to the best of our knowledge, there is currently no existing literature in the field of text-based video editing utilizing masked generative transformers. Besides, there is a notable absence of research that delves into explicit structural control within this area.

Diffusion Models in Image Editing. Built upon the latent diffusion models (LDM) [44], numerous studies have achieved significant success in text-based image editing [12, 21, 28, 35, 36, 38, 51, 61, 63]. For example, ControlNet [61], T2I-Adapter [36], and Composer [25] finetune LDM with spatial condition such as depth maps and edge maps, enabling text-to-image synthesis with the same structure as the input image. The PNP [51] incorporates DDIM inversion features [49] from the input image into the T2I generation process, enabling zero-shot image editing with pre-trained T2I models. InstructPix2Pix [6] trains a conditional diffusion model for text-guided image editing using synthetic paired examples, which avoid the tedious inversion. Nevertheless, applying these T2I models on each frame independently usually leads to inconsistencies and flickering.

Diffusion Models in Video Editing. Diffusion models also dominate the field of video generation and video editing [9, 23, 47, 56]. Due to the lack of billion-scale paired video datasets, these works usually finetune existing T2I foundation models [44] with additional temporal modules on paired text-video datasets [17] or overfit a single video [56]. Meanwhile, several works explore using T2I models in a zero-shot manner for video editing [8, 29, 41, 55, 56, 58, 62]. To enable a cohesive appearance, a common approach in these studies involves extending the spatial self-attention module to spatial-temporal attention module and conducting cross-frame attention. In detail, Text2Video-Zero [29] performs cross-frame attention of each frame on the first frame to preserve appearance consistency. ControlVideo [62] extends ControlNet [61] with fully cross-frame attention to joint edit all frames and further improves the performance with interleaved-frame smoother. TokenFlow [19] enhance PNP [51] with extended-attention to jointly edit a few keyframes at each denoising step and propagate the edit frames to the entire video. However, conducting hundreds of denoising steps with extended attention across all frames is time-consuming. In contrast, we only utilize diffusion models to edit a few keyframes, rather than all frames.

Video Frame Interpolation (VFI). VFI aims to generate intermediate images between a pair of frames, which can be applied to creating slow-motion videos and enhancing refresh rate. Advanced methods typically estimate dense motions between frames, like optical flow, and subsequently warping the provided frames to generate intermediate ones [26, 27, 34, 37, 43, 46]. However, these regression methods are most effective with simple or monotonous motion. In our work, we design a generative model to perform frame interpolation with additional structural signals, enabling interpolation with large motion.

3. Preliminaries

3.1. Masked Generative Transformers

Masked generative transformers [10] follow a two-stage pipeline. In the first stage, an image is quantized into a sequence of discrete tokens via a Vector-Quantized (VQ) auto-encoder [16]. In detail, given an image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, an encoder \mathcal{E} encodes it into a series of latent vectors, which are then discretized via nearest neighbor lookup in a codebook of quantized embeddings with size M . Thus, an image can be represented with a sequence of codebook’s indices $\mathbf{Z} = [z_i]_{i=1}^{h \times w}$, $z_i \in \{1, 2, \dots, M\}$, where h and w is the resolution of latent features. A decoder \mathcal{D} can reconstruct the indices back to image $\mathcal{D}(\mathbf{Z}) \approx \mathbf{I}$. In the second stage, a bidirectional transformer model [53] is learned with Masked Token Modeling (MTM). Specifically, during training, a random mask ratio $r \in (0, 1)$ is selected and

$[\gamma(r) \cdot h \times w]$ tokens in \mathbf{Z} are replaced with a special [MASK] token, where $\gamma(r)$ is a mask scheduling function [10]. We denote the corrupted sequence with masked tokens as $\bar{\mathbf{Z}}$ and conditions such as class labels or text prompt as \mathbf{c} . Given the training dataset \mathbb{D} , a BERT [14] parameterized by Φ is learned to minimize the cross-entropy loss between the predicted and the ground truth token at each masked position:

$$\mathcal{L}_{MTM} = \mathbb{E}_{\mathbf{Z} \in \mathbb{D}} \left[\sum_{z_i = [\text{MASK}]} -\log p_{\Phi}(z_i | \bar{\mathbf{Z}}, \mathbf{c}) \right]. \quad (1)$$

During inference, non-autoregressive decoding is applied to generate images. Specifically, all tokens of an image or video are initialized as [MASK] tokens. At step k , all tokens are predicted in parallel while only tokens with the highest prediction scores are kept. The remaining tokens with least prediction scores are masked out and regenerated in the next iteration. The mask ratio is determined by $\gamma(\frac{k}{K})$ at step k , where K is the total number of iteration steps.

3.2. ControlNet

Latent Diffusion Models Denoising Diffusion Probabilistic Models [22] generate images through a progressive noise removal process applied to an initial Gaussian noisy image, carried out over a span of T time steps. To enable efficient high-resolution image generation, LDM [44] operates the diffusion process in the latent space of an auto-encoder. In detail, an encoder \mathcal{E}' first compresses an image \mathbf{I} to a low-resolution latent code $x = \mathcal{E}'(\mathbf{I}) \in \mathbb{R}^{h \times w \times c}$. Second, a U-Net ϵ_{θ} with attention modules [53] is trained to remove the noise with loss

$$\mathcal{L}_{LDM} = \mathbb{E}_{x_0, \epsilon \sim N(0, I), t \sim \cdot} \|\epsilon - \epsilon_{\theta}(x_t, t, \tau)\|_2^2 \quad (2)$$

, where τ is the text prompt and x_t is the noisy latent sample at timestep t .

ControlNet Since accurately describing the layout of generated images using only text prompts is challenging, ControlNet [61] extends LDM by incorporating spatial layout conditions such as edge maps, depth maps, and human poses. In detail, ControlNet train the same U-Net architecture as LDM and finetune it with additional conditions. We denote it as $\epsilon'_{\theta}(x_t, t, \tau, s)$, where s is the spatial layout.

4. Methodology

4.1. Overview

Fig. 2 shows an overview of our MaskINT. MaskINT disentangle the video editing task into keyframe editing stage and structure-aware frame interpolation stage. Specifically, given a video clip with N frames $\{\mathbf{I}^i\}_{i=1}^N$, in the first stage, with the input text prompt τ , we simultaneously edit two

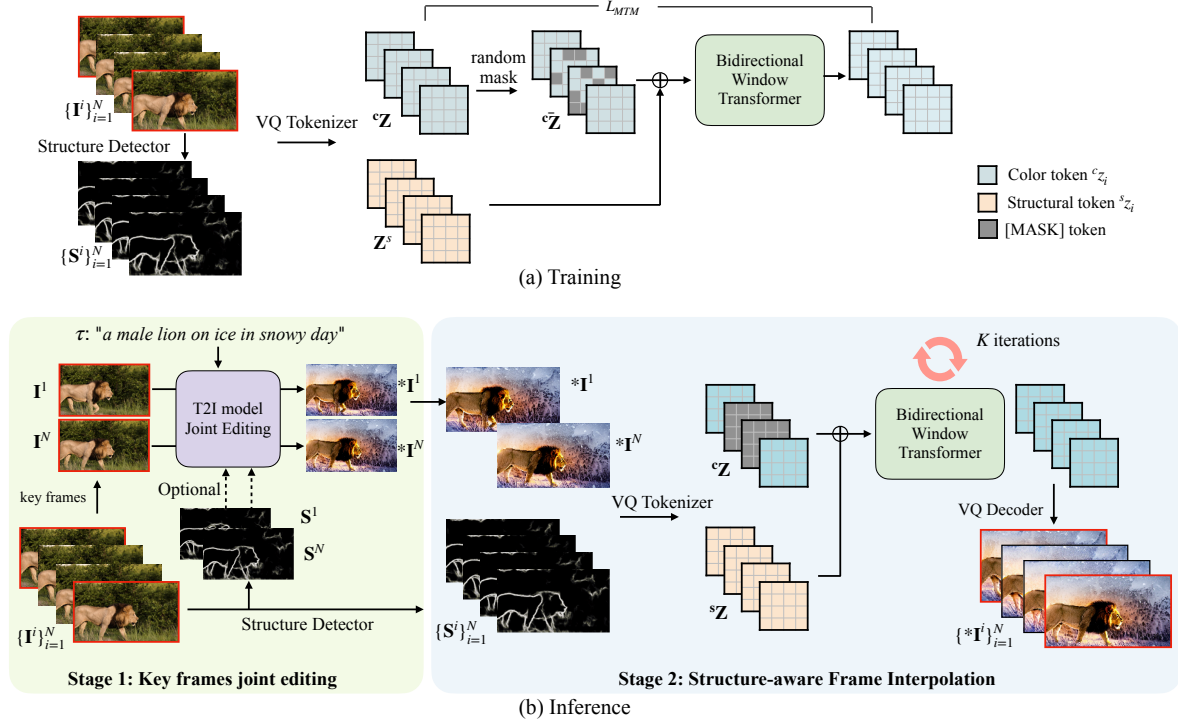


Figure 2. Overview of MaskINT. MaskINT disentangle the video editing task into two separate stages, i.e., keyframes joint editing and structure-aware frame interpolation.

keyframes, i.e., the initial frame \mathbf{I}^1 and last frame \mathbf{I}^N , using existing image-editing model $g(\cdot)$ in a zero-shot way. This process results in the creation of high-quality coherent edited frames $*\mathbf{I}^1, *\mathbf{I}^N = g(\mathbf{I}^1, \mathbf{I}^N, \tau)$. In the second stage, MaskINT performs structure-aware frame interpolation via non-autoregressive transformers. The entire edited video frames are generated by $\{*\mathbf{I}^i\}_{i=1}^N = f_\Phi(*\mathbf{I}^1, *\mathbf{I}^N, \{\mathbf{S}^i\}_{i=1}^N)$, where $\mathbf{S}^i \in [0, 1]^{H \times W \times 1}$ is the structural condition (i.e., the HED edge map [57]). The structure-aware interpolation module is trained with masked token modeling (MTM) on video only datasets, conditioning on the structural signal $\{\mathbf{S}^i\}_{i=1}^N$ as well as the initial frame \mathbf{I}^1 and last frame \mathbf{I}^N .

4.2. Keyframes Joint Editing

To maintain the structure layout, we take ControlNet to edit \mathbf{I}^1 and \mathbf{I}^N based on text prompt τ as well as their edge maps \mathbf{S}^1 and \mathbf{S}^N . However, even with the identical noise, applying ControlNet to each keyframes individually cannot guarantee the appearance consistency. To address this issue, following previous work [19, 56, 62], we extend the self-attention to simultaneously process two keyframes. In detail, the self-attention module projects the noisy feature map x_t^j of j^{th} frame at time step t into query \mathbf{Q}^j , key \mathbf{K}^j , value \mathbf{V}^j in the original U-Net. We extend the self-attention block to perform attention across all selected keyframes by

concatenating their keys and values:

$$\text{Softmax}\left(\frac{\mathbf{Q}^j [\mathbf{K}^1, \mathbf{K}^N]^T}{\sqrt{c}}\right) [\mathbf{V}^1, \mathbf{V}^N]. \quad (3)$$

Since each frame queries all frames and aggregates information from all of them, the appearance consistency can be achieved. Although only two frames were used in Eq. (3), this joint editing can be seamlessly extended to any number of frames, but at the cost of increased processing time and significant GPU memory demand.

4.3. Structure-Aware Frame Interpolation

Structural-aware embeddings. Non-autoregressive masked generative transformers [20, 60] can effectively perform frame prediction and interpolation tasks. However, these works lack explicit control of structure, making it difficult to follow the motion of input videos. Thus, we explicitly introduce structural condition of each frame into the generation process. In detail, we tokenize both RGB frames $\{\mathbf{I}^i\}_{i=1}^N$ and structure maps $\{\mathbf{S}^i\}_{i=1}^N$ with an off-the-shelf 2D VQ tokenizer [16]. We utilize 2D VQ rather than 3D VQ [60] to accommodate diverse numbers of frames and frame rate. We denote tokens from RGB frames as ${}^c\mathbf{Z} = \{c^z_i\}_{i=1}^N \times h \times w$ (color token) and tokens from edge maps as ${}^s\mathbf{Z} = \{s^z_i\}_{i=1}^N \times h \times w$ (structure token), where $s^z_i, c^z_i \in \{1, 2, \dots, M\}$ and M is the codebook size. Subsequently, two distinct embedding layers $e^C(\cdot)$ and $e^S(\cdot)$ are

learned to map tokens ${}^c\mathbf{Z}$ and ${}^s\mathbf{Z}$ into their respective embedding spaces. Learnable 2D spatial positional encoding \mathbf{P}^S and temporal positional encoding \mathbf{P}^T are also added [4]. Thus, the input to the following transformer layers can be formulated by $\mathbf{X} = e^c({}^c\mathbf{Z}) + e^s({}^s\mathbf{Z}) + \mathbf{P}^S + \mathbf{P}^T \in \mathbb{R}^{N \times h \times w \times c}$.

Transformer with Window-Restricted Attention. Previous masked generative transformers [10, 60] employ the standard transformer with global attention [53]. Given that there is no substantial motion between consecutive frames, we adopt self-attention within a restricted window [20] to further mitigate computational overhead. In detail, our approach involves two distinct stages of attention. First, we conduct spatial window attention, where attention is calculated within each frame of dimensions $1 \times h \times w$. We then conduct spatial-temporal window attention, where attention is calculated within a tube of dimensions $N \times h_w \times w_w$, where h_w and w_w is the window size. Besides, to further reduce computation of transformer, we also add a shallow convolution layers to downsample \mathbf{X} before the transformer encoder layers and an upsample layer at the end.

Training. We train MaskINT using video-only datasets. Denote the color token of i^{th} RGB frame as ${}^c\mathbf{Z}^i = \{c_{z^i}\}_{i=1}^{h \times w}$. During the training time, we keep color tokens of the initial frame ${}^c\mathbf{Z}^1$ and last frame ${}^c\mathbf{Z}^N$, and randomly replace $[\gamma(r) \cdot (N - 2) \cdot N]$ color tokens of intermediate frames with the [MASK] tokens. We denote this corrupted video color tokens as ${}^c\bar{\mathbf{Z}} = \{{}^c\mathbf{Z}^1, {}^c\bar{\mathbf{Z}}^2, \dots, {}^c\bar{\mathbf{Z}}^{N-1}, {}^c\mathbf{Z}^N\}$. The structure-aware window-restricted transformer with parameters Θ is trained by

$$\mathcal{L}_{MTM} = \mathbb{E}_{\{c\mathbf{Z}, s\mathbf{Z}\} \in \mathbb{D}} \left[\sum_{c\bar{z}_i = [\text{MASK}]} -\log p_{\Theta}(c_{z_i} | c\bar{\mathbf{Z}}, s\mathbf{Z}) \right]. \quad (4)$$

Inference. During inference, MaskINT can seamlessly generalize to perform frame interpolation between the jointly edited frames, although it is only trained with regular videos. Specifically, we tokenize the the initial and last edited frames $*\mathbf{I}^1$ and $*\mathbf{I}^N$ from Stage 1 into color tokens ${}^c\mathbf{Z}^1$ and ${}^c\mathbf{Z}^N$, and initialize color tokens of all intermediate frames $\{{}^c\bar{\mathbf{Z}}^2, \dots, {}^c\bar{\mathbf{Z}}^{N-1}\}$ with [MASK] tokens. We follow the iterative decoding in MaskGIT [10] with a total number of K steps. At step k , we predict all color tokens in parallel and keep tokens with the highest confidence score.

5. Experiments

5.1. Settings

Implementation Details. We train our model with 100k videos from Shutterstock [1]. During training, we random select a video clip with $T = 16$ frames and frame interval of

1, 2, 4 from each video. All frames are resized to 384×672 . We utilize Retina-VQ [15] with 8 downsample ratio, i.e., each frame has 48×84 tokens. We employ Transformer-Base as our MaskINT and optimized it from scratch with the AdamW optimizer [33] for a duration of 100 epochs. The initial learning rate is set to $1e - 4$ and decayed with cosine schedule. We set the number of decoding step K to 32 and the temperature t to 4.5 in inference.

Evaluation. We evaluate our method on 40 selected object-centric videos [56] from the DAVIS dataset [40], as well as 30 unseen videos from the Shutterstock [1], covering animals, vehicles, etc. For each video, we manually design 5 edited prompts, including object editing, background changes and style transfers. Following previous works [8, 19, 41, 62], we evaluate the quality of the generated videos with the following metrics: 1) Prompt consistency, which calculates the CLIP [42] embedding similarity between text prompt and all video frames. 2) Temporal consistency, which calculates the CLIP embedding similarity of all pairs of consecutive frames. 3) Warping-error [30], which computes the optical flow between consecutive frames using RAFT [50], and warps the edited frames to the next according it. 4) Long-term temporal consistency. Inspired by the global coherence metric in [52], we also calculate the warp error between all possible pairs of frames as the long-term consistency. As for efficiency, we report the running time for generating a 16-frame video clip on a single NVIDIA A6000 GPU.

5.2. Results

We apply ControlNet to each frame individually with the same noise as the baseline. We select methods that built upon T2I diffusion models for comparison, including Tune-A-Video [56], TokenFlow [19], Text-to-video zero [29], and ControlVideo [62]. Besides, we compare with VFI methods such as FILM [43] using the same edited keyframes.

Method	DAVIS				Shutterstock				Time
	P.C.↑	T.C.↑	W.E.↓ / $\times 10^{-3}$	L.C.↓ / $\times 10^{-3}$	P.C.↑	T.C.↑	W.E.↓ / $\times 10^{-3}$	L.C.↓ / $\times 10^{-3}$	
ControlNet per frame [61]	0.314	0.914	36.6	58.6	0.304	0.942	25.9	45.1	50s
Tune-a-Video [56]	0.299	0.966	20.4	48.2	0.292	0.979	13.4	34.9	20min
Text2Video-zero [29]	0.312	0.964	20.7	42.5	0.304	0.981	16.0	33.0	60s
ControlVideo-edge [62]	0.314	0.975	6.9	23.9	0.303	0.986	7.4	20.5	120s
TokenFlow [19]	0.317	0.977	7.0	19.6	0.313	0.987	5.4	15.8	150s
MaskINT (ours)	0.311	0.952	9.5	27.7	0.304	0.971	8.6	22.3	22s

Table 1. Quantitative comparisons. ‘‘T.C.’’ stands for ‘‘temporal consistency’’, ‘‘P.C.’’ stands for ‘‘prompt consistency’’, ‘‘W.E.’’ stands for ‘‘warping-error’’, and ‘‘L.C.’’ stands for ‘‘long-term temporal consistency’’.

Quantitative Comparisons. Tab. 1 summarize the performance of these methods on both DAVIS and Shutterstock datasets. Notably, our method achieves comparable performance with diffusion methods, in terms of all four evaluation metrics, while brings a significant acceleration in processing speed. In detail, the test-time finetuning of TAV

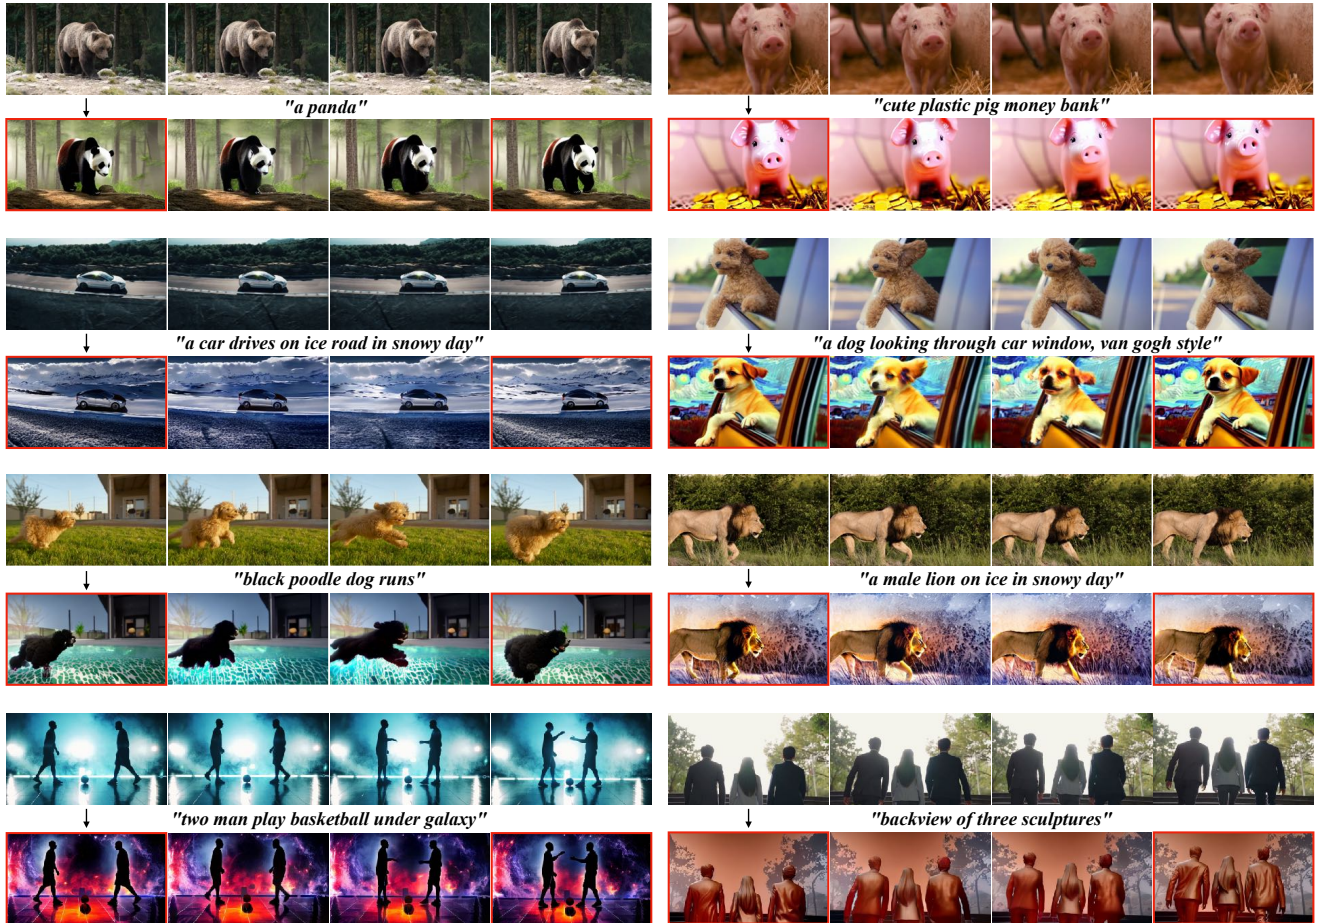


Figure 3. Examples of video editing with MaskINT. Frames with red bounding box are jointly edited keyframes.

[56] is extremely time-consuming, making it a less practical solution. MaskINT is almost 5.5 times faster than ControlVideo [62], whose fully cross-frame attention is computationally extensive. Moreover, MaskINT is nearly 7 times faster than TokenFlow [19], whose DDIM inversion is time-consuming. The efficiency of our method is derived from a combination of a lightweight network and a reduced number of decoding steps with masked generative transformers.

Qualitative Comparisons. Fig. 1 and Fig. 3 show edited video samples using MaskINT. Our method is able to generate temporally consistent videos in accordance with text prompts, enabling diverse applications including stylization, background and foreground editing, and beyond. It is also effective in videos with complex motion, such as jumping and running, as well as scenarios with multiple subjects. Besides, although the structure-aware interpolation module is trained with realistic videos only, it still generalizes to diverse synthesized images. For example, we can successfully generate a video in the style of Van Gogh even in the absence of videos for this particular style in the training set.

Fig. 4 provides a comparison of MaskINT to other methods. Remarkably, diffusion-based methods with extended attention [19, 29, 56, 62] can maintain the overall appearance consistency, but sometimes result in inconsistent details. For example, TokenFlow [19] and Text2Video-Zero[29] exhibit noticeable artifacts in the leg region of the human subjects and ControlVideo [62] produces inconsistent hats. The potential explanation lies in the fact that these methods offer control over temporal consistency implicitly. Besides, FILM [43] generates videos that cannot follow the original motions. Our MaskINT consistently interpolates the intermediate frames based on the structure condition and even maintain better consistency in local regions.

Extension on Long Video Editing. Given that the non-autoregressive pipeline generates all frames simultaneously, it's challenging to edit an entire long video at once due to GPU memory limitation. Nevertheless, our framework can be extended to handle long videos by dividing them into short clips and progressively performing frame interpolation within each segment. For instance, given a video with



Figure 4. Qualitative comparisons with diffusion-based methods. More examples are shown in Supplementary.

60 frames, we select the 1st, 16th, 31st, 46th, and 60th frames as keyframes. We jointly edit these selected 5 frames together and then perform structure-aware frame interpolation within each pair of consecutive keyframes. As shown in Fig. 5, with this design, our method can still generate consistent long videos. Moreover, in this proposed extension, the generation of later frames is dissociated from that of early frames, which differs from the autoregressive generation pipeline in Phenaki [54]. Consequently, even if some early frames encounter difficulties, the generation of later frames can still proceed successfully.

5.3. Ablation Studies

Video Frame Interpolations. In this ablation study, we evaluate the performance of the structure-aware interpolation module. In detail, we perform frame interpolation using original keyframes, and compare the interpolated frames with the original intermediate video frames. We use same testing samples from DAVIS and Shutterstock datasets, employing peak signal-to-noise ratio (PSNR), learned perceptual image patch similarity (LPIPS), and structured similarity (SSIM) as the evaluation metrics. We

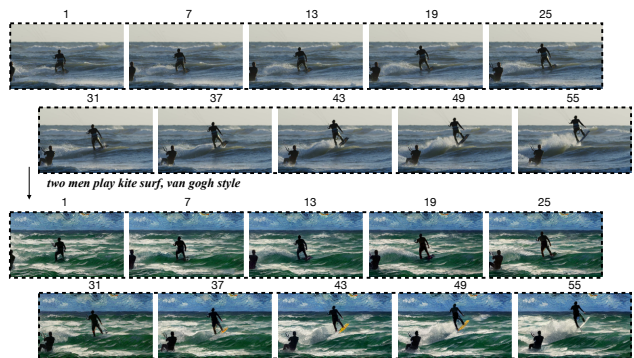


Figure 5. Examples of long video generation. The number indicates the index of frame. More examples are in Supplementary.

benchmark our method against two state-of-the-art VFI methods, namely FILM [43] and RIFE [26]. We also showcase the performance of applying VQ-GAN [16] to all video frames, serving as an upper bound for our method.

As in Tab. 2, our method significantly outperforms VFI

methods on all evaluation metrics, with the benefit of the structure guidance from the intermediate frames. Furthermore, Fig. 6 shows qualitative comparison between FILM [43] and MaskINT. Even when confronted with significant motion between two frames, our method successfully reconstructs the original video, maintaining consistent motion through the aid of structural guidance. In contrast, FILM introduces undesirable artifacts, including disorted background, duplicated cat hands, and the absence of a camel’s head, etc. The major reason is that current VFI models mainly focus on generating slow-motion effects and enhancing frame rate, making them less effective in handling frames with complex motions, which usually requires a better semantic understanding. Besides, the absence of structural guidance poses a challenge for them in accurately aligning generated videos with the original motion.




Figure 6. Qualitative comparisons on video reconstruction with original RGB frames. Frames with red bounding box are given.

Method	DAVIS			ShutterStock		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
RIFE [26]	17.31	0.5195	0.2512	20.44	0.7210	0.1533
FILM [43]	17.00	0.5011	0.2363	20.90	0.7453	0.1246
Ours	22.15	0.6332	0.1483	24.19	0.7616	0.1097
VQGAN (ground truth)	25.66	0.7429	0.0784	27.81	0.8327	0.0561

Table 2. Quantitative comparisons on video frame interpolation with original keyframes.

Number of of keyframes Although our model is trained with frame interpolation by default, MaskINT can seamlessly generalize to an arbitrary number of keyframes without finetuning. We assess the impact of varying the quantity of keyframes on the generation performance. As shown in the left part of Table 3, with an increase in the number

of keyframes, the model exhibits an improvement in performance. Generally, with more information, performing frame interpolation is easier. However, simultaneously editing more frames requires longer time due to the global attention among them.

Decoding steps We also explore the number of decoding steps K for the masked generative transformers in the second stage. The right part of Tab. 3 shows that more decoding steps can bring slight improvement on the temporal consistency, but requires more time. Considering the trade-off between performance and efficiency, we chose $K = 32$ steps by default in all experiments.

# keyframes	T.C.	P.C.	Time	#decoding step K	T.C.	P.C.	Time
1	0.9690	0.2984	19s	16	0.9691	0.3038	15s
2	0.9714	0.3038	22s	32	0.9714	0.3038	22s
3	0.9721	0.3051	26s	64	0.9719	0.3040	33s
4	0.9728	0.3069	29s	128	0.9720	0.3041	62s
6	0.9737	0.3035	35s				

Table 3. Ablation study on the number of keyframes and the number of decoding steps K . “T.C.” stands for “temporal consistency” and “P.C.” stands for “prompt consistency”.

6. Limitation and Future Work

One limitation of MaskINT is that it can only perform structure-preserving video editing, such as altering style or appearance. Thus, it cannot handle edits that require structural changes, for example, change a dog to a horse, a limitation shared with TokenFlow [19]. It also requires that no new objects should appear in the intermediate frames. Besides, the performance of MaskINT highly relies on the tedious image-editing model and structure detector. When these models fail, the structure-aware interpolation becomes meaningless, resulting in artifacts. In the future, given that our approach disentangles video editing into two distinct stages, we intend to explore the integration of token-based methods, like Muse [11], for image editing. This endeavor aims to further enhance the efficiency of the initial stage.

7. Conclusion

We propose MaskINT towards consistent and efficient video editing with prompt. MaskINT disentangles this task into keyframes joint editing with T2I model and structure-aware frame interpolation with non-autoregressive masked transformers. Experimental results demonstrate that MaskINT achieves comparable performance with pure diffusion-based methods while significantly accelerate the inference up to 7x. We believe our work demonstrates the substantial promise of non-autoregressive generative transformers within the realm of video editing.

References

- [1] Stock footage video, royalty-free hd, 4k video clips. 2023. 5, 1

- [2] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 1
- [3] Dina Bashkirova, José Lezama, Kihyuk Sohn, Kate Saenko, and Irfan Essa. Masksketch: Unpaired structure-guided masked image generation. In *CVPR*, pages 1879–1889, 2023. 2
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 5
- [5] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 1
- [6] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, pages 18392–18402, 2023. 1, 2
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020. 2
- [8] Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. Pix2video: Video editing using image diffusion. In *ICCV*, pages 23206–23217, 2023. 1, 3, 5
- [9] Wenhao Chai, Xun Guo, Gaoang Wang, and Yan Lu. Stable-video: Text-driven consistency-aware diffusion video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23040–23050, 2023. 1, 3
- [10] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *CVPR*, pages 11315–11325, 2022. 1, 2, 3, 5
- [11] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *ICML*, 2023. 1, 2, 8
- [12] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *ICLR*, 2023. 2
- [13] Xiaoliang Dai, Ji Hou, Chih-Yao Ma, Sam Tsai, Jiali Wang, Rui Wang, Peizhao Zhang, Simon Vandenhende, Xiaofang Wang, Abhimanyu Dubey, et al. Emu: Enhancing image generation models using photogenic needles in a haystack. *arXiv preprint arXiv:2309.15807*, 2023. 1
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL*, 2019. 2, 3
- [15] Abhimanyu Dubey, Filip Radenovic, Dhruv Mahajan, and Vignesh Ramanathan. Retina vq. 2023. 5
- [16] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pages 12873–12883, 2021. 2, 3, 4, 7
- [17] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. *ICCV*, 2023. 1, 3
- [18] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In *ECCV*, pages 102–118. Springer, 2022. 2
- [19] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. *ICLR*, 2024. 1, 3, 4, 5, 6, 8
- [20] Agrim Gupta, Stephen Tian, Yunzhi Zhang, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Maskvit: Masked visual pre-training for video prediction. In *ICLR*, 2023. 2, 4, 5
- [21] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020. 1, 3
- [23] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 1, 3
- [24] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *ICLR*, 2023. 2
- [25] Lianghua Huang, Di Chen, Yu Liu, Yujun Shen, Deli Zhao, and Jingren Zhou. Composer: Creative and controllable image synthesis with composable conditions. *arXiv preprint arXiv:2302.09778*, 2023. 2
- [26] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Real-time intermediate flow estimation for video frame interpolation. In *ECCV*, pages 624–642, 2022. 3, 7, 8
- [27] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *CVPR*, pages 9000–9008, 2018. 3
- [28] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *CVPR*, pages 6007–6017, 2023. 2
- [29] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. *ICCV*, 2023. 1, 3, 5, 6
- [30] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning blind video temporal consistency. In *ECCV*, 2018. 5
- [31] Guillaume Le Moing, Jean Ponce, and Cordelia Schmid. Ccvs: context-aware controllable video synthesis. *NeurIPS*, 34:14042–14055, 2021. 2
- [32] Jun Hao Liew, Hanshu Yan, Jianfeng Zhang, Zhongcong Xu, and Jiashi Feng. Magicedit: High-fidelity and temporally coherent video editing. *arXiv preprint arXiv:2308.14749*, 2023. 1

- [33] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [34] Liying Lu, Ruizheng Wu, Huaijia Lin, Jiangbo Lu, and Jiaya Jia. Video frame interpolation with transformer. In *CVPR*, pages 3532–3542, 2022. 3
- [35] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *ICLR*, 2022. 2
- [36] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhonggang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. 2
- [37] Simon Niklaus, Long Mai, and Feng Liu. Video frame interpolation via adaptive separable convolution. In *ICCV*, pages 261–270, 2017. 3
- [38] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2
- [39] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *JCCV*, pages 4195–4205, 2023. 1
- [40] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017. 5
- [41] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. *ICCV*, 2023. 1, 3, 5
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 5
- [43] Fitsum Reda, Janne Kontkanen, Eric Tabellion, Deqing Sun, Caroline Pantofaru, and Brian Curless. Film: Frame interpolation for large motion. In *ECCV*, 2022. 3, 5, 6, 7, 8
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 1, 2, 3
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 35:36479–36494, 2022. 1
- [46] Hyeonjun Sim, Jiyoung Oh, and Munchurl Kim. Xvfi: extreme video frame interpolation. In *ICCV*, pages 14489–14498, 2021. 3
- [47] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *ICLR*, 2023. 3
- [48] Kihyuk Sohn, Nataniel Ruiz, Kimin Lee, Daniel Castro Chin, Irina Blok, Huiwen Chang, Jarred Barber, Lu Jiang, Glenn Entis, Yuanzhen Li, et al. Styledrop: Text-to-image generation in any style. *arXiv preprint arXiv:2306.00983*, 2023. 2
- [49] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *ICLR*, 2021. 2
- [50] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 5
- [51] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *CVPR*, pages 1921–1930, 2023. 1, 2, 3
- [52] Rotem Tzaban, Ron Mokady, Rinon Gal, Amit Bermano, and Daniel Cohen-Or. Stitch it in time: Gan-based facial editing of real videos. In *SIGGRAPH Asia*, 2022. 5
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 3, 5
- [54] Ruben Villegas, Mohammad Babaeizadeh, Pieter-Jan Kindermans, Hernan Moraldo, Han Zhang, Mohammad Taghi Saffar, Santiago Castro, Julius Kunze, and Dumitru Erhan. Phenaki: Variable length video generation from open domain textual descriptions. In *ICLR*, 2023. 2, 7
- [55] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *NeurIPS*, 2023. 3
- [56] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Weixian Lei, Yuchao Gu, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. *ICCV*, 2023. 1, 3, 4, 5, 6
- [57] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *ICCV*, pages 1395–1403, 2015. 4
- [58] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. *ACM SIGGRAPH Asia*, 2023. 3
- [59] Jiahui Yu, Xin Li, Jing Yu Koh, Han Zhang, Ruoming Pang, James Qin, Alexander Ku, Yuanzhong Xu, Jason Baldridge, and Yonghui Wu. Vector-quantized image modeling with improved VQGAN. In *ICLR*, 2022. 2
- [60] Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video transformer. In *CVPR*, pages 10459–10469, 2023. 1, 2, 4, 5
- [61] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *ICCV*, 2023. 1, 2, 3, 5
- [62] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. *ICLR*, 2024. 1, 3, 4, 5, 6
- [63] Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris N Metaxas, and Jian Ren. Sine: Single image editing with text-to-image diffusion models. In *CVPR*, pages 6027–6037, 2023. 2