

TDM: Temporally-Consistent Diffusion Model for All-in-One Real-World Video Restoration

Yizhou Li¹[0000–0002–7122–2087], Zihua Liu¹[0000–0002–3908–0517], Yusuke Monno¹[0000–0001–6733–3406], and Masatoshi Okutomi¹[0000–0001–5787–0742]

Institute of Science Tokyo, Tokyo, Japan
{yli, zliu, ymonno}@ok.sci.titech.ac.jp,
mxo@ctrl.titech.ac.jp

Abstract. In this paper, we propose the first diffusion-based all-in-one video restoration method that utilizes the power of a pre-trained Stable Diffusion and a fine-tuned ControlNet. Our method can restore various types of video degradation with a single unified model, overcoming the limitation of standard methods that require specific models for each restoration task. Our contributions include an efficient training strategy with Task Prompt Guidance (TPG) for diverse restoration tasks, an inference strategy that combines Denoising Diffusion Implicit Models (DDIM) inversion with a novel Sliding Window Cross-Frame Attention (SW-CFA) mechanism for enhanced content preservation and temporal consistency, and a scalable pipeline that makes our method all-in-one to adapt to different video restoration tasks. Through extensive experiments on five video restoration tasks, we demonstrate the superiority of our method in generalization capability to real-world videos and temporal consistency preservation over existing state-of-the-art methods. Our method advances the video restoration task by providing a unified solution that enhances video quality across multiple applications.

Keywords: Multi-task Video Restoration · Diffusion Models · ControlNet.

1 Introduction

Video restoration is crucial as videos often lose quality due to factors like adverse weather, noise, compression, and limited sensor resolution, which can significantly hamper computer vision tasks such as object detection and video surveillance. Existing video restoration methods [4, 5, 9, 17, 18, 30, 34, 35] have shown progress but are limited to specific degradations, requiring separate models for each restoration task. This is costly and impractical for real-world applications where multiple degradations may occur. A unified all-in-one video restoration model is highly demanded for practical use. Recent attempts to create a single model for all-in-one image restoration have been promising [16, 25], but they fall short in video applications due to the challenge of maintaining temporal consistency. The main demands for video restoration are (1) developing a single model that can handle various degradations, (2) ensuring temporal consistency of restored videos, and (3) achieving robust real-world performance.

In recent years, diffusion models have markedly advanced image generation [27] and video generation [14, 32, 38], becoming potential keystones in vision-based AI. Ex-

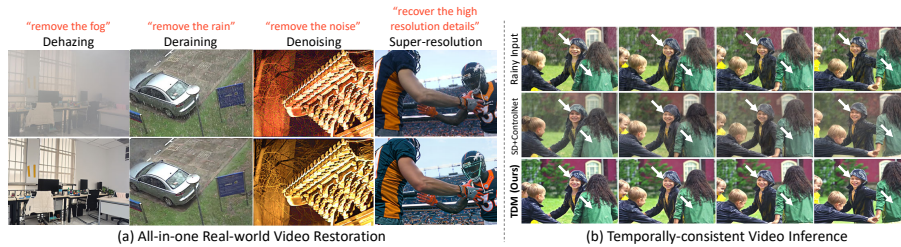


Fig. 1. Our Temporally-consistent Diffusion Model (TDM) has two main features: (a) Our model is all-in-one and can restore various real-world video degradation with a single diffusion model under the guidance of task prompts. (b) Our model can generate temporally consistent video frames with better preservation of original contents included in the input video.

tensive research has demonstrated their ability to parse and encode diverse visual representations from massive text-to-image datasets, enriching downstream applications with strong real-world generalization for tasks such as image editing [3, 22] and classification [6]. However, the random nature of diffusion models often disrupts the preservation of original image contents, presenting challenges for image restoration tasks. Some works have developed a diffusion model for each specific task [26, 28] by training it from scratch on a certain dataset, limiting its utility to a single task. Alternatively, approaches like Denoising Diffusion Implicit Models Inversion (DDIM Inversion) [8, 22] and a content-preserving module with a task-wise plugin [21] guide the diffusion process using input images for multi-target image editing and restoration. However, when these methods are directly applied to video restoration through image-by-image inference, they show insufficient temporal consistency.

To address the above issues, we propose the first diffusion-based all-in-one video restoration method, named Temporally-consistent Diffusion Model (TDM). TDM utilizes the capabilities of the pre-trained text-to-image Stable Diffusion (SD) model [27] alongside a fine-tuned ControlNet [36]. TDM is designed to offer a simple, efficient, and easy-to-extend framework for video restoration, focusing on improving training and inference strategies for a diffusion model without introducing complex modules.

For the training phase, we fine-tune a single-image-based ControlNet [36] on multiple tasks using text prompts referring to the task name, such as "remove the noise" for video denoising and "remove the rain" for video deraining. We refer to this as Task Prompt Guidance (TPG), which directs the diffusion process toward a specific restoration task. This approach leverages the strong zero-shot classification capability of the pre-trained SD model [6]. This simple strategy enables robust all-in-one restoration on multiple tasks without additional computation time and parameters, as demonstrated in Fig. 1(a). By focusing on fine-tuning ControlNet rather than training an entire diffusion model, we harness the generative power of the pre-trained SD model. This approach improves the quality of video restoration and enhances the robustness against real-world data. Furthermore, because ControlNet is fine-tuned using single-image inputs, our strategy eliminates the extreme memory requirements needed to process mul-

multiple video frames together during training, which is common in earlier methods [2, 10, 31] for ensuring temporal consistency. This allows our model to be trained on a single GPU using common single-image restoration datasets, simplifying data preparation and enhancing accessibility for future research.

In the inference phase, we aim to ensure content preservation and temporal consistency. Previous text-to-video studies [14, 38] addressed temporal consistency by replacing the self-attention in U-Net with cross-frame attention during inference. However, the method [14] struggles with large motions between video frames and the method [38] requires extensive memory and computational resources. Also, both methods do not consider content preservation. To address these issues, we propose a Sliding Window Cross-Frame Attention (SW-CFA) mechanism, combined with DDIM Inversion [8, 22]. This combination effectively tackles both content preservation and temporal consistency. Specifically, SW-CFA extends the concept of reference frames to a sliding window around the current frame, employing mean-based temporal smoothing in the attention calculation with minimal computational increase. This allows SW-CFA to handle large motions better than [14], achieving effective zero-shot image-to-video adaptation by adjusting the attention mechanism during inference. As discussed in [21], adding random Gaussian noise to the input latent can reduce the fidelity of diffusion-generated images. Therefore, we introduce DDIM Inversion for better content preservation. The combination of SW-CFA and DDIM Inversion maintains a deterministic and close distribution of input noises across adjacent frames, resulting in a more coherent video output. As illustrated in Fig. 1(b), compared with a straightforward combination of SD and ControlNet, our proposed TDM with SW-CFA and DDIM Inversion generates more temporally consistent video frames.

In summary, we build the first all-in-one video restoration diffusion model in this work. Our main contributions are as follows:

- Training strategies: We achieve efficient training by fine-tuning a single ControlNet using single-image inputs without requiring video inputs. The proposed TPG offers a simple way to achieve cross-task robustness.
- Inference strategies: We incorporate DDIM Inversion and introduce a novel SW-CFA mechanism for zero-shot video inference. This combination ensures accurate content preservation and robust temporal consistency in restored videos.
- Scalability: Our proposed TDM can be trained with a single GPU and is adaptable for video inference after being trained on single-image restoration datasets. This flexibility allows straightforward expansion to other video restoration tasks using only single-image datasets.
- Generalization performance: We conduct extensive experiments on five restoration tasks, demonstrating strong generalization performance of our TDM to real-world data over existing regression-based and diffusion-based methods.

2 Methodology

In this section, we first introduce the preliminary on latent diffusion models in Sec. 2.1. Then, we introduce the details of our method in Sec 2.2 (Training) and Sec 2.3 (Inference).

2.1 Preliminary

Latent Diffusion Model (LDM) [27] significantly advances traditional diffusion models by operating in a latent space. LDM employs an encoding mechanism to transform an image x into a compressed latent representation $z = E(x)$, thereby facilitating the learning of the latent codes' distribution, denoted as $z_0 \sim p_{\text{data}}(z_0)$ where p_{data} denotes the latent distributions of training images, following the Denoising Diffusion Probabilistic Model (DDPM) strategy introduced by [12]. First, LDM follows a forward phase that introduces Gaussian noise progressively over time steps t to derive z_t :

$$q(z_t|z_{t-1}) = \mathcal{N}(z_t; \sqrt{1 - \beta_t}z_{t-1}, \beta_t I), \quad (1)$$

with the noise scale represented by $\{\beta_t\}_{t=1}^T$, the gaussian distribution by \mathcal{N} , an all-one tensor with the same shape as z_t expressed by I , and the total diffusion steps denoted by T . Then, LDM follows a backward phase, where the model strives to reconstruct the preceding less noisy state z_{t-1} as follows:

$$p_\theta(z_{t-1}|z_t) = \mathcal{N}(z_{t-1}; \mu_\theta(z_t, t, \tau), \Sigma_\theta(z_t, t, \tau)), \quad (2)$$

where μ_θ and Σ_θ are mean and variance of the current state characterized by learnable parameters θ , and implemented with a noise prediction model ϵ_θ . To generate novel samples, initialization starts with a Gaussian sample $z_T \sim \mathcal{N}(0, 1)$, followed by a DDIM backward process of z_{t-1} for preceding time steps:

$$z_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{z_t - \sqrt{1 - \alpha_t} \epsilon_\theta(z_t, t, \tau)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \cdot \epsilon_\theta(z_t, t, \tau), \quad (3)$$

where the cumulative product $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$ is used for simplicity to indicate the transition towards z_0 at step t . The most well-known latent diffusion model is SD [27], which exemplifies text-to-image LDMs trained on a vast corpus of image-text pairs, with τ representing the text prompt. Moreover, it demonstrates its scalability across various other tasks [20, 19, 13].

ControlNet [36] enhances the capabilities of SD for more precise text-to-image synthesis by incorporating additional inputs such as depth maps, body poses, and images. While retaining the U-Net architecture identical to that of SD, ControlNet is fine-tuned to accommodate specific conditional inputs, modifying the function from $\epsilon_\theta(z_t, t, \tau)$ to $\epsilon_\theta(z_t, t, c, \tau)$, where c represents these additional conditional inputs.

DDIM Inversion [8, 29] is an inverse process of the DDIM forward process, based on the assumption that the ordinary differential equation can be reversed in the limit of small steps:

$$z_{t+1} = \sqrt{\alpha_{t+1}} \left(\frac{z_t - \sqrt{1 - \alpha_t} \epsilon_\theta(z_t, t, \tau)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t+1}} \cdot \epsilon_\theta(z_t, t, \tau). \quad (4)$$

In other words, the diffusion process is performed in the reverse direction, that is z_0 to z_T instead of z_T to z_0 , where z_0 in our case is set to be the encoding of the given input degraded image. Despite using ControlNet, we observe that employing noises generated through DDIM Inversion, rather than relying on random Gaussian noises, provides stable structural guidance that better preserves the contents of the input image.

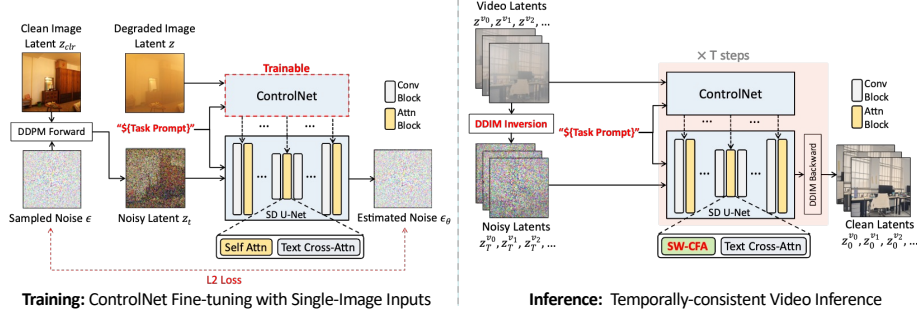


Fig. 2. Overall architecture of our proposed temporally-consistent diffusion model (TDM).

2.2 Training: Task Prompt Guided ControlNet Fine-Tuning with Single-Image Inputs

The pre-trained SD model excels at generating high-quality images without degradation, a feature we strive to maintain, particularly in the face of severe degradations such as sensor noise, low resolution, and haze. To this end, as shown in Fig. 2, we leverage ControlNet [36] which is a copy of the SD U-Net architecture. We fine-tune the ControlNet model initialized as the official tile resample model from ControlNet for SD version 1.5 because it is a solid foundation for image-to-image tasks. This ControlNet is then fine-tuned on our datasets, which include pairs of degraded images and their clean ground-truth images, focusing on image restoration.

Task Prompt Guidance (TPG). In our experiments, we consider five video restoration tasks: dehazing, deraining, denoising, MP4 compression artifact removal, and super-resolution. Leveraging the capacity of SD to generate diverse image styles through different text prompts, we recognize their potential to address a variety of tasks within a single model framework. To enhance this versatility, we introduce TPG, a method that employs specific task descriptions as input text prompts (τ) for both training and inference, as shown in Fig. 1 and Fig. 2. Unlike the approach in [21], which relies on separate plugin modules for each task and a classifier to select the appropriate plugin based on the prompt, we utilize the proven proficiency of the pre-trained SD model to interpret and classify directly from text prompts as noted in studies [6, 7].

Single-Image Inputs During Training. As illustrated in Fig. 2, our proposed method requires only single-image inputs for video restoration tasks during the fine-tuning of the ControlNet. Drawing inspiration from previous works [14, 38], we achieve significant temporal consistency through training-free operations at the inference stage, thereby circumventing the high memory demands of processing multiple video frames simultaneously during training, which is a common requirement in earlier approaches [2, 10, 31] for achieving temporal consistency. This strategy allows our model to be trained efficiently on a single GPU using single-image restoration datasets, significantly reducing data preparation complexity and making it more accessible for future research. This approach paves the way for future research to easily append novel tasks with training on corresponding single-image datasets.

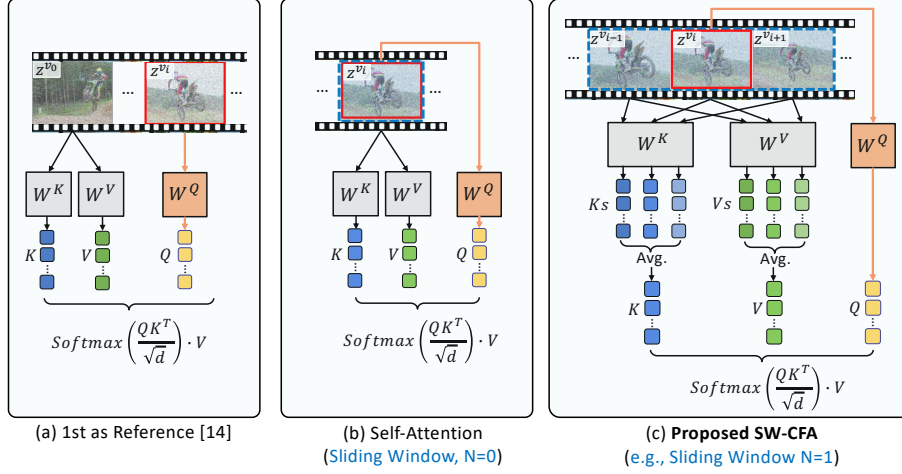


Fig. 3. Proposed SW-CFA compared with existing cross-frame attention.

2.3 Inference: Training-Free Content-Preserved Temporal Consistency for Larger Motion

As mentioned earlier, we only use single-image inputs in the training phase, so temporal consistency is addressed through a training-free approach during the inference. Previous studies [14, 38] achieve this by replacing self-attention layers in denoising U-Nets with cross-frame attention layers. LDM models utilize a U-Net architecture with downscaling and upscaling phases, enhanced with skip connections. The architecture includes 2D convolutional residual blocks and transformer blocks, each containing a self-attention layer, a cross-attention layer, and a feed-forward network. Self-attention formulates spatial correlations within the feature map, while cross-attention maps relationships between the feature map and external conditions like text prompts. For a video frame v_i with latent representation z^{v_i} , as shown in Fig. 3(b), self-attention is defined as: $Attention(Q, K, V) = Softmax\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V$, where

$$Q = W^Q z^{v_i}, K = W^K z^{v_i}, V = W^V z^{v_i}, \quad (5)$$

with W^Q , W^K , and W^V being trainable matrices for query, key, and value projections, respectively, and d is the dimension of key and query features. Previous studies [14] reorganize self-attention into cross-frame attention using a single reference latent $z_{v_{ref}}$ as both the key (K) and value (V). Typically, this reference frame is set as the first frame (v_0) as shown in Fig. 3(a):

$$Q = W^Q z^{v_i}, K = W^K z^{v_0}, V = W^V z^{v_0}, \quad (6)$$

This captures the spatial relationship between the current frame’s query and the reference frame’s key, maintaining the appearance, structure, and identities across frames, thus enhancing temporal consistency. However, using only v_0 as the reference limits

motion range, as large motions may result in limited overlap between v_0 and subsequent frames. Additionally, the consistency between later consecutive frames is not explicitly addressed, further limiting overall temporal consistency even with small motions.

Sliding Window Cross-Frame Attention (SW-CFA). We propose a novel cross-frame attention mechanism named SW-CFA to accommodate a wider range of motions. As shown in Fig. 3(c), instead of relying on a single fixed reference frame for the key (K) and value (V), we extend the reference frames to include frames within a local window. Specifically, we average the keys and values inside the window, which includes both the preceding and succeeding N frames. The formulation is as follows:

$$Q = W^Q z^{v_i}, K = \frac{1}{2N+1} \sum_{j=i-N}^{i+N} W^K z^{v_j}, V = \frac{1}{2N+1} \sum_{j=i-N}^{i+N} W^V z^{v_j}, \quad (7)$$

which can be further simplified as

$$Q = Q_i, K = \frac{1}{2N+1} \sum_{j=i-N}^{i+N} K_j, V = \frac{1}{2N+1} \sum_{j=i-N}^{i+N} V_j, \quad (8)$$

and our SW-CFA is formulated as follows:

$$\text{Attention}(Q, K, V) = \text{Softmax} \left(Q_i \frac{1}{2N+1} \sum_{j=i-N}^{i+N} K_j^T \right) \left(\frac{1}{2N+1} \sum_{j=i-N}^{i+N} V_j \right). \quad (9)$$

By averaging K_j and V_j across a local window of frames, SW-CFA smooths over rapid variations, focusing on stable and consistent features. This is not merely filtering but a principled integration of temporal information, enhancing the model’s ability to prioritize relevant information across temporal axis.

Mathematically, this averaging process reduces variance in the attention mechanism’s input, effectively functioning as a temporal low-pass filter. This reduction allows the attention mechanism to produce outputs less sensitive to frame-to-frame fluctuations, increasing temporal consistency. Within the softmax function, averaging indirectly weights each nearest N frames’ K and V matrices according to their temporal proximity and similarity to the current frame i , leveraging the softmax function’s property of amplifying significant signals while attenuating weaker ones. Thus, SW-CFA captures both spatial dependencies within frames and temporal dependencies. This results in a robust attention mechanism that enhances temporal consistency through a simple yet effective modification of the previous cross-frame attention paradigm.

Combination with DDIM Inversion. Our proposed SW-CFA mechanism significantly enhances temporal consistency. However, random Gaussian noise added to input latents, as noted in [21], can reduce the fidelity of diffusion-generated images. To address this, we integrate DDIM Inversion [8, 29] to provide stable input noise for each video frame’s latents, serving as a solid structural guide, as shown in Fig. 2. DDIM Inversion ensures the noise added to each input frame’s latents is deterministic, consistent, and derived from the latents themselves. Since these input latents originate from temporally coherent video frames, the noise introduced by DDIM Inversion also exhibits temporal coherence. This synergy eliminates the disruption caused by random

Gaussian noise, enabling SW-CFA to more precisely capture temporal relationships. Therefore, combining DDIM Inversion not only improves content preservation but also further strengthens temporal consistency.

3 Experiments

3.1 Settings

Here, we outline our experimental settings, including the datasets, implementation details, and evaluation metrics.

Datasets. We evaluated the proposed TDM for five video restoration tasks. To train the TDM and the compared methods, we utilized representative datasets for each task, dehazing: REVIDE [37], deraining: NTURain-syn [35], denoising: DAVIS [15], MP4 compression artifacts removal: MFQEv2 [9], and super-resolution (SR): REDS [23]. Only REVIDE is a real-world dataset and the others are synthetic datasets. To avoid unbalanced training due to largely different amounts of training images for each task, we adjusted the number of images from each dataset to around 5,000-6,000 images. Then, we combined all of those images to construct the training dataset, resulting in a total of 27,843 training images. For testing, we used real-world benchmark datasets for each task, dehazing: REVIDE [37] (284 images), denoising: CRVD [34] (560 images), deraining: NTURain-real [35] (658 images), MP4 compression artifacts removal: MFQEv2 [9] (1,080 images), and SR: UDM10 [33] (320 images).

Implementation. During the training and the inference phase, we adjusted the size of each image by resizing it to ensure that the shorter side is 512 pixels. For training, we randomly cropped patches of 512x512 pixels from resized images. We utilized the AdamW optimizer with its default settings, including betas and weight decay. The training of our ControlNet was carried out with a constant learning rate of $1e-5$ and a batch size of 4, on a single RTX 4090 GPU, and we incorporated gradient checkpoints for efficiency. The training process lasts for 25 epochs. In the inference stage, we set the window radius N for SW-CFA as 3 for all experiments. For the sampling technique, we initially applied DDIM Inversion [8] with 10 timesteps, followed by DDIM backward sampling [29] using 32 timesteps. Our proposed TDM model can process a 15-frame video with resolutions of 512x896 pixels in under 30 seconds with a single RTX 4090 GPU while requiring 10GB of GPU memory.

Metrics. We followed [21, 27] to employ widely adopted non-reference perceptual metrics, FID [11] and KID [1], to evaluate our TDM, as the ground-truth images are not always available for real-world datasets. For easier view, the KID value is scaled by $100\times$. For temporal consistency evaluation, we followed [38] to estimate (i) Frame consistency (FC): the average cosine similarity between all pairs of consecutive frames, and (ii) Warping Error (WE): the average mean squared error of consecutive frames after aligning the next frame to the current one using estimated optical flow. For easier view, FC is scaled by $10\times$ and WE is scaled by $1000\times$.

3.2 Comparison with State-of-the-Art Methods

We compared our TDM with six state-of-the-art methods for image/video restoration: AirNet [16], PromptIR [25], VRT [17], RVRT [18], WeatherDiff [24], and InstructP2P [3].

Table 1. Quantitative comparison with state-of-the-art methods (Red: rank 1st; Blue: rank 2nd).

Method Types	Methods	Dehazing		Deraining		Denoising		MP4		SR $\times 4$		Average	
		FID↓	KID↓	FID↓	KID↓	FID↓	KID↓	FID↓	KID↓	FID↓	KID↓	FID↓	KID↓
Single-image regression	AirNet [16]	83.34	7.77	90.86	4.24	80.63	4.41	104.48	6.56	88.09	3.07	89.47	5.21
	PromptIR [25]	75.21	5.73	88.47	5.28	82.11	4.95	102.73	6.68	89.37	3.12	87.57	5.15
Video regression	VRT [17]	79.88	7.03	88.36	4.69	81.94	4.88	107.00	6.56	88.97	2.71	89.23	5.17
	RVRT [18]	79.27	6.55	94.13	5.44	83.59	4.57	107.37	6.55	88.51	2.93	90.57	5.20
Single-image diffusion	WeatherDiff [24]	74.42	6.22	88.20	4.85	80.11	4.18	102.24	6.67	88.07	2.77	86.28	4.93
	InstructP2P [3]	74.02	6.01	81.51	3.36	79.71	3.90	101.62	7.70	88.21	2.62	85.01	4.71
Video diffusion	TDM (Ours)	73.68	6.42	81.27	3.79	78.63	4.09	100.91	6.59	88.04	2.55	84.50	4.68

As the category of each method, AirNet, PromptIR, VRT, and RVRT are regression-based methods, such as based on CNNs or Transformers, while WeatherDiff, InstructP2P, and our TDM are diffusion-based methods. As the input for the inference, AirNet, PromptIR, WeatherDiff, InstructP2P use a single-image input, while VRT, RVRT, and our TDM use a video input. We trained all methods using the same training dataset and tested them for real-world testing datasets, as explained in the previous subsection.

Quantitative results. We provide a quantitative comparison in Table 1. It demonstrates the superiority of diffusion-based methods in producing higher-quality images, as evidenced by better FID and KID scores compared to regression-based methods. Compared with other regression- and diffusion-based methods, our proposed TDM achieves the best results on average. Although InstructP2P also generates high-quality images with low FID and KID, it tends to alter the original contents, as seen in Fig. 4. Compared to the regression-based methods including AirNet and PromptIR, which are designed for generalization across different tasks, our TDM still demonstrates robust cross-task performance with the guidance of the proposed TPG. Furthermore, while video restoration methods such as VRT and RVRT are designed to utilize video temporal information effectively during the training phase, they struggle with multi-task handling. In contrast, our TDM consistently delivers state-of-the-art performance in video restoration tasks, even though it is trained on single-image inputs.

Qualitative results. Figure 4 illustrates the robust performance of our TDM across five challenging real-world video restoration tasks. In the denoising task (1st row), while most regression-based methods trained on synthesized Gaussian noise struggle to recognize real-world noise patterns, all diffusion-based methods including TDM successfully eliminate the real-world noise. However, TDM distinctly outperforms WeatherDiff and InstructP2P, which both leave artifacts post-denoising. In the dehazing task (2nd row), TDM excels by clearing heavy fog to reveal the sharpest and the most haze-free images. For the MP4 compression artifact removal task (3rd row), TDM outperforms regression-based methods in detail restoration. Compared with diffusion-based InstructP2P, which inaccurately alters hairstyles, our TDM can restore sharp details and maintain consistency with the original input. In the SR task (4th row), TDM achieves the clearest detail enhancement, whereas other methods produce comparatively blurred results. Finally, in the challenging heavy rain scenario (5th row), while other models fail to detect and remove the rain, TDM significantly reduces rain artifacts, demonstrating its robustness in severe weather conditions. Overall, our TDM consistently delivers



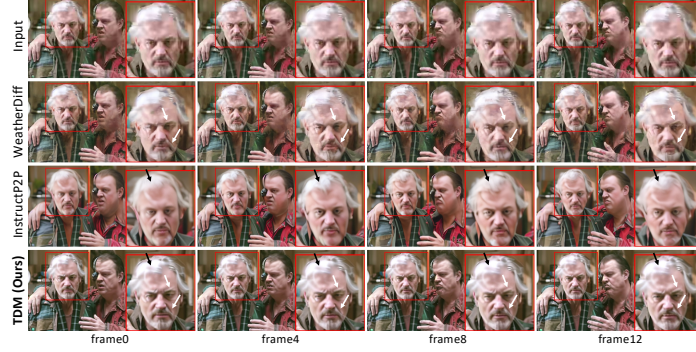
Fig. 4. Qualitative comparison with state-of-the-art methods.

superior restoration quality with remarkable detail preservation and consistency across different video restoration challenges.

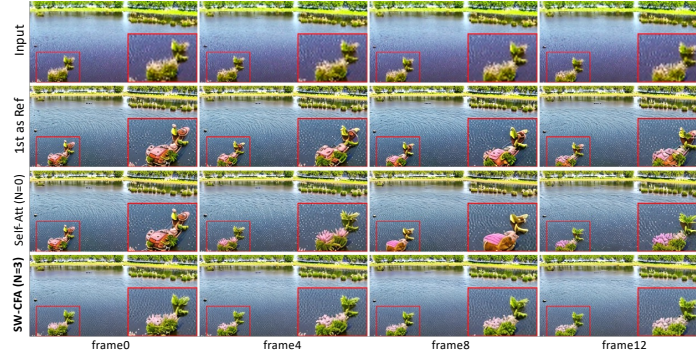
Consistency Evaluation. Although diffusion-based methods typically yield images of high visual quality, their inherent randomness often leads to poor temporal consistency in video processing. In Table 2, we compare our TDM against other diffusion-based techniques using two temporal consistency metrics: WE and FC. The results show that TDM consistently outperforms the others in maintaining temporal consistency according to both metrics. Figure 5 provides a visual comparison that supports these results. While the methods WeatherDiff and InstructP2P can remove the noise, they exhibit notable inconsistencies between the frames. WeatherDiff (2nd row) exhibits heavy changes in the appearance of building frames in the background. For InstructP2P (3rd row), despite using image-based classifier-free guidance, it still displays fluctuation across frames. In contrast, TDM, utilizing DDIM inversion to eliminate randomness and SW-CFA to bolster temporal stability, maintains consistent features across all frames, leading to stable denoising results. This robustness underscores TDM’s superior temporal consistency compared to other diffusion-based methods. Also, in Table 3 and Fig. 6, we evaluate our proposed SW-CFA against the cross-frame attention from Text2Video-zero [14] (referred to as "1st as Ref") and the standard self-attention ($N=0$). While "1st as Ref" enhances consistency by using the first frame as a reference, it overlooks the consistency of consecutive frames. In contrast, our SW-CFA, which averages key-value pairs within a sliding window, achieves more substantial consistency

Table 2. Image quality and consistency comparison with other diffusion-based methods.

Methods	Dehazing			Deraining			Denoising			MP4			SR $\times 4$		
	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓
WeatherDiff [24]	74.42	9.759	8.537	88.20	9.682	3.652	80.11	9.439	2.182	102.24	9.908	1.282	88.07	9.682	5.751
InstructP2P [3]	74.02	9.663	6.994	81.51	9.676	4.339	79.71	9.436	2.628	101.62	9.877	2.126	88.21	9.628	5.929
TDM (Ours)	73.68	9.849	5.464	81.27	9.700	3.415	78.63	9.296	2.008	100.91	9.921	1.138	88.04	9.651	5.741

**Fig. 5.** Consistency comparison (MP4) with other diffusion-based methods.**Table 3.** Image quality and consistency comparison with state-of-the-art zero-shot cross-frame attention (by replacing the self-attention with according cross-frame attention during inference).

Methods	Dehazing			Deraining			Denoising			MP4			SR $\times 4$		
	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓	FID↓	FC↑	WE↓
1st as Ref. [14]	74.37	9.712	6.432	82.39	9.693	3.526	80.12	9.229	2.219	99.63	9.901	1.387	88.62	9.622	6.350
Self-Attn. ($N=0$)	73.36	9.703	6.801	83.06	9.690	3.596	80.17	9.213	2.368	100.56	9.896	1.459	88.30	9.618	6.769
SW-CFA ($N=3$)	73.68	9.849	5.464	81.27	9.700	3.415	78.63	9.296	2.008	100.91	9.921	1.138	88.04	9.651	5.741

**Fig. 6.** Consistency comparison (SR $\times 4$) with other zero-shot cross-frame attention.

improvements. It demonstrates a superior ability to maintain uniformity across frames, effectively handling larger motions and providing robust temporal stability.

Table 4. Ablation study.

Methods	Proposals			Dehazing			Deraining			Denoising			MP4			SR $\times 4$		
	TPG	Inv.	SW-CFA	FID \downarrow	FC \uparrow	WE \downarrow	FID \downarrow	FC \uparrow	WE \downarrow	FID \downarrow	FC \uparrow	WE \downarrow	FID \downarrow	FC \uparrow	WE \downarrow	FID \downarrow	FC \uparrow	WE \downarrow
T+I	✓	✓		73.36	9.703	6.801	83.06	9.690	3.596	80.17	9.213	2.368	100.56	9.896	1.459	88.30	9.618	6.769
T+S	✓		✓	73.45	9.738	7.121	82.79	9.673	4.125	80.07	9.207	3.792	101.47	9.866	1.841	88.35	9.552	8.418
I+S		✓	✓	77.56	9.838	5.653	87.28	9.698	3.437	82.68	9.273	2.296	102.44	9.909	1.190	91.02	9.630	5.874
Ours	✓	✓	✓	73.68	9.849	5.464	81.27	9.700	3.415	78.63	9.296	2.008	100.91	9.921	1.138	88.04	9.651	5.741

3.3 Ablation Study

In Table 4, we present an ablation study to assess the contribution of each component of our proposed TDM. We test three configurations: T+I, which omits the SW-CFA; T+S, which excludes DDIM Inversion (labeled as Inv.); and I+S, which lacks the TGP. The results show that removing either SW-CFA or DDIM Inversion diminishes temporal consistency, while their combination significantly enhances it. This indicates that DDIM Inversion, by stabilizing the input noise across frames, substantially supports SW-CFA in maintaining consistency. Furthermore, the setup without TGP achieves comparable consistency scores, but there is a noticeable degradation in image quality, as evidenced by the increased FID scores. This underscores the critical role of TGP in boosting cross-task generalization and improving overall image quality in various restoration tasks.

4 Conclusion

In conclusion, we introduced the Temporally-consistent Diffusion Model (TDM) for all-in-one video restoration, utilizing a pre-trained Stable Diffusion model and fine-tuned ControlNet. Our method handles various video degradations with a single model, using Task Prompt Guidance (TPG) for training and combining DDIM Inversion with Sliding Window Cross-Frame Attention (SW-CFA) for enhanced temporally-consistent video inference. Experiments across five tasks exhibits proposed TDM’s superior generalization over existing methods, setting a new standard for video restoration. However, our method still falls short of regression-based methods in temporal consistency. Future work will focus on addressing this gap to further improve effectiveness.

References

1. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd gans. arXiv preprint arXiv:1801.01401 (2018)
2. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
3. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18392–18402 (2023)
4. Cao, J., Wang, Q., Liang, J., Zhang, Y., Zhang, K., Timofte, R., Van Gool, L.: Learning task-oriented flows to mutually guide feature alignment in synthesized and real video denoising. arXiv preprint arXiv:2208.11803 (2022)

5. Chan, K.C., Zhou, S., Xu, X., Loy, C.C.: Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5972–5981 (2022)
6. Chen, H., Dong, Y., Wang, Z., Yang, X., Duan, C., Su, H., Zhu, J.: Robust classification via a single diffusion model. *arXiv preprint arXiv:2305.15241* (2023)
7. Clark, K., Jaini, P.: Text-to-image diffusion models are zero shot classifiers. *Advances in Neural Information Processing Systems* **36** (2024)
8. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)
9. Guan, Z., Xing, Q., Xu, M., Yang, R., Liu, T., Wang, Z.: Mfqc 2.0: A new approach for multi-frame quality enhancement on compressed video. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(3), 949–963 (2019)
10. 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)
11. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems* **30** (2017)
12. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
13. Ke, B., Obukhov, A., Huang, S., Metzger, N., Daudt, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024)
14. Khachatryan, L., Movsisyan, A., Tadevosyan, V., Henschel, R., Wang, Z., Navasardyan, S., Shi, H.: Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 15954–15964 (2023)
15. Khoreva, A., Rohrbach, A., Schiele, B.: Video object segmentation with language referring expressions. In: *Asian Conference on Computer Vision*. pp. 123–141 (2019)
16. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17452–17462 (2022)
17. Liang, J., Cao, J., Fan, Y., Zhang, K., Ranjan, R., Li, Y., Timofte, R., Van Gool, L.: Vrt: A video restoration transformer. *arXiv preprint arXiv:2201.12288* (2022)
18. Liang, J., Fan, Y., Xiang, X., Ranjan, R., Ilg, E., Green, S., Cao, J., Zhang, K., Timofte, R., Gool, L.V.: Recurrent video restoration transformer with guided deformable attention. *Advances in Neural Information Processing Systems* **35**, 378–393 (2022)
19. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9298–9309 (2023)
20. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453* (2023)
21. Liu, Y., Liu, F., Ke, Z., Zhao, N., Lau, R.W.: Diff-plugin: Revitalizing details for diffusion-based low-level tasks. *arXiv preprint arXiv:2403.00644* (2024)
22. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6038–6047 (2023)
23. Nah, S., Baik, S., Hong, S., Moon, G., Son, S., Timofte, R., Lee, K.M.: Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (2019)

24. Özdenizci, O., Legenstein, R.: Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
25. Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.S.: Promptir: Prompting for all-in-one blind image restoration. *arXiv preprint arXiv:2306.13090* (2023)
26. Ren, M., Delbracio, M., Talebi, H., Gerig, G., Milanfar, P.: Multiscale structure guided diffusion for image deblurring. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10721–10733 (2023)
27. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10684–10695 (2022)
28. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: *Proceedings of ACM SIGGRAPH Conference*. pp. 1–10 (2022)
29. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
30. Tassano, M., Delon, J., Veit, T.: Fastdvdnet: Towards real-time deep video denoising without flow estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1354–1363 (2020)
31. Wang, J., Yuan, H., Chen, D., Zhang, Y., Wang, X., Zhang, S.: Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571* (2023)
32. Wu, J.Z., Ge, Y., Wang, X., Lei, S.W., Gu, Y., Shi, 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. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7623–7633 (2023)
33. Yi, P., Wang, Z., Jiang, K., Jiang, J., Ma, J.: Progressive fusion video super-resolution network via exploiting non-local spatio-temporal correlations. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3106–3115 (2019)
34. Yue, H., Cao, C., Liao, L., Chu, R., Yang, J.: Supervised raw video denoising with a benchmark dataset on dynamic scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2301–2310 (2020)
35. Yue, Z., Xie, J., Zhao, Q., Meng, D.: Semi-supervised video deraining with dynamical rain generator. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 642–652 (2021)
36. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3836–3847 (2023)
37. Zhang, X., Dong, H., Pan, J., Zhu, C., Tai, Y., Wang, C., Li, J., Huang, F., Wang, F.: Learning to restore hazy video: A new real-world dataset and a new method. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9239–9248 (2021)
38. Zhang, Y., Wei, Y., Jiang, D., Zhang, X., Zuo, W., Tian, Q.: Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077* (2023)