

VipDiff: Towards Coherent and Diverse Video Inpainting via Training-free Denoising Diffusion Models

Chaohao Xie Kai Han* Kwan-Yee K. Wong*

The University of Hong Kong

viousxie@connect.hku.hk, kaihanx@hku.hk, kykwong@cs.hku.hk

Abstract

*Recent video inpainting methods have achieved encouraging improvements by leveraging optical flow to guide pixel propagation from reference frames, either in the image space or feature space. However, they would produce severe artifacts when the masked area is too large and no pixel correspondences could be found. Recently, denoising diffusion models have demonstrated impressive performance in generating diverse and high-quality images, and have been exploited in a number of works for image inpainting. These methods, however, cannot be applied directly to videos to produce temporal-coherent inpainting results. In this paper, we propose a training-free framework, named **VipDiff**, for conditioning diffusion model on the reverse diffusion process to produce temporal-coherent inpainting results without requiring any training data or fine-tuning the pre-trained models. **VipDiff** takes optical flow as guidance to extract valid pixels from reference frames to serve as constraints in optimizing the randomly sampled Gaussian noise, and uses the generated results for further pixel propagation and conditional generation. **VipDiff** also allows for generating diverse video inpainting results over different sampled noise. Experiments demonstrate that our **VipDiff** outperforms state-of-the-art methods in terms of both spatial-temporal coherence and fidelity.*

1. Introduction

Video inpainting aims to generate spatial-temporal coherent contents for the masked areas in corrupted video frames. Existing video inpainting methods can be broadly classified into 1) end-to-end synthesis methods and 2) flow-based pixel propagation methods.

End-to-end synthesis methods typically adopt 3D convolutions [3, 4, 12], attention modules [17, 18, 41], or multi-frame transformer-based networks [19, 20, 41, 47]. They take corrupted frames and their corresponding masks as in-

put, and train a video completion network to directly output the completed frames. Since they have fixed temporal receptive windows (i.e., only a fixed length of reference frames can be fed to the input), severe artifacts and temporal incoherence often occur when the input frames fail to provide useful texture and content hints for the frame completion models (see Fig. 1(b)).

Flow-based pixel propagation methods first conduct optical flow completion and then utilize the completed flows to guide a pixel propagation step, which warps valid pixels from reference frames to a target frame. An additional synthesis network is trained to fill the remaining masked pixels. Although these flow-based methods [8, 14, 39, 43] can achieve better temporal coherence over the generated frames, their results still show obvious artifacts in the mask center when the masks are large (see Fig. 1(c)). Undoubtedly, there is still a large room for improvement.

Recently, denoising diffusion models [11, 28] have attracted enormous attention due to their impressive performance in generating diverse and high-quality images from Gaussian noise through a series of denoising steps. There are several methods, like DPS [5] and ReSampling [31], which adopt pre-trained diffusion model for image inpainting via sampling posterior distributions. RePaint [21] introduces a time travel strategy to improve spatial coherence between masked and unmasked areas, and CoPaint [42] proposes optimizing the randomly sampled Gaussian noise using the unmasked pixels as constraints. Latent Diffusion Models (LDM) [27] map the input to a low-dimensional space with an encoder and feed the encoded features for conditional generation in the latent feature space. These diffusion-based methods work well for the image inpainting task. However, they cannot be applied directly to videos since they do not consider temporal coherence. For instance, we show the inpainting results of LDM on three frames randomly selected from two videos in Fig. 1(d). Even though the surrounding patterns are similar, LDM generates different contents in the masked area in different frames and fails to maintain temporal coherence.

Inspired by the above diffusion-based image inpainting

* Corresponding author



(a) Input (b) ProPainter [47] (c) ECFVI [14] (d) LDM [27] (e) Ours-Sample1 (f) Ours-Sample2 (g) Ours-Sample3

Figure 1. Video inpainting results on Davis dataset, we randomly select 3 frames from 'bear' and 'cows' videos. LDM generates different contents, ProPainter and ECFVI contain artifacts in the mask center, while our samples are more sharp and temporal-coherent.

methods, we propose a training-free framework for conditioning pre-trained image-level diffusion models to generate spatial-temporal coherent and diverse video inpainting results. Given a corrupted video, we first adopt a flow completion model to predict all the optical flows between frames. We then use the completed flows to propagate valid pixels from reference frames to a target frame. Next, we carry out the reverse diffusion process, where valid pixels (i.e., both pixels in the unmasked area and pixels propagated from other frames) in the target frame serve as constraints in optimizing the sampled noise through backpropagation.

After optimization, a spatial-temporal coherent frame is generated. Its pixels can be propagated to other frames again and the process is repeated for the next target frame. Note that once we generate one completed frame and propagate its pixels to other frames, the number of remaining invalid (masked) pixels in most frames, especially those nearby ones, would become zero. Since we only need to perform conditional generation on those frames with a non-zero size mask, our framework is therefore very time efficient. **VipDiff** can generate both spatial-temporal coherent results and diverse inpainted videos. In Fig. 1(e), (f) and (g), we show three different results on two videos. All of our generated results are coherent with the video motions.

VipDiff has the following advantages. 1) It eliminates the efforts in collecting large-scale video datasets and training a large video diffusion model, which may only be possible to be conducted by large companies or foundations. 2) It leverages the diverse generation capability of pre-trained

image-level diffusion models and generates results far better than SOTA methods, while still allows users to choose from diverse video inpainting results. 3) It just takes several minutes to inpaint a video sequence of about 100 to 200 frames on a single RTX 3090 GPU, demonstrating that it is both device friendly and time efficient. Hence our method enables users with limited GPU resources to utilize the ability of pre-trained diffusion models for video inpainting.

We summarize our contributions as follows:

- A training-free framework, named **VipDiff**, for conditioning pre-trained image-level diffusion model to generate spatial-temporal coherent video inpainting results. On large mask cases, it is capable of generating diverse results for users.
- **VipDiff** saves the extensive computation and time costs in specifically training a video denoising diffusion model for video inpainting task.
- To the best of our knowledge, we present the first work that successfully tames pre-trained image-level diffusion models for video inpainting task.
- Experiments on public datasets demonstrate superior performance of **VipDiff** over existing state-of-the-arts.

2. Related Works

In this section, we give a brief review of video inpainting methods, including recent methods based on diffusion

probabilistic model [28].

2.1. Patch-based Methods

Traditional inpainting methods [1] analyze local structures and fluid dynamics [2] along the mask boundary, and exploit spatial-temporal similarities [24, 36] between patches from undamaged areas to sample patches for filling the masked regions. To improve the performance on scenes with occluded objects, additional energy functions [9, 23, 25] and optical-flow constraints [13, 22] have been proposed to guide the optimization process. Although these methods can produce plausible results, their lack of understanding of global semantics results in poor visual quality for complex scenes. They also require substantial amount of time for computation, which limits their potentials in practical applications.

2.2. Deep Learning Methods

With the great progress in deep learning, deep neural networks for video inpainting have become popular, and large performance gains have been reported. To capture spatial-temporal relationships among masked frames, 3D convolutions have been widely adopted to combine temporal information and local spatial structures [3, 4, 12, 33]. To improve visual coherence over the generated sequences, optical flows have been utilized to guide the propagation of pixels from nearby frames into the masked areas as additional prior [8, 39, 43]. To obtain optical flows in the masked areas, an additional flow completion network is trained to complete the global flows. ECFVI [14] introduces an error correcting model for reducing color discrepancy when propagating pixels from far neighbors. Other methods [17, 18, 41] utilize different attention modules for aggregating valid features from different frames.

Recently, Transformer [32] and Vision Transformer [7] have shown great potential in video inpainting. STTN [41] introduces spatial-temporal transformer to combine multi-scale information across different frames and fill masked areas in all input frames simultaneously. FuseFormer [20] adopts overlapping patch splitting strategy for learning fine-grained features to enhance inpainting quality. FAST [40] combines information from the frequency domain. E2FGVI [19], FGT [44, 45] and ProPainter [47] incorporate optical flows in training the transformer structures. Although these transformer-based methods improve the overall visual quality, they still suffer from performance degradation when facing larger masks or when neighbouring frames fail to provide useful information.

2.3. Diffusion-based Methods

More recently, diffusion-based methods [11, 28, 29] have made significant progress in image generation. To generate coherent inpainted images, DPS [5] and ReSampling [31]

treat the image inpainting problem as a process of sampling images from the posterior distributions conditioned on the masked images. RePaint [21] adopts a time travel strategy which combines the current step t denoising result with the noise degraded image based on the mask and uses it to generate the step $t + 1$ image by a one-step forward process. CoPaint [42] improves the spatial coherence by taking the unmasked pixels as constraints and optimizing the noise z to match the generated image with the masked image. Latent Diffusion Models (LDM) [27] utilize an auto-encoder structure to project the masked image into a latent space with reduced resolution for conditional generation. These methods can produce diverse and high-quality inpainted images. However, they cannot be directly applied to video inpainting, since they only consider spatial coherence within an image but not temporal coherence among frames. Other methods try to train large video diffusion models for video sequences editing [10, 46], or combining language models with text prompts [37] to act as agents. Even trained with large models, they do not explicitly address the temporal inconsistent issues. Our proposed method combines the advantages of diffusion-based methods and flow-based methods, and generates spatial-temporal coherent and diverse results for video inpainting, which only requires on image-level pre-trained diffusion models.

3. Proposed Methods

In this section, we first give a brief overview of the Denoising Diffusion Probabilistic Models (DDPMs) and describe the notations used in this paper. We then introduce our training-free framework for conditioning pre-trained image-level diffusion models to generate spatial-temporal coherent video inpainting results. We also recommend readers to refer to the original papers [11, 28] of the diffusion models for details and derivations if needed.

Let $\mathcal{X} = \{x^k | k \in [1, N]\}$ denote a corrupted video with N frames, and $\mathcal{M} = \{m^k | k \in [1, N]\}$ be the corresponding per-frame masks with the same spatial dimensions as \mathcal{X} . A video inpainting method takes the image frames \mathcal{X} and the masks \mathcal{M} as input, examines the frame core, and outputs the recovered video $\hat{\mathcal{Y}} = \{\hat{y}^k | k \in [1, N]\}$, which should be spatial-temporal coherent.

3.1. Preliminaries

Denoising diffusion probabilistic models (DDPMs) [11] can be viewed as a type of generative models which are trained to approximate a data distribution $x_0 \sim p(x)$ through a series of intermediate variables $x_{1:T}$ that are degraded states of x_0 . These models comprise two processes, namely the *forward diffusion process* which defines how the original image x_0 becomes the noise image x_T , and the *reverse diffusion process* which recovers the clean image x_0 .

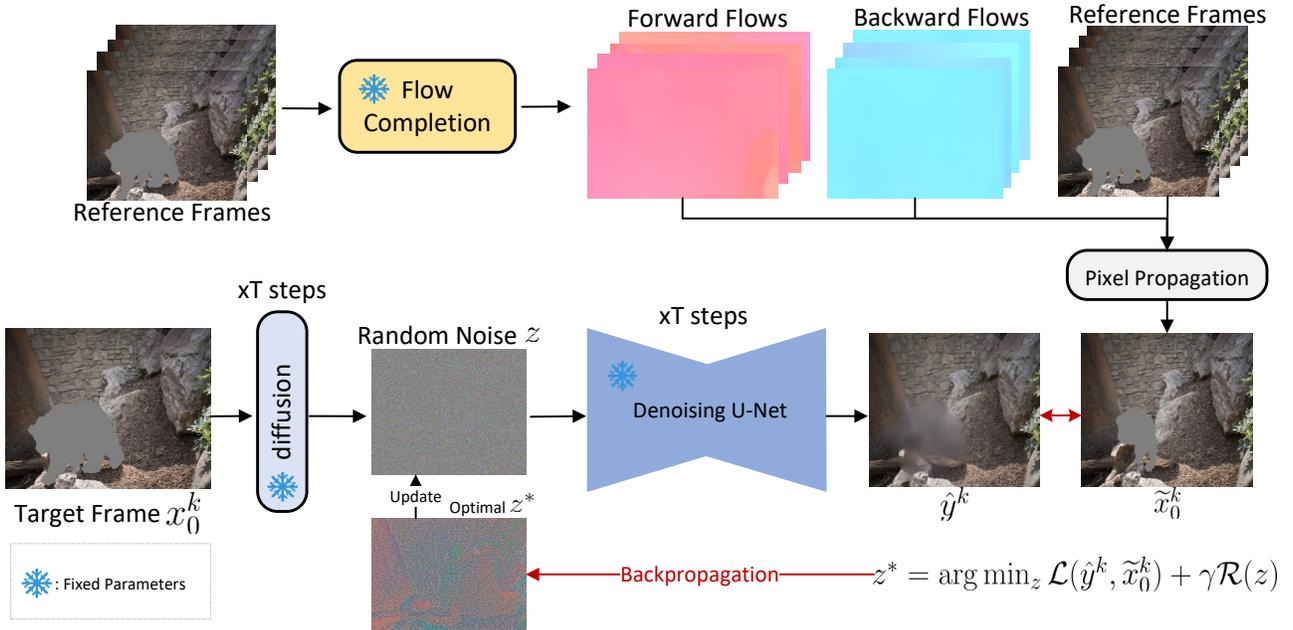


Figure 2. Overall framework of our **VipDiff**. Given a target frame x_0^k , we first adopt a flow completion model to predict the optical flows. Then the flows are utilized for pixel propagation to extract temporal prior from reference frames to get partially inpainted image \tilde{x}_0^k . Next, \tilde{x}_0^k will act as constrains for optimizing the random sampled Gaussian noise z which is feed for the reverse denoising U-Net. Through backpropagation, we optimize the noise z at each time step and finally find an optimal z^* for filling the target frame. Note that all the model parameters are frozen during the noise optimization process.

Given the k -th masked input frame $x_0^k = x^k * (1 - m^k)$ from the corrupted video \mathcal{X} , the forward diffusion process defines a Markov chain q of T steps which progressively adds Gaussian noise to the damaged image x_0^k by a variance schedule $\{\beta_t\}_{t=1}^T$, i.e.,

$$q(x_t^k | x_{t-1}^k) = \mathcal{N}(x_t^k; \sqrt{1 - \beta_t} x_{t-1}^k, \beta_t \mathbf{I}) \quad (1)$$

where each $\beta_t \in (0, 1)$. With a sufficiently large T , the final state $q(x_T^k | x_0^k)$ will be close to real Gaussian noise.

For the reverse diffusion process, a diffusion model with learnable parameters θ is adopted to estimate the amount of noise at each time step (x_t^k),

$$p_\theta(x_{t-1}^k | x_t^k) = \mathcal{N}(x_{t-1}^k; \mu_\theta(x_t^k, t), \Sigma_\theta(x_t^k, t)). \quad (2)$$

To perform unconditional generation from a sampled Gaussian noise x_T^k , a shared-weight denoising auto-encoder $\epsilon_\theta(x_t^k, t)$ is trained to predict the denoised version of x_t^k . The objective can be formulated as

$$\mathcal{L}_{DDPM} = \mathbb{E}_{x_0^k, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), t \sim \mathcal{U}(1, T)} [\|\epsilon - \epsilon_\theta(x_t^k, t)\|^2], \quad (3)$$

where $\mathcal{U}(1, T)$ is a uniform distribution on $\{1, 2, \dots, T\}$ and ϵ is the noise level at current time step.

Latent Diffusion Models (LDMs) [27] are variants of DDPMs which utilize an encoder \mathcal{E} to project the input image x^k into low-resolution latent representations and a

decoder \mathcal{D} to reconstruct the image. The forward diffusion process and reverse diffusion process are carried out in the latent space. LDMs allow for conditional generation by introducing additional prior in the reverse process $p_\theta(x_{t-1}^k | x_t^k, c)$, where c denotes the condition which can be texts, object segmentation mask, or other useful information. For inpainting, the condition c would be provided by the masked input frame x_0^k . The denoising network $\epsilon_\theta(x_t^k, t, c)$ can be trained under conditional reverse process.

Note that conditional LDMs only inpaint a single frame at a time without enforcing any temporal coherence between frames. Naïvely applying conditional LDMs to videos will inevitably lead to temporal incoherence in the generated sequence (see Fig. 1). In this paper, we propose a training-free framework that successfully tames image-level diffusion models for generating spatial-temporal coherent inpainting results, by utilizing optical flow to induce prior from pixels in nearby frames.

3.2. Optical Flow Guided Pixel Propagation

Since pre-trained models [6] or LDMs [27] have demonstrated strong capability in producing high-quality inpainting results from masked input images, we propose to directly make use of them in our method without modifying any of its network structure (attention modules, etc.) or fine-tuning its parameters. Unlike other methods which train a video diffusion model for specific video generation

Table 1. Quantitative comparison with state-of-the-art video inpainting methods on the YouTube-VOS and DAVIS datasets. \uparrow (\downarrow) indicates the higher (lower) the better. We highlight the best results in **bold**, and the second best in underline.

Methods	YouTube-VOS				DAVIS				Speed
	PSNR \uparrow	SSIM \uparrow	VFID \downarrow	$E_{warp} \downarrow$	PSNR \uparrow	SSIM \uparrow	VFID \downarrow	$E_{warp} \downarrow$	s/frame
LGTSM [4]	29.74	0.9504	0.070	0.1859	28.57	0.9409	0.170	0.1640	0.19
VINet [15]	29.20	0.9434	0.072	0.1490	28.96	0.9411	0.199	0.1785	-
DFVI [39]	29.16	0.9429	0.066	0.1509	28.81	0.9404	0.187	0.1608	1.96
CAP [17]	31.58	0.9607	0.071	0.1470	30.28	0.9521	0.182	0.1533	0.27
FGVC [8]	29.67	0.9403	0.064	0.1022	30.80	0.9497	0.165	0.1586	1.75
STTN [41]	32.34	0.9655	0.053	0.0907	30.67	0.9560	0.149	0.1449	0.06
FuseFormer [20]	33.29	0.9681	0.053	0.0900	32.54	0.9700	0.138	0.1362	0.15
E2FGVI [19]	33.71	0.9700	<u>0.046</u>	<u>0.0864</u>	33.01	0.9721	0.116	0.1315	0.13
FGT [44]	33.81	0.9711	0.066	0.0893	33.79	<u>0.9737</u>	0.106	0.1331	0.79
ECFVI [14]	33.77	0.9703	0.053	0.0913	33.52	0.9732	0.105	<u>0.1292</u>	1.34
ProPainter [47]	34.23	<u>0.9764</u>	0.054	0.0971	34.27	0.9731	<u>0.104</u>	0.1312	0.09
Ours	<u>34.21</u>	0.9773	0.041	0.0828	<u>34.23</u>	0.9745	0.102	0.1280	2.73

tasks, we optimize the randomly sampled Gaussian noise z to maximize the spatial-temporal coherence by leveraging both the unmasked pixels in the target frame and pixels propagated from reference frames as constraints.

Figure 2 shows the overall generation pipeline of our **VipDiff**, the key to making the inpainted frames temporal-coherent lies in the success of deriving proper prior for the masked areas from as many valid pixels from nearby frames as possible. To this end, we adopt a pre-trained RAFT model [30] as our flow estimator. Pre-trained flow completion RAFT is capable of predicting complete optical flows for images with missing areas. The completed flow from frame k to frame j is denoted as

$$\tilde{f}_{k \rightarrow j} = F(x_0^k, x_0^j, m^k, m^j), \quad (4)$$

where $F(\cdot)$ is the flow estimator. After obtaining $\tilde{f}_{k \rightarrow j}$, we can propagate valid pixels from frame j (referred to as the *reference frame*) to frame k (referred to as the *target frame*) to provide temporal constraints for the diffusion model.

3.2.1 Pixel Propagation

Let x_0^j be a source frame of x_0^k that contains valid pixel information for the masked area of x_0^k . We define a backward warping function $\omega(\cdot)$ to propagate valid pixels from the source frame x_0^j to the masked area of the target frame x_0^k based on the completed flow $\tilde{f}_{k \rightarrow j}$. Written as

$$\begin{aligned} \tilde{x}_0^k &= x_0^k + m^{j \rightarrow k} \odot \omega(x_0^j, \tilde{f}_{k \rightarrow j}), \\ m^{j \rightarrow k} &= m^k \odot (1 - \omega(m^j, \tilde{f}_{k \rightarrow j})), \end{aligned} \quad (5)$$

where $m^{j \rightarrow k}$ denotes a mask for the propagated valid pixels falling within the masked area of x_0^k . After pixel propagation, we update the invalid mask of x_0^k to $\tilde{m}^k = m^k -$

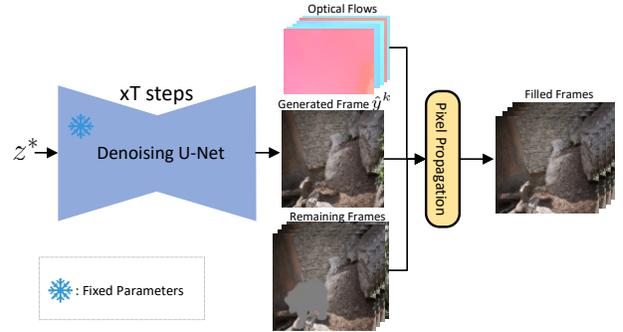


Figure 3. Process for propagating constrained generated pixels to other frames.

$m^{j \rightarrow k}$. We repeat the above pixel propagation step for different j , starting from the nearest to the furthest neighboring frame, until the invalid mask becomes all zeros or all neighboring frames have been exhausted. As pointed out in [14], directly propagating pixels from other frames may result in color discrepancy due to misalignment or brightness inconsistency. The authors introduced an error compensation network to reduce this problem. We adopt the same error compensation network in our work to reduce the color discrepancy in pixel propagation.

3.3. Noise Optimized Reverse Diffusion Process

After applying the above optical flow guided pixel propagation, we obtain a (partially) filled frame \tilde{x}_0^k which contains fewer invalid pixels compared with x_0^k . We can then proceed to generate an inpainted frame \hat{y}^k for \tilde{x}_0^k based on the reverse diffusion process. The goal is to sample a Gaussian noise z that maximizes the similarity between the generated image \hat{y}^k and \tilde{x}_0^k . This can be achieved by optimizing the noise z through backpropagation with the following objective

$$\mathcal{L}_{cond} = \|\hat{y}^k \odot (1 - \tilde{m}^k) - \tilde{x}_0^k \odot (1 - \tilde{m}^k)\|^2. \quad (6)$$

Specifically, $\hat{y}^k = \epsilon_\theta^d(z, t, x_0^k)$ with ϵ_θ^d being the pre-trained latent diffusion model. In general, our method can work for unconditional diffusion models with $\hat{y}^k = \epsilon_\theta^d(t, x_t^k)$. Following [42], we also add a regularization term $\mathcal{R}(z) = \|z - z_0\|^2$ to constrain z with the original sampled noise z_0 to follow a Gaussian distribution. Hence, the overall optimization objective can be written as

$$z^* = \arg \min_z \mathcal{L}_{cond}(\hat{y}^k, \tilde{x}_0^k) + \gamma \mathcal{R}(z), \quad (7)$$

where γ is the weight of the regularization term.

With the optimal z^* , we can generate a spatial-temporal coherent inpainting result for the target frame. The spatial coherence can be achieved by the valid pixels existing in the target frame, and the temporal consistency comes from the constraints of propagated pixels. The pre-trained diffusion models provide overall coherence for the unconstrained pixels. Note that we do not tune any parameters of the pre-trained diffusion model but only optimize z throughout the entire process, making our method device friendly and time efficient. The generation framework is shown in Fig. 2.

3.3.1 Inpaint the Entire Sequence

After obtaining the inpainted frame \hat{y}^k , we update its invalid mask to all zeros. We then proceed to the next frame and, likewise, carry out optical flow guided pixel propagation. Note that once we have a completed frame with no invalid pixels, in most cases, many subsequent frames can already be completed by pixel propagation alone (see Fig. 3). We only need to carry out the noise optimized reverse diffusion process for those frames which remain incomplete after pixel propagation (the number of such frames is usually relatively small compared to the video length). The whole process is repeated for each frame in the sequence until all frames have been completed.

4. Experiments

4.1. Experimental Settings

4.1.1 Dataset

Following common practice, we adopt two widely used video object segmentation datasets, namely YouTube-VOS [38] and DAVIS [26], for evaluation. YouTube-VOS comprises 3471, 474, and 508 video clips for training, validation, and testing respectively. DAVIS consists of 150 video clips, with 60 for training, 30 for validation and 90 for testing. Since we do not need any training data, we follow the same setting as FuseFormer [20] and E2FGVI [19] to use test sets from YouTube-VOS and DAVIS for quantitative and qualitative evaluations. For the testing masks, we

adopt the same stationary masks as E2FGVI [19] for computing objective metrics and object shape masks for qualitative comparisons. Our code will be made publicly available at: <https://vivos.github.io/projects/VipDiff>.

4.1.2 Quantitative Metrics

For a fair comparison, we select PSNR, SSIM [35], VFID [34], and flow warping error E_{warp} [16] to evaluate the performance of our method quantitatively with other state-of-the-art video inpainting methods.

4.1.3 Implementation Details

We adopt the pre-trained image-level diffusion model from LDMs [27], and the number of reverse optimization steps is set to 50, with an adaptive learning rate initialized at $\eta_0 = 0.01$, and $\eta_t = 0.9\eta_{t-1}$ for optimizing the noise. We set the weight for the regularization term to $\gamma = 0.001$. We adopt the pre-trained RAFT provided by [14] as our flow completion model, and also utilize their error compensation model for reducing color discrepancy in pixel propagation. The whole generating process is performed on a single GTX 3090 GPU. For a corrupted video with 100 frames, our method takes about 4 to 6 minutes to complete a video, which is also efficient since our **VipDiff** does not require any training time and training data.

To select the first frame x_0^k for constrained generation, we warp the masks of each corrupted frames using the completed flows, and then select the frame which holds the most overlapping mask region as the starting frame.

4.2. Main Comparisons

4.2.1 Quantitative Results

We compute the quantitative results on the YouTube-VOS and DAVIS datasets, and compare with existing state-of-the-art methods, including LGTSM [4], ViNet [15], DFVI [39], CAP [17], STTN [41], FGVC [8], FuseFormer [20], E2FGVI [19], ProPainter [47], FGT [44] and ECFVI [14]. Following E2FGVI and ECFVI, we test all the frames at the resolution of 432×240 . As shown in Tab. 1, although our **VipDiff** does not require any training on those datasets, it achieved the most highest metric scores compared with other SOTA methods. As for the video related metrics, we also achieved the highest E_{warp} and VFID both on YouTube-VOS and on DAVIS datasets. This demonstrates the superior performance of **VipDiff**.

Inference time. We compared the running time per image on the DAVIS validation dataset using a single RTX 3090 GPU and Intel 4214R CPU, as shown in Table 1 (last column). On average, our **VipDiff** is only about 1.4s slower than optical-flow guided methods like ECFVI, but providing better visual quality. Diffusion models themselves are time-consuming (LDM takes about 4.3s just for inference), our framework can do object removal for a video sequence

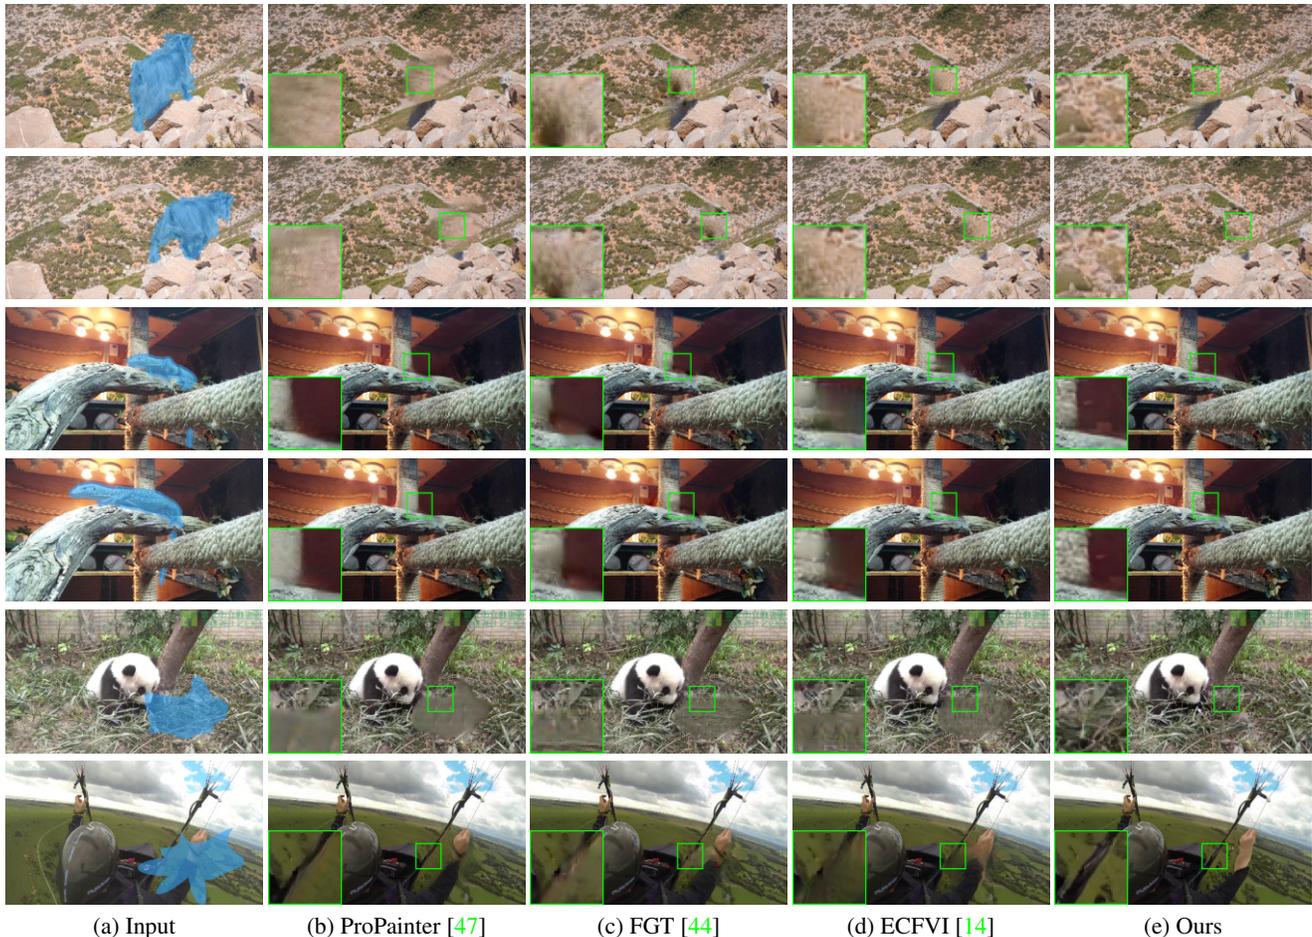


Figure 4. Qualitative comparisons with SOTA video inpainting methods. Best viewed in PDF with zoom. Please refer to the supplementary video for a comprehensive comparison.

about 100 frames in about 4 minutes, which is time-efficient for video inpainting. For inpainting task, users focus more on visual quality, though feed-forward models like E2FGVI or Propainter takes less time, they produce severe artifacts. With the faster modern GPU techniques and better acceleration algorithms of diffusion models in future, our **VipDiff** is possible to close the speed gap to those feed-forward models, while providing sharper and high-detailed results.

4.2.2 Qualitative Results

For qualitative comparisons, we select three recent SOTA methods, namely ProPainter [47], FGT [44], and ECFVI [14], to conduct visual comparisons. We select two videos with object-shaped masks (top 4 rows in Fig. 4) and two videos with stationery masks (bottom 2 rows in Fig. 4). One can see that for object-removal task (top 4 rows), our **VipDiff** can generate both sharp and temporal-coherent results, while the competing methods fail to generate temporal-coherent and sharp contents in the mask center. As for the stationery masks (bottom 2 rows), our **VipDiff**

largely surpasses the SOTA video inpainting methods since they cannot find useful contents in the reference frames, while our training-free framework can generate meaningful details in the masked region. We refer readers to our video results in the supplementary material for a better evaluation of the temporal-coherent performance of our **VipDiff**.

4.3. Ablation Study

We conduct an ablation study on the DAVIS dataset with the following variants: (i) direct output of LDM [27] inpainting model for all corrupted frames, (ii) iteratively select a corrupted frame to inpaint using LDM and use completed flows to propagate pixels to other neighbouring frames, until there is no invalid mask, (iii) without noise optimization in the reverse diffusion process (i.e., use optical flow to guide pixel propagation first and then randomly sample a fixed Gaussian noise for inpainting all the corrupted frames). We denote these three variants as ‘LDM’, ‘LDM+PP’, and ‘w/o Opt’ respectively.

We show the qualitative results of the ablation study in



Figure 5. Ablation study on different variants. From top row to the bottom, they are (a) input, (b) output of LDM, (c) LDM combined pixel propagation process, (d) our framework without reverse noise optimization, and (e) ours.

Fig. 5, with three consecutive frames randomly selected from the ‘bear’ video. One can clearly see that when directly applying LDM inpainting model without considering any temporal prior, though the generated results are sharp for every frame, they are not temporal-coherent (see Fig. 5(b)). For LDM+PP, the inside inpainted contents may be temporal-coherent, but directly warping large patches yield severe color discrepancy issues due to the brightness of different frames. From Fig. 5(d), one can see that removing the noise optimization step and using fixed Gaussian noise for all corrupted frames would also generate different contents inside the mask center. Our **VipDiff**, on the other hand, can generate both sharp and temporal-coherent contents, demonstrating its effectiveness.

We show the quantitative results of the ablation study in Tab. 2. One can note that our full framework yielded the highest quantitative performance over other variants. LDM achieved the worst results in all metrics since it does not consider any temporal coherence. LDM+PP improved over LDM by adopting optical flow for pixel propagation, but it faces color discrepancy and structure incoherent issues. Removing the noise optimization step (w/o Opt) largely decreased the video-related quantitative performance. Furthermore, we conducted additional experiments by performing noise optimized reverse diffusion process for every single frame. However, we observed that this approach was considerably slow and resulted in significant frame flickering issues. Hence, we omit this result in comparison.

Ablation on $\gamma R(z)$. We conducted experiments on differ-

Table 2. Ablation studies on different variants.

Variant	PSNR \uparrow	SSIM \uparrow	VFID \downarrow	$E_{warp}\downarrow$
LDM	32.64	0.9673	0.206	0.1833
LDM+PP	33.14	0.9701	0.133	0.1379
w/o Opt	32.80	0.9698	0.187	0.1620
Ours	34.23	0.9745	0.102	0.1280

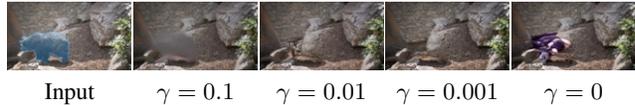


Figure 6. Ablation study on regularization weight γ .

ent weights of γ , ranging from 0.1 to 0, as illustrated in Fig 6. As the weight increases, we observed the blurriness and spatial distortions. This can be attributed to the heavier burden imposed by large Gaussian constraints on the later reverse diffusion steps, resulting in a higher magnitude of noise level. On the other hand, setting $\gamma = 0$, implying the absence of Gaussian constraints, may lead to the occurrence of artifacts. This is likely because the reconstruction loss induces rapid changes in the noise distribution during the early denoising steps, where it should still resemble Gaussian noise. Based on our experiments, we recommend selecting a slightly small γ ranging from 0.01 to 0.001 (avoiding values that are too large or too small, such as 0) to achieve satisfactory results.

We refer readers to our supplementary material for visual comparisons in videos and more ablation study, which more clearly delineate the strength of our method. Further, our training-free framework possesses strong generalization capability. It can be applied to any pre-trained image-level unconditional denoising diffusion models for video inpainting task. We provide video results by adopting other diffusion models in the suppl.

5. Conclusion

In this paper, we have proposed the first training-free framework, named **VipDiff**, that effectively tames a pre-trained image-level diffusion model for the video inpainting task. By introducing an optical flow guided reverse noise optimization process, our framework successfully generates sharp and temporally coherent video completion results. Our **VipDiff** further allows for providing diverse video outputs, and it also saves significant amount of efforts in collecting extensive video data for training a video inpainting diffusion model. Experiments have shown that our method achieves the state-of-the-art performance on benchmark datasets, and generates superior video completion results than other competing methods.

Acknowledgement This work is partially supported by the Hong Kong Research Grants Council - General Research Fund (Grant No.: 17211024) and HKU Seed Fund for Basic Research.

References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (TOG)*, 2009. 3
- [2] Marcelo Bertalmio, Andrea L Bertozzi, and Guillermo Sapiro. Navier-stokes, fluid dynamics, and image and video inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2001. 3
- [3] Ya-Liang Chang, Zhe Yu Liu, Kuan-Ying Lee, and Winston Hsu. Free-form video inpainting with 3d gated convolution and temporal patchgan. In *IEEE/CVF International Conference on Computer Vision*, 2019. 1, 3
- [4] Ya-Liang Chang, Zhe Yu Liu, Kuan-Ying Lee, and Winston Hsu. Learnable gated temporal shift module for deep video inpainting. In *British Machine Vision Conference*, 2019. 1, 3, 5, 6
- [5] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023. 1, 3
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 4
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 3
- [8] Chen Gao, Ayush Saraf, Jia-Bin Huang, and Johannes Kopf. Flow-edge guided video completion. In *European Conference on Computer Vision*, 2020. 1, 3, 5, 6
- [9] Miguel Granados, James Tompkin, K Kim, Oliver Grau, Jan Kautz, and Christian Theobalt. How not to be seen—object removal from videos of crowded scenes. In *Computer Graphics Forum*, 2012. 3
- [10] Bohai Gu, Yongsheng Yu, Heng Fan, and Libo Zhang. Flow-guided diffusion for video inpainting. *arXiv preprint arXiv:2311.15368*, 2023. 3
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 2020. 1, 3
- [12] Yuan-Ting Hu, Heng Wang, Nicolas Ballas, Kristen Grauman, and Alexander G Schwing. Proposal-based video completion. In *European Conference on Computer Vision*, 2020. 1, 3
- [13] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Temporally coherent completion of dynamic video. *ACM Transactions on Graphics (ToG)*, 2016. 3
- [14] Jaeyeon Kang, Seoung Wug Oh, and Seon Joo Kim. Error compensation framework for flow-guided video inpainting. In *European Conference on Computer Vision*, 2022. 1, 2, 3, 5, 6, 7
- [15] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Deep video inpainting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 5, 6
- [16] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning blind video temporal consistency. In *European conference on computer vision*, 2018. 6
- [17] Sungho Lee, Seoung Wug Oh, DaeYeun Won, and Seon Joo Kim. Copy-and-paste networks for deep video inpainting. In *IEEE/CVF international conference on computer vision*, 2019. 1, 3, 5, 6
- [18] Ang Li, Shanshan Zhao, Xingjun Ma, Mingming Gong, Jianzhong Qi, Rui Zhang, Dacheng Tao, and Ramamohanarao Kotagiri. Short-term and long-term context aggregation network for video inpainting. In *European Conference on Computer Vision*, 2020. 1, 3
- [19] Zhen Li, Cheng-Ze Lu, Jianhua Qin, Chun-Le Guo, and Ming-Ming Cheng. Towards an end-to-end framework for flow-guided video inpainting. In *IEEE/CVF conference on computer vision and pattern recognition*, 2022. 1, 3, 5, 6
- [20] Rui Liu, Hanming Deng, Yangyi Huang, Xiaoyu Shi, Lewei Lu, Wenxiu Sun, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fuseformer: Fusing fine-grained information in transformers for video inpainting. In *IEEE/CVF international conference on computer vision*, 2021. 1, 3, 5, 6
- [21] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1, 3
- [22] Yasuyuki Matsushita, Eyal Ofek, Weina Ge, Xiaou Tang, and Heung-Yeung Shum. Full-frame video stabilization with motion inpainting. *IEEE Transactions on pattern analysis and Machine Intelligence*, 2006. 3
- [23] Alasdair Newson, Andrés Almansa, Matthieu Fradet, Yann Gousseau, and Patrick Pérez. Video inpainting of complex scenes. *Siam journal on imaging sciences*, 2014. 3
- [24] Kedar A Patwardhan, Guillermo Sapiro, and Marcelo Bertalmio. Video inpainting of occluding and occluded objects. In *IEEE International Conference on Image Processing 2005*, 2005. 3
- [25] Kedar A Patwardhan, Guillermo Sapiro, and Marcelo Bertalmio. Video inpainting under constrained camera motion. *IEEE Transactions on Image Processing*, 2007. 3
- [26] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *IEEE conference on computer vision and pattern recognition*, 2016. 6
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF conference on computer vision and pattern recognition*, 2022. 1, 2, 3, 4, 6, 7
- [28] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using

- nonequilibrium thermodynamics. In *International conference on machine learning*, 2015. 1, 3
- [29] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 3
- [30] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European Conference on Computer Vision*, 2020. 5
- [31] Brian L. Trippe, Jason Yim, Doug Tischer, David Baker, Tamara Broderick, Regina Barzilay, and Tommi S. Jaakkola. Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. In *International Conference on Learning Representations*, 2023. 1, 3
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 2017. 3
- [33] Chuan Wang, Haibin Huang, Xiaoguang Han, and Jue Wang. Video inpainting by jointly learning temporal structure and spatial details. In *AAAI Conference on Artificial Intelligence*, 2019. 3
- [34] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, 2018. 6
- [35] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 2004. 6
- [36] Yonatan Wexler, Eli Shechtman, and Michal Irani. Space-time completion of video. *IEEE Transactions on pattern analysis and machine intelligence*, 2007. 3
- [37] Jianzong Wu, Xiangtai Li, Chenyang Si, Shangchen Zhou, Jingkang Yang, Jiangning Zhang, Yining Li, Kai Chen, Yunhai Tong, Ziwei Liu, et al. Towards language-driven video inpainting via multimodal large language models. In *IEEE conference on computer vision and pattern recognition*, 2024. 3
- [38] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *European conference on computer vision*, 2018. 6
- [39] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 1, 3, 5, 6
- [40] Bingyao Yu, Wanhua Li, Xiu Li, Jiwen Lu, and Jie Zhou. Frequency-aware spatiotemporal transformers for video inpainting detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 3
- [41] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *European Conference on Computer Vision*, 2020. 1, 3, 5, 6
- [42] Guanhua Zhang, Jiabao Ji, Yang Zhang, Mo Yu, Tommi Jaakkola, and Shiyu Chang. Towards coherent image inpainting using denoising diffusion implicit models, 2023. 1, 3, 6
- [43] Haotian Zhang, Long Mai, Ning Xu, Zhaowen Wang, John Collomosse, and Hailin Jin. An internal learning approach to video inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 1, 3
- [44] Kaidong Zhang, Jingjing Fu, and Dong Liu. Flow-guided transformer for video inpainting. In *European Conference on Computer Vision*, 2022. 3, 5, 6, 7
- [45] Kaidong Zhang, Jialun Peng, Jingjing Fu, and Dong Liu. Exploiting optical flow guidance for transformer-based video inpainting. *arXiv preprint arXiv:2301.10048*, 2023. 3
- [46] Zhixing Zhang, Bichen Wu, Xiaoyan Wang, Yaqiao Luo, Luxin Zhang, Yinan Zhao, Peter Vajda, Dimitris Metaxas, and Licheng Yu. Avid: Any-length video inpainting with diffusion model. *arXiv preprint arXiv:2312.03816*, 2023. 3
- [47] Shangchen Zhou, Chongyi Li, Kelvin C.K Chan, and Chen Change Loy. ProPainter: Improving propagation and transformer for video inpainting. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 3, 5, 6, 7