Be-Your-Outpainter: Mastering Video Outpainting through Input-Specific Adaptation

Fu-Yun Wang¹⁽⁶⁾ Xiaoshi Wu¹ Zhaoyang Huang²⁽⁶⁾ Xiaoyu Shi¹⁽⁶⁾ Dazhong Shen³⁽⁶⁾ Guanglu Song⁴⁽⁶⁾ Yu Liu³ Hongsheng Li^{1,3,5}⁽⁶⁾

¹MMLab, CUHK ²Avolution AI ³Shanghai AI Lab ⁴SenseTime Research ⁵ CPII under InnoHK {fywang@link, hsli@ee}.cuhk.edu.hk

Abstract. Video outpainting is a challenging task, aiming at generating video content outside the viewport of the input video while maintaining inter-frame and intra-frame consistency. Existing methods fall short in either generation quality or flexibility. We introduce MOTIA (Mastering Video Outpainting Through Input-Specific Adaptation), a diffusion-based pipeline that leverages both the intrinsic data-specific patterns of the source video and the image/video generative prior for effective outpainting. MOTIA comprises two main phases: input-specific adaptation and pattern-aware outpainting. The input-specific adaptation phase involves conducting efficient and effective pseudo outpainting learning on the single-shot source video. This process encourages the model to identify and learn patterns within the source video, as well as bridging the gap between standard generative processes and outpainting. The subsequent phase, pattern-aware outpainting, is dedicated to the generalization of these learned patterns to generate outpainting outcomes. Additional strategies are proposed to better leverage the diffusion model's generative prior and the acquired video patterns from source videos for inference. Extensive evaluations underscore MOTIA's superiority, outperforming existing state-of-the-art methods in widely recognized benchmarks. Notably, these advancements are achieved without necessitating extensive, task-specific tuning. More details are available at https://be-your-outpainter.github.io/.

1 Introduction

Video outpainting aims to expand the visual content out of the spatial boundaries of videos, which has important real-world applications [3, 5, 6]. For instance, in practice, videos are usually recorded with a fixed aspect ratio, such as in movies or short clips. This becomes an issue when viewing these videos on smartphones, which often have varying aspect ratios, resulting in unsightly black bars on the screen that detract from the viewing experience. Proper ways for video outpainting are crucial in solving this issue. By expanding the visual content beyond the original frame, it adapts the video to fit various screen sizes seamlessly. This process ensures that the audience enjoys a full-screen experience without any



Fig. 1: MOTIA is a high-quality flexible video outpainting pipeline, leveraging the intrinsic data-specific patterns of source videos and image/video generative prior for state-of-the-art performance. Quantitative metric improvement of MOTIA is significant (Table 1).

compromise in visual integrity. However, the challenges associated with video outpainting are significant. It requires not just the expansion of each frame's content but also the preservation of temporal (inter-frame) and spatial (intra-frame) consistency across the video.

Currently, there are two primary approaches to video outpainting. The first employs optical flows and specialized warping techniques to extend video frames, involving complex computations and carefully tailored hyperparameters to ensure the added content remains consistent [5, 7]. However, their results are far from satisfactory, suffering from blurred content. The other type of approach in video outpainting revolves around training specialized models tailored for video inpainting and outpainting with extensive datasets [6, 35]. However, they have two notable limitations: 1) An obvious drawback of these models is their dependency on the types of masks and the resolutions of videos they can handle, which significantly constrains their versatility and effectiveness in real-world applications, as they may not be adequately equipped to deal with the diverse range of video formats and resolutions commonly encountered in practical scenarios. 2) The other drawback is their inability to out-domain video outpainting, even intensively trained on massive video data. Fig. 5 shows a failure example of most advanced previous work [6] that the model faces complete outpainting failure, with only blurred corners. We show that a crucial reason behind this is that the model fails at capturing the intrinsic data-specific patterns from out-domain source (input) videos.

In this work, we propose **MOTIA**: Mastering Video Outpainting Through Input-Specific Adaptation, a diffusion-based method for open-domain video outpainting with arbitrary types of mask, arbitrary video resolutions and lengths, and arbitrary styles. At the heart of MOTIA is treating the source video itself as a rich source of information [18, 22], which contains key motion and content patterns (intrinsic data-specific patterns) necessary for effective outpainting. We demonstrate that the patterns learned from the source video, combined with the generative capabilities of diffusion models, can achieve surprisingly great outpainting performance.

Fig. 2 illustrates the workflow of MOTIA. MOTIA consists of two stages: input-specific adaptation and pattern-aware outpainting. During the input-specific adaptation stage, we conduct pseudo video outpainting learning on the source video (videos to be outpainted) itself. Specifically, at each iteration, we heuristically add random masks to the source video and prompt the base diffusion model to recover the masked regions by learning to denoise the video corrupted by white noise, relying on the extracted information from unmasked regions. This process not only allows the model to capture essential data-specific patterns from the source video but also narrows the gap between standard generation and outpainting. We insert trainable lightweight adapters into the diffusion model for tuning to keep the efficiency and stability. In the pattern-aware outpainting stage, we combine the learned patterns from the source video and the generation prior of the diffusion model for effective outpainting. To better leverage the generation ability of the pretrained diffusion model and the learned pattern from the source video, we propose spatial-aware insertion (SA-Insertion) of the tuned adapters for outpainting. Specifically, the insertion weights of adapters gradually decay as the spatial position of features away from the known regions. In this way, the outpainting of pixels near the known regions is more influenced by the learned patterns, while the outpainting of pixels far from the known regions relies more on the original generative prior of diffusion model. To further mitigate potential denoising conflicts and enhance the knowledge transfer between known regions and unknown regions, we incorporate noise regret that we add noise and denoise periodically at early inference steps, which works for more harmonious outpainting results.

Extensively quantitative and qualitative experiments verify the effectiveness of our proposed method. MOTIA overcomes many limitations of previous methods and outperforms the state-of-the-art intensively trained outpainting method in standard widely used benchmarks. In summary, our contribution is three-fold: 1) We show that the data-specific patterns of source videos are crucial for effective outpainting, which is neglected by previous work. 2) We introduce an adaptation strategy to effectively capture the data-specific patterns and then propose novel strategies to better leverage the captured patterns and pretrained image/video generative prior for better results. 3) Vast experiments verify that

4 Fu-Yun Wang et al.



Fig. 2: Workflow of MOTIA. Blue lines represent the workflow of input-specific adaptation, and green lines represent the workflow of pattern-aware outpainting.

our performance in video outpainting is great, significantly outperforming previous state-of-the-art methods in both quantitative metrics and user studies.

2 Related Works

In this section, we discuss related diffusion models and outpainting methods. **Diffusion models.** Diffusion models (a.k.a., score-based models) [9,10,17,20,26] have gained increasing attention due to its amazing ability to generate highly-detailed images. Current video diffusion models [4, 9, 15, 16, 24, 25, 28, 30] are generally built upon the extension of image diffusion models through inserting temporal layers. They are either trained with image-video joint tuning [12, 25] or trained with spatial weights frozen [4] to mitigate the negative influence of the poor captions and visual quality of video data.

Ooutpainting methods. Video outpainting is largely built upon the advancements in image outpainting, where techniques ranged from patch-based methods (e.g., PatchMatch [3]) to more recent deep learning approaches like GANs [1,34]. Diffusion models [2,14], benefiting from the learned priors on synthesis tasks and the process of iterative refinement, achieve state-of-the-art performance on image outpainting tasks. The research focusing on video outpainting is comparatively few. Previous works typically apply optical flow for outpainting, which warps content from adjacent frames to the outside corners, but their results are far from satisfactory. Recently, M3DDM [6] trained a large 3D diffusion models with specially designed architecture for outpainting on massive video data, achieving much better results compared to previous methods, showcasing the huge potential of diffusion models on video outpainting. However, as we claimed, they have two main limitations: 1) The inflexibility for mask types and video resolutions. They can only outpaint video with resolution 256×256 with square type of masking. 2) Inability for out-domain video outpainting. As shown in Fig. 5, they encounter outpainting failure when processing out-domain videos even intensively trained on massive video data.

3 Preliminaries

Diffusion models [10] add noise to data through a Markov chain process. Given initial data $x_0 \sim q(x_0)$, this process is formulated as:

$$q(\boldsymbol{x}_{1:T}|\boldsymbol{x}_0) = \prod_{t=1}^T q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}), \quad q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) = \mathcal{N}(\boldsymbol{x}_t|\sqrt{\alpha_t}\boldsymbol{x}_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

where β_t is the noise schedule and $\alpha_t = 1 - \beta_t$. The data reconstruction, or denoising process, is accomplished by the reverse transition modeled by $p_{\theta}(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t)$:

$$q(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t, \boldsymbol{x}_0) = \mathcal{N}(\boldsymbol{x}_{t-1}; \tilde{\boldsymbol{\mu}}_t(\boldsymbol{x}_t, \boldsymbol{x}_0), \hat{\beta}_t \mathbf{I}),$$
(2)

with $\tilde{\boldsymbol{\mu}}_t(\boldsymbol{x}_t, \boldsymbol{x}_0) = \frac{1}{\sqrt{\alpha_t}} \boldsymbol{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}\sqrt{\alpha_t}} \boldsymbol{\epsilon}, \ \bar{\alpha}_t = \prod_{s=1}^t \alpha_s, \ \tilde{\beta}_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t} \beta_t$, and $\boldsymbol{\epsilon}$ is the noise added to \boldsymbol{x}_0 to obtain \boldsymbol{x}_t .

Diffusion-based outpainting aims to predict missing pixels at the corners of the masked region with the pre-trained diffusion models. We denote the ground truth as \boldsymbol{x} , mask as \boldsymbol{m} , known region as $(\boldsymbol{1} - \boldsymbol{m}) \odot \boldsymbol{x}$, and unknown region as $\boldsymbol{m} \odot \boldsymbol{x}$. At each reverse step in the denoising process, we modify the known regions by incorporating the intermediate noisy state of the source data from the corresponding timestep in the forward diffusion process (which adds noise), provided that this maintains the correct distribution of correspondences. Specifically, each reverse step can be denoted as the following formulas:

$$\boldsymbol{x}_{t-1}^{\text{known}} \sim \mathcal{N}\left(\sqrt{\bar{\alpha}_{t}}\boldsymbol{x}_{0}, (1-\bar{\alpha}_{t})\,\mathbf{I}\right), \quad \boldsymbol{x}_{t-1}^{\text{unknown}} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\theta}\left(\boldsymbol{x}_{t}, t\right), \boldsymbol{\Sigma}_{\theta}\left(\boldsymbol{x}_{t}, t\right)\right), \quad (3)$$

$$\boldsymbol{x}_{t-1} = \boldsymbol{m} \odot \boldsymbol{x}_{t-1}^{\text{known}} + (\boldsymbol{1} - \boldsymbol{m}) \odot \boldsymbol{x}_{t-1}^{\text{unknown}}, \qquad (4)$$

where the x_{t-1}^{known} is padded to the target resolution before the masked merging.

4 Methodology

This section presents MOTIA, a method demonstrating exceptional performance in video outpainting tasks. We begin by defining the concept of video outpainting and describing the foundational model in Section 4.1. and Section 4.2. Subsequently, we delve into the specifics of input-specific adaptation and pattern-aware outpainting in Sections 4.3 and 4.4, respectively. Moreover, we show that our approach has great promise in extending its application to long video outpainting, which will be explored in Section 4.5.



Fig. 3: Sample results of quantitative experiments. All videos are outpainted with a horizontal mask ratio of 0.66. Contents outside the yellow lines are outpainted by MOTIA.

4.1 Problem Formulation

For a video represented as $\boldsymbol{v} \in \mathbb{R}^{t \times h \times w \times d}$, where t denotes the number of frames, h denotes frame height, w denotes frame width, and d denotes channel depth. Video outpainting model $f(\boldsymbol{v})$ is designed to generate a spatially expanded video $\boldsymbol{v}' \in \mathbb{R}^{t \times h' \times w' \times d}$. This process not only increases the spatial dimensions (h' > h, w' > w), but also requires to ensure continuity and harmony between the newly expanded regions and the known regions. The transformation maintains the known regions unchanged, with $f(\boldsymbol{v})$ acting as an identity in these regions.

4.2 Network Expansion

Network inflation. MOTIA leverages the pre-trained text-to-image (T2I) model, Stable Diffusion. In line with previous video editing techniques [32], we transform 2D convolutions into pseudo 3D convolutions and adapt 2D group normalizations into 3D group normalizations to process video latent features. Specifically, the 3×3 kernels in convolutions are replaced by $1 \times 3 \times 3$ kernels, maintaining identical weights. Group normalizations are executed across both temporal and spatial dimensions, meaning that all 3D features within the same group are normalized simultaneously, followed by scaling and shifting.

Masked video as conditional input. Additionally, we incorporate a ControlNet [36], initially trained for image inpainting, to manage additional mask inputs. Apart from noise input, ControlNet can also process masked videos to extract effective information for more controllable denoising. In these masked videos, known regions have pixel values ranging from 0 to 1, while values of masked regions are set to -1.

Temporal consistency prior. To infuse the model with temporal consistency priors, we integrate temporal modules pre-trained on text-to-video (T2V) tasks. Note that although MOTIA relies on pre-trained video diffusion modules, applying these pre-trained temporal modules directly for video outpainting yields rather bad results, significantly inferior to all baseline methods (Table. 2). However, when equipped with our proposed MOTIA, the model demonstrates superior performance even in comparison to models specifically designed and trained for video outpainting, underscoring the efficacy of MOTIA.

4.3 Input-Specific Adaptation

The input-specific adaptation phase is crucial in our video outpainting method, aiming to tailor the model for the specific challenges of outpainting. This phase involves training on the source video with a pseudo-outpainting task, importantly, enabling the model to learn intrinsic content and motion patterns (dataspecific patterns) within the source video as well as narrowing the gap between the standard generation process and outpainting.

Video augmentation. Initially, we augment the source video. Transformations like identity transformation, random flipping, cropping, and resizing can be employed. This step can potentially help the model better learn and adapt to diverse changes in video content. For longer video outpainting, as we will discuss later, instead of taking it as a whole, we randomly sample short video clips from it to reduce the cost of the adaptation phase.

Video masking. We then add random masks to the video. We adopt a heuristic approach that uniformly samples edge boundaries of 4 sides within given limits. The area enclosed by these boundaries is considered the known region, while the rest is the unknown region. This masked video serves as the conditional input for the ControlNet, simulating the distribution of known and unknown areas in actual outpainting scenarios.

Video noising. Additionally, we apply noise to the video following the DDPM [10] by randomly selecting diffusion timesteps. This noisy video serves as an input for both the ControlNet and the Stable Diffusion model.

Optimization. Finally, we optimize the model. To ensure efficiency, low-rank adapters are inserted into the layers of the diffusion model. We optimize only the parameters of these adapters while keeping the other parameters frozen. The loss function is

$$\mathcal{L} = \left\| \boldsymbol{\epsilon} - \hat{\boldsymbol{\epsilon}}_{\bar{\boldsymbol{\theta}}_l, \bar{\boldsymbol{\theta}}_c, \boldsymbol{\theta}_a}(\boldsymbol{v}_{\text{noisy}}, \boldsymbol{v}_{\text{masked}}, t) \right\|_2, \qquad (5)$$



Fig. 4: Spatial-aware insertion scales the insertion weights of adapters for better leveraging of learned patterns and generative prior.



Fig. 5: Failure example of previous methods. Many previous methods including the intensively trained models on video outpainting still might suffer from generation failure, that the model simply generates blurred corners. MOTIA never encounters this failure.

where t represents the timestep in the process, $\boldsymbol{\epsilon}$ is the added noise, $\boldsymbol{v}_{\text{noisy}}$ refers to the video perturbed by $\boldsymbol{\epsilon}$, and $\boldsymbol{v}_{\text{masked}}$ denotes the masked video. The parameters $\boldsymbol{\theta}_l$, $\boldsymbol{\theta}_c$, and $\boldsymbol{\theta}_a$ correspond to the Diffusion Model, ControlNet, and adapters, respectively. The bar over these parameters indicates they are frozen during the optimization. This optimization process, including the steps of augmentation, masking, and noising, is repeated to update the lightweight adapters to capture the data-specific patterns from the source video.

4.4 Pattern-Aware Outpainting

Following the initial phase of input-specific adaptation, our model shows promising results in video outpainting using basic pipelines as outlined in Eq. 3 and Eq. 4, achieving commendable quality. However, we here introduce additional inference strategies that can be combined to better leverage the learned dataspecific patterns from the input-specific adaptation phase for better outpainting results. We call the outpainting process that incorporates these strategies pattern-aware outpainting.

Spatial-aware insertion. It is important to acknowledge that in the inputspecific adaptation phase, the model is fine-tuned through learning outpainting within the source video. However, at the outpainting phase, the model is expected to treat the entire source video as known regions and then fill the unknown regions at edges (*i.e.*, generating a video with a larger viewport and resolution). This specificity may lead to a noticeable training-inference gap during outpainting, potentially affecting the outpainting quality. To balance the fine-tuned patterns with the diffusion model's inherent generative prior, we introduce the concept of spatial-aware insertion (SA-Insertion) of adapters as shown in Fig. 4. The adaptation involves adjusting the insertion weight of tuned low-rank adapters based on the feature's spatial position. We increase insertion weight near known areas to utilize the learned patterns while decreasing it in farther regions to rely more on the original generative capacity of the diffusion model. To be specific,

$$\mathbf{W}_{\text{adapted}}^{\top} \boldsymbol{x}_{\boldsymbol{p}} = \mathbf{W}^{\top} \boldsymbol{x}_{\boldsymbol{p}} + \alpha(\boldsymbol{p}) \left(\mathbf{W}_{\text{up}} \mathbf{W}_{\text{down}} \right)^{\top} \boldsymbol{x}_{\boldsymbol{p}}.$$
 (6)

Here, \boldsymbol{p} signifies the spatial position of $\boldsymbol{x}, \mathbf{W} \in \mathbb{R}^{d_{\text{in}} \times d_{\text{out}}}$ denotes the linear transformation in layers of diffusion model, $\mathbf{W}_{\text{down}} \in \mathbb{R}^{d_{\text{in}} \times r}$ and $\mathbf{W}_{\text{up}} \in \mathbb{R}^{r \times d_{\text{out}}}$ are the linear components of the adapter with rank $r \ll \min(d_{\text{in}}, d_{\text{out}})$. The function $\alpha(\boldsymbol{p})$ is defined as:

$$\alpha(\boldsymbol{p}) = \exp(-\frac{K\|\boldsymbol{p} - \boldsymbol{p}_c\|}{\max_{\bar{\boldsymbol{p}}} \|\bar{\boldsymbol{p}} - \boldsymbol{p}_c\|}),\tag{7}$$

where K is a constant for controlling decay speed, and p_c represents the nearest side of the known region to p.

Noise regret. In the context of Eq. 3, merging noisy states from known and unknown regions in video outpainting tasks is similar to sampling from two different vectors, which can potentially disrupt the denoising direction. Given the significant impact of early steps on the generation's structure, later denoising may not rectify these initial discrepancies. Inspired by DDPM-based image inpainting methods [14,21], we propose to re-propagate the noisy state into a noisier state by adding noise when denoising and then provide the model a second chance for re-denoising. This helps integrate known region data more effectively and reduces denoising direction conflicts. In detail, after obtaining v_t during denoising, we conduct

$$\boldsymbol{v}_{t+L} = \sqrt{\Pi_{i=t+1}^{t+L}\alpha_i}\boldsymbol{v}_t + \sqrt{1 - \Pi_{i=t+1}^{t+L}\alpha_i}\boldsymbol{\epsilon}, \qquad (8)$$

where $\alpha_i = 1 - \beta_i$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Then we restart the denoising process. We repeat this progress for M times. We only conduct it in the early denoising steps.

4.5 Extension to Long Video Outpainting

We show that our method can be easily extended for long video outpainting. Specifically, for the stage of input-specific adaptation, instead of taking the long video as a whole for adaptation (Direct adaptation on long videos is costly and does not align with the video generation prior of the pretrained modules), we randomly sample short video clips from the long video for tuning to learn global patterns without requiring more GPU memory cost. For the stage of pattern-aware outpainting, we split the long video into short video clips with temporal overlapping (*i.e.*, some frames are shared by different short video clips), and then conduct temporal co-denoising following Gen-L [29]. Specifically, the denoising result for j^{th} frame of the long video at timestep t is approximated by the weighted sum of all the corresponding frames in short video clips that contain it,

$$\boldsymbol{v}_{t-1,j} = \frac{\sum_{i \in \mathcal{I}^j} \left(\left(W_{i,j^*} \right)^2 \otimes \boldsymbol{v}_{t-1,j^*}^i \right)}{\sum_{i \in \mathcal{I}^j} \left(W_{i,j^*}^2 \right)^2}, \qquad (9)$$



Fig. 6: Qualitative comparison. Other methods outpainting the source video with a mask ratio of 0.6. MOTIA outpainting the source video with a larger mask ratio of 0.66 while achieving obviously better outpainting results.

where \otimes denotes element-wise multiplication, $v_{t-1,j}$ denotes the noisy state of the j^{th} frame at timestep t, v_{t-1,j^*}^i is the noisy state of j^{th} frame predicted with only information from the i^{th} video clip at timestep t, W_{i,j^*} is the per-pixel weight, which is as **1** as default.

5 Experiments

5.1 Experimental Setup

Benchmarks. To verify the effectiveness of MOTIA, we conduct evaluations on DAVIS [19] and YouTube-VOS [33], which are widely used benchmarks for video outpainting. Following M3DDM [6], we compare the results of different methods in the horizontal direction, using mask ratios of 0.25 and 0.66.

Evaluation metrics. Our evaluation approach utilizes four well-established metrics: Peak Signal to Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [31], Learned Perceptual Image Patch Similarity (LPIPS) [37], and Frechet Video Distance (FVD) [27]. For assessing PSNR, SSIM, and FVD, the generated videos are converted into frames within a normalized value range of [0, 1]. LPIPS is evaluated over a range of [-1, 1]. About the FVD metric, we adopt a uniform frame sampling, with 16 frames per video for evaluation following M3DDM.

Implementation details. Our method is built upon Stable Diffusion v1-5. We add the ControlNet pretrained on image inpainting to enable the model to accept additional masked image inputs. The temporal modules are initialized with the weights from pretrained motion modules [8] to obtain additional motion priors. The motion modules are naive transformer blocks trained with solely text-to-video tasks on WebVid. For the input-specific adaptation, the low-rank adapters are trained using the Adam optimizer. We set the learning rate to 10^{-4} , and set the weight decay to 10^{-2} . The LoRA rank and α_{lora} are set to 16 and,

Table 1: Quantitative comparison of video outpainting methods on DAVIS and YouTube-VOS datasets. \uparrow means 'better when higher', and \downarrow indicates 'better when lower'.

Mathad	DAVIS				YouTube-VOS				
Method	$PSNR \uparrow$	SSIM \uparrow	LPIPS \downarrow	$\mathrm{FVD}\downarrow$	PSNR ↑	\uparrow SSIM \uparrow	LPIPS \downarrow	$\mathrm{FVD}\downarrow$	
VideoOutpainting [5]	17.96	0.6272	0.2331	363.1	18.25	0.7195	0.2278	149.7	
SDM [6]	20.02	0.7078	0.2165	334.6	19.91	0.7277	0.2001	94.81	
M3DDM [6]	20.26	0.7082	0.2026	300.0	20.20	0.7312	0.1854	66.62	
MOTIA	20.36	0.7578	0.1595	286.3	20.25	0.7636	0.1727	58.99	

8, respectively. The number of training steps is set to 1000. We do not apply augmentation for simplicity. For both mask ratios of 0.66 and 0.25, we simply apply the same random mask strategy, which uniformly crops a square in the middle as the known regions. For the pattern-aware outpainting, the diffusion steps are set to 25 and the classifier-free guidance (CFG) [11,23] scale is set to 7.5 and we only apply CFG at the first 15 inference steps. When adding noise regret to further improve the results, we set jump length L = 3, and repeat time M = 4. We only apply noise regret in the first half inference steps. Note that our method is built upon LDM, which requires text-conditional inputs. For a fair comparison and to remove the influence of the choice of text prompt, we apply Blip [13] to select the prompt automatically. We observe dozens of prompt mistakes but do not modify them to avoid man-made influence.

5.2 Qualitative Comparison

Fig. 6 showcases a qualitative comparison of MOTIA against other methods. Outpainting a narrow video into a square format. MOTIA employs a mask ratio of 0.66, surpassing the 0.6 ratio utilized by other methods, and demonstrates superior performance even with this higher mask ratio. The SDM method only manages to blur the extremities of the video's background, egregiously overlooking the primary subject and resulting in the outpainting failure as previously highlighted in Fig. 5. Dehan's approach effectively outpaints the background but utterly fails to address the foreground, leading to notable distortions. In contrast, the M3DDM method adeptly handles both subject and background integration but is marred by considerable deviations in subject characteristics, such as incorrect brown coloration in the dog's fur across several frames. Our method stands out by achieving optimal results, ensuring a harmonious and consistent outpainting of both the foreground and background.

5.3 Quantitative Comparison

Table 1 summarizes the evaluation metrics of our method compared to other approaches. Our method achieves comparable results to the best method in PSNR. It shows significant improvements in video quality (SSIM), perceptual metric (LPIPS), and distribution similarity (FVD). Specifically, our SSIM, LPIPS,



Fig. 7: Visual examples of ablation study on the proposed input-specific adaptation.



Fig. 8: Visual examples of ablation study on pattern-aware outpainting.

and FVD metrics show improvements of 7.00%, 21.27%, and 4.57% respectively on the DAVIS dataset, and 4.43%, 6.85%, and 11.45% on the YouTube-VOS dataset compared to the best-performing method.

5.4 Ablation Study

Ablation study on input-specific adaptation. We conducted the ablation study on input-specific adaptation with the DAVIS dataset to verify its effectiveness, as shown in Fig. 7 and Table 2. "SD+T" represents the result of directly combining the temporal module with Stable Diffusion, which led to a complete outpainting failure. "SD+T+C" indicates the additional use of ControlNet, resulting in similarly poor outcomes. "Direct-tune" refers to the approach of directly fitting the original video without outpainting training; in this case, we observed a very noticeable color discrepancy between the outpainted and known areas. In contrast, our method achieved the best results, ensuring consistency in both the visual and temporal aspects. The metrics shown in Table 2 also support this observation, with MOTIA significantly outperforming the other baselines.

Ablation study on pattern-aware outpainting. Table 3 summarizes our ablation experiments for the pattern-aware outpainting part. We conducted extensive validation on the YouTube-VOS dataset. "Direct" refers to performing outpainting according to Eq. 3 directly after input-specific adaptation. "SA" denotes spatially-aware insertion, and "SA+NR" indicates the combined use of spatially-aware insertion and noise regret. The experimental results demonstrate that each of our components effectively enhances performance. Specifically, Combining both SA-Insertion and Noise regret, the PSNR, SSIM, LPIPS, and FVD metrics show improvements of 2.69%, 0.90%, 3.95%, and 11.32% respectively than directly applying Eq. 3. Fig. 8 presents the visual examples of ablation study on our proposed pattern-aware outpainting part. When removing NR, it

Table 2: Ablation study on input- specific adaptation.					Table 3: Ablation study on the proposed pattern-aware outpainting.					
Component	PSNR^{\uparrow}	$SSIM\uparrow$	LPIPS↓	$\mathrm{FVD}\downarrow$	Component	$PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	LPIPS↓	FVD ↓	
SD+T	15.59	0.6640	0.2979	672.7	Direct	19.72	0.7568	0.1798	66.52	
$_{\rm SD+T+C}$	16.81	0.6961	0.2338	515.4	SA	19.97	0.7608	0.1752	58.40	
Direct-tune	19.42	0.7375	0.1784	312.1	SA+NR	20.25	0.7636	0.1727	58.99	
MOTIA	20.36	0.7578	0.1595	286.3						

might fail to align the texture colors or produce unreasonable details (*e.g.*, arms in the middle of Fig. 8). When further removing SA, it could potentially generate unrealistic results caused by the overfitting to the target video (*e.g.*, the white collar on the left of Fig. 8). Note that even though the FVD degrades in a very slight manner, all the other metrics increase and we qualitatively find it to be helpful for improving results.

5.5 Discussions

Model and computation complexity. Model Complexity: The original model has 1.79 billion (including the auto-encoder and text encoder) parameters in total, while the added adapters contain 7.49 million parameters, leading to an increase of 0.42% in memory usage. Computation Complexity: We report the peak GPU VRAM and the time required for outpainting a target video from 512×512 to 512×1024 with 16 frames at two stages in Table 4. For longer videos, as described in Section 4.5, instead of processing the long video as a whole, we adapt only to short video clips sampled from the long video. This approach does not require additional time or GPU VRAM during the input-specific adaptation phase. Additionally, with temporal co-denoising [29], the GPU VRAM usage remains the same as that for short video during the pattern-aware outpainting phase, while the required time increases linearly with the video length.

User study. We conducted a user study between MOTIA and M3DDM, utilizing the DAVIS dataset with a horizontal mask of 0.66 as source videos. Preferences were collected from 10 volunteers, each evaluating 50 randomly selected sets of results based on visual quality (such as clarity, color fidelity, and texture detail) and realism (including motion consistency, object continuity, and integration with the background). Table 2 demonstrates that the outputs from MOTIA are preferred over those from M3DDM in both visual quality and realism.

Why MOTIA outperforms (Why previous methods fail). 1) Flexibility. Current video diffusion models are mostly trained with fixed resolution and length, lacking the ability to tackle videos with various aspect ratios and lengths. In contrast, the adaptation phase of MOTIA allows the model to better capture the size, length, and style distribution of the source video, greatly narrowing the gap between pretrained weights and the source video. 2) Ability for capturing intrinsic patterns from source video. A crucial point for successful outpainting is the predicted score of diffusion models should be well-compatible with the

Table 4: Computation complexity.

tween M3DDM and MOTIA. Phase VRAM↓ Time↓ Method Visual-Quality Realism Adaptation 12.70 GB 309 Seconds Outpainting 5.80 GB 58 Seconds M3DDM 42.8%27.4%MOTIA MOTIA (In total) 12.70 GB 367 Seconds 72.6%57.2%

Table 5: User study comparison be-

original known regions of the source video. To achieve this, the model should effectively extract useful information from the source video for denoising. For instance, M3DDM concatenates local frames of source video at the input layers and incorporates the global frames through the cross-attention mechanism after passing light encoders. However, the information might not be properly handled especially for out-domain inputs, thus leading to outpainting failure. Instead, by conducting input-specific adaptation on the source video, the model can effectively capture the data-specific patterns in the source videos through gradient. Through this, MOTIA better leverage the data-specific patterns of the source video and image/video generative prior for outpainting. We hope this work inspires following research to exploit more from the source video itself instead of purely relying on the generative prior from intensive training on videos.

6 Conclusion

We present MOTIA, relying on a combination of input-specific adaptation for capturing inner video patterns and pattern-aware outpainting to generalize these patterns for effective outpainting. Extensive experiments validate the effectiveness.

Limitations: MOTIA requires learning necessary patterns from the source video, it faces limitation for outpainting videos with limited information. We provide an example of failure case. In Fig. 9, the original source video just contains grass and a woman with a blue shirt. We find our method produces a woman with *very similar appearance* when outpainting, which might be regarded as a "copy" phenomenon.



Fig. 9: Example of failure case. Contents outside the red lines are outpainted.

Acknowledge:

This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)'s InnoHK, by Smart Traffic Fund PSRI/76/2311/PR, by RGC General Research Fund Project 14204021. Hongsheng Li is a PI of CPII under the InnoHK.

References

- Arora, R., Lee, Y.J.: Singan-gif: Learning a generative video model from a single gif. In: CVPR. pp. 1310–1319 (2021)
- Avrahami, O., Lischinski, D., Fried, O.: Blended diffusion for text-driven editing of natural images. In: CVPR. pp. 18208–18218 (2022)
- Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. ACM Trans. Graph. 28(3), 24 (2009)
- 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. pp. 22563–22575 (2023)
- Dehan, L., Van Ranst, W., Vandewalle, P., Goedemé, T.: Complete and temporally consistent video outpainting. In: CVPRW. pp. 687–695 (2022)
- Fan, F., Guo, C., Gong, L., Wang, B., Ge, T., Jiang, Y., Luo, C., Zhan, J.: Hierarchical masked 3d diffusion model for video outpainting. In: ACM MM. pp. 7890–7900 (2023)
- Gao, C., Saraf, A., Huang, J.B., Kopf, J.: Flow-edge guided video completion. In: ECCV. pp. 713–729. Springer (2020)
- Guo, Y., Yang, C., Rao, A., Wang, Y., Qiao, Y., Lin, D., Dai, B.: Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. arXiv preprint arXiv:2307.04725 (2023)
- 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 preprint arXiv:2210.02303 (2022)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS 33, 6840–6851 (2020)
- Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
- Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. arXiv:2204.03458 (2022)
- Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML. pp. 12888– 12900. PMLR (2022)
- Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: CVPR. pp. 11461–11471 (2022)
- Ma, Y., He, Y., Cun, X., Wang, X., Shan, Y., Li, X., Chen, Q.: Follow your pose: Pose-guided text-to-video generation using pose-free videos. arXiv preprint arXiv:2304.01186 (2023)
- Ma, Y., He, Y., Wang, H., Wang, A., Qi, C., Cai, C., Li, X., Li, Z., Shum, H.Y., Liu, W., et al.: Follow-your-click: Open-domain regional image animation via short prompts. arXiv preprint arXiv:2403.08268 (2024)
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
- Nikankin, Y., Haim, N., Irani, M.: Sinfusion: Training diffusion models on a single image or video. arXiv preprint arXiv:2211.11743 (2022)
- Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675 (2017)

- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
- Rout, L., Parulekar, A., Caramanis, C., Shakkottai, S.: A theoretical justification for image inpainting using denoising diffusion probabilistic models. arXiv preprint arXiv:2302.01217 (2023)
- Shaham, T.R., Dekel, T., Michaeli, T.: Singan: Learning a generative model from a single natural image. In: ICCV. pp. 4570–4580 (2019)
- Shen, D., Song, G., Xue, Z., Wang, F.Y., Liu, Y.: Rethinking the spatial inconsistency in classifier-free diffusion guidance. In: CVPR. pp. 9370–9379 (June 2024)
- Shi, X., Huang, Z., Wang, F.Y., Bian, W., Li, D., Zhang, Y., Zhang, M., Cheung, K.C., See, S., Qin, H., et al.: Motion-i2v: Consistent and controllable imageto-video generation with explicit motion modeling. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024)
- 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 textvideo data. arXiv preprint arXiv:2209.14792 (2022)
- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
- Unterthiner, T., Van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717 (2018)
- Voleti, V., Jolicoeur-Martineau, A., Pal, C.: Mcvd: Masked conditional video diffusion for prediction, generation, and interpolation. NeurIPS (2022)
- Wang, F.Y., Chen, W., Song, G., Ye, H.J., Liu, Y., Li, H.: Gen-l-video: Multi-text to long video generation via temporal co-denoising. arXiv preprint arXiv:2305.18264 (2023)
- Wang, F.Y., Huang, Z., Shi, X., Bian, W., Song, G., Liu, Y., Li, H.: Animatelcm: Accelerating the animation of personalized diffusion models and adapters with decoupled consistency learning. arXiv preprint arXiv:2402.00769 (2024)
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13(4), 600–612 (2004)
- Wu, J.Z., Ge, Y., Wang, X., Lei, W., Gu, Y., Hsu, W., Shan, Y., Qie, X., Shou, M.Z.: Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. arXiv preprint arXiv:2212.11565 (2022)
- Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: Youtube-vos: A large-scale video object segmentation benchmark. arXiv preprint arXiv:1809.03327 (2018)
- 34. Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X.: Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In: CVPR. pp. 1316–1324 (2018)
- 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. pp. 10459–10469 (2023)
- Zhang, L., Agrawala, M.: Adding conditional control to text-to-image diffusion models. arXiv preprint arXiv:2302.05543 (2023)
- 37. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)