

NightRain: Nighttime Video Deraining via Adaptive-Rain-Removal and Adaptive-Correction

Beibei Lin¹, Yeying Jin¹, Wending Yan², Wei Ye², Yuan Yuan², Shunli Zhang³, Robby Tan¹

¹National University of Singapore

²Huawei International Pte Ltd

³Beijing Jiaotong University

{beibei.lin, e0178303}@u.nus.edu, {yan.wending, yewei10, yuanyuan10}@huawei.com, slzhang@bjtu.edu.cn, robby.tan@nus.edu.sg

Abstract

Existing deep-learning-based methods for nighttime video deraining rely on synthetic data due to the absence of real-world paired data. However, the intricacies of the real world, particularly with the presence of light effects and low-light regions affected by noise, create significant domain gaps, hampering synthetic-trained models in removing rain streaks properly and leading to over-saturation and color shifts. Motivated by this, we introduce NightRain, a novel nighttime video deraining method with adaptive-rain-removal and adaptive-correction. Our adaptive-rain-removal uses unlabeled rain videos to enable our model to derain real-world rain videos, particularly in regions affected by complex light effects. The idea is to allow our model to obtain rain-free regions based on the confidence scores. Once rain-free regions and the corresponding regions from our input are obtained, we can have region-based paired real data. These paired data are used to train our model using a teacher-student framework, allowing the model to iteratively learn from less challenging regions to more challenging regions. Our adaptive-correction aims to rectify errors in our model’s predictions, such as over-saturation and color shifts. The idea is to learn from clear night input training videos based on the differences or distance between those input videos and their corresponding predictions. Our model learns from these differences, compelling our model to correct the errors. From extensive experiments, our method demonstrates state-of-the-art performance. It achieves a PSNR of 26.73dB, surpassing existing nighttime video deraining methods by a substantial margin of 13.7%.

Introduction

Nighttime rain videos are adversely affected by several factors, including low light, nighttime light effects (e.g., glow, glare, and uneven light distribution), rain streaks, etc. Existing nighttime video deraining methods (i.e., (Patil et al. 2022b,a)) rely on paired data for training. Due to the lack of real-world paired data, these methods are trained using synthetic datasets. However, complex nighttime light effects introduce a significant domain gap between synthetic and real-world rain streak distribution. Consequently, these methods struggle to effectively remove real-world rain streaks, particularly in the presence of complex light effects. Moreover,

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Qualitative results on real-world nighttime rain videos. First column: Input image. Second column: MetaRain’s results (Patil et al. 2022b). Third column: Our results. Zoom-in for better visualization.

models trained on synthetic datasets tend to produce over-saturated and color-shifted results. As shown in Fig. 1, existing nighttime video deraining methods yield over-saturated light regions and color-shifts in some regions.

Existing daytime video deraining methods (Zhang et al. 2022; Kulkarni, Patil, and Murala 2021; Wang et al. 2022a) can be retrained to remove nighttime rain streaks. However, these methods are designed for daytime scenes that do not have complex lighting conditions and the significant presence of noise in low-light regions. Thus, they struggle to handle rain streaks in nighttime scenes. Existing single-image deraining methods (Shi et al. 2020; Tang et al. 2020; Chen et al. 2023; Jiang et al. 2020) can be adapted to remove rain streaks by incorporating alignment and temporal smoothness priors. Nonetheless, they also encounter challenges from nighttime light effects and low-light regions. Additionally, alignment becomes problematic in the presence of intricate lighting conditions, rain streaks, and noise in low-light regions.

To address these challenges, we introduce *NightRain*, a novel nighttime video deraining method with two main mod-

ule: **adaptive-rain-removal** and **adaptive-correction**. Note that, we exclude night videos with the rain veiling effect (or rain accumulation) since these videos typically do not exhibit prominent rain streaks.

In our NightRain, our adaptive-rain-removal leverages unlabeled rain videos to enhance our model’s deraining ability in real-world videos, particularly in low-light regions or regions affected by complex light effects. Specifically, we employ a pretrained model as a teacher model to generate rain-free predictions based on confidence scores. These high-confidence predictions primarily arise from less challenging regions, such as less pronounced low-light regions or regions with slight light effects. In contrast, low-confidence predictions tend to result in inaccurate restorations because their corresponding inputs are more likely influenced by complex light effects or contain significantly low light regions. High-confidence predictions and their corresponding inputs are then used as region-based paired real data to train a student model. Through iteratively transferring knowledge from the student model to the teacher model, our teacher model gains the ability to handle more challenging inputs, resulting in improved performance. As shown in Fig. 1, our model effectively removes rain streaks even in the presence of light effects. Our adaptive-rain-removal operates as an unsupervised method as it does not require ground-truths.

Our adaptive-correction aims to rectify errors in our model’s predictions, such as over-saturation and color shifts. To achieve this, we utilize our video diffusion model to generate predictions from real-world nighttime clear videos. The discrepancies between clear input videos and our predictions essentially represent these errors. Thus, we take these predictions and their corresponding inputs as paired data to continually refine our model. Through iterative error correction, our model is compelled to learn a more comprehensive nighttime clear distribution, leading to more precise predictions. As shown in Fig. 1, our method removes real-world rain streaks without introducing over-saturation and color-shifted problems. Our adaptive-correction operates as an unsupervised method as we only use nighttime clear videos as inputs.

As a summary, our main contributions as follows:

1. We introduce adaptive-rain-removal to enable our model to derain real-world nighttime rain videos, especially in low-light regions or regions affected by complex light effects. By learning high-confidence predictions, our adaptive-rain-removal process forces our model to iteratively address more challenging regions, leading to effective real-world rain streak removal.
2. We propose adaptive-correction to address over-saturated and color-shifted problems. By learning from differences between clear night videos and their corresponding predictions, our adaptive-correction continually rectifies errors in our models’ predictions, resulting in more accurate restoration.
3. Experimental results demonstrate that our method not only achieves state-of-the-art restoration performance on the SynNightRain dataset, improving by 13.7%, but also effectively removes real-world rain streaks even in the

presence of light effects.

Related Work

Video Deraining Existing nighttime video deraining methods (Patil et al. 2022a,b) remove rain streaks by networks trained on synthetic datasets. They synthesize rain streaks on clear nighttime videos, and then train CNNs on the synthetic datasets. However, their methods suffer from the domain gap problem since there exists a significant domain gap between synthetic and real-world rain streak dsi. Another problem of existing video deraining methods is that they tend to over-suppress the input videos, leading to color-shifted and over-saturated results.

Self-learning daytime video deraining methods (Yang et al. 2020; Yan et al. 2021) do not require paired data to train their network and may be applied for nighttime video deraining. Specifically, their methods are based on an assumption that adjacent frames may contain rain-free information and can be used to estimate a rain-free current frame. However, these methods cannot remove dense rain streaks since adjacent frames in heavy rain do not contain enough rain-free information. Some supervised daytime video deraining methods (Zhang et al. 2022; Kulkarni, Patil, and Murala 2021; Wang et al. 2022a,b; Zhuang et al. 2022; Xue et al. 2020; Wang et al. 2019; Islam and Paul 2021; Su et al. 2023; Li et al. 2023) can be directly used to address nighttime rain streaks. However, these methods also suffer from the domain gap problem since they rely on synthetic datasets for training.

Diffusion Models Existing diffusion models (Özdenizci and Legenstein 2023; Harvey et al. 2022; Peebles and Xie 2022; Saharia et al. 2022; Ma, Fang, and Wang 2023; Fang et al. 2023) show a remarkable success in different generative tasks. Ozdenizci et al. (Özdenizci and Legenstein 2023) propose a 2D UNet-based diffusion model to address multiple weather degradation, i.e., snow, raindrop and rain. Harvey et al. (Harvey et al. 2022) present a 3D UNet-based diffusion model for video completions. These diffusion models can be re-trained on nighttime rain streak removal datasets and then used to remove real-world nighttime rain streaks. However, due to the lack of real-world paired data, the network training depends on synthetic datasets. As a result, their predictions also suffer from the domain gap problem, leading to inaccurate restoration.

Proposed Method: NightRain

Figs. 2 and 3 display the two main training processes for our NightRain: adaptive-rain-removal and adaptive-correction. The goal of our adaptive-rain-removal is to enable our model to remove rain streaks from real night videos without requiring ground-truths in its training phase. The objective of our adaptive-correction is to enable our model to rectify errors in its predictions, such as over-saturation or color shifts, by exploiting nighttime clear videos. These training processes are separate from each other.

Initially, our adaptive-rain-removal and adaptive-correction rely on a pre-trained video diffusion model w trained on a synthetic dataset $\mathbf{D}_s = \{\mathbf{x}_i^s, \mathbf{y}_i^s\}_{i=1}^{N_s}$, where

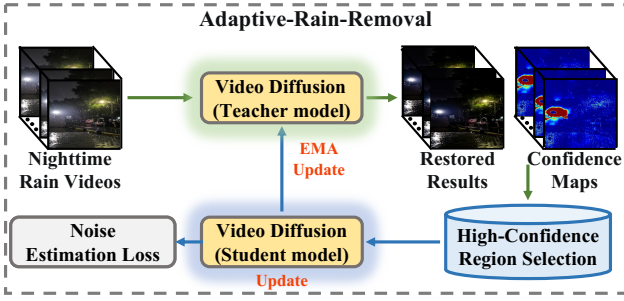


Figure 2: Overview of our adaptive-rain-removal. We pre-train a video diffusion model on synthetic nighttime deraining datasets as a teacher model. Our adaptive-rain-removal utilizes the teacher model to generate predictions from real-world nighttime rain videos. We also generate confidence maps of rain streak removal within these predictions (green lines). The high-confidence predictions with their corresponding inputs are then selected to train a student model, thus reducing the domain gap (blue lines). Finally, we utilize Exponential Moving Average (EMA) to update our teacher model.

\mathbf{x}_i^s and \mathbf{y}_i^s are the i -th input and ground-truth videos, respectively. N_s is the number of supervised videos. The pre-trained video diffusion model is implemented by a conditional video diffusion model. During the training phase, instead of sampling the input data $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ (standard diffusion model), we sample $(\mathbf{x}_0, \tilde{\mathbf{x}}) \sim q(\mathbf{x}_0, \tilde{\mathbf{x}})$ to train our diffusion model, where \mathbf{x}_0 and $\tilde{\mathbf{x}}$ denote clear and rain videos. The reverse process can be formulated as:

$$p_w(\mathbf{x}_{0:T} | \tilde{\mathbf{x}}) = p(\mathbf{x}_T) \prod_{t=1}^T p_w(\mathbf{x}_{t-1} | \mathbf{x}_t, \tilde{\mathbf{x}}). \quad (1)$$

Here, the inputs \mathbf{x} and $\tilde{\mathbf{x}}$ are concatenated channel-wise. During the inference phase, given a random Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and a degraded video $\tilde{\mathbf{x}}$, the clear video \mathbf{x}_0 can be generated by gradually reversing a diffusion process. $\mathbf{x}_{t-1} \sim p_w(\mathbf{x}_{t-1} | \mathbf{x}_t, \tilde{\mathbf{x}})$ is computed as:

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \cdot \boldsymbol{\epsilon}_w(\mathbf{x}_t, \tilde{\mathbf{x}}, t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \cdot \boldsymbol{\epsilon}_w(\mathbf{x}_t, \tilde{\mathbf{x}}, t). \quad (2)$$

where $\boldsymbol{\epsilon}_w(\cdot)$ and $\boldsymbol{\epsilon}_w(\mathbf{x}_t, \tilde{\mathbf{x}}, t)$ are the noise estimation network and predicted noise vectors, respectively.

In our method, we use a video transformer to implement our noise estimation network, which is shown in Fig. 4. Given a sequence $X_s \in \mathbb{R}^{C_s \times T_s \times H_s \times W_s}$ and a time step $T_s \in \mathbb{R}^1$, where C_s is the number of channels, T_s is the length of sequences and (H_s, W_s) is the image size of each frame, we utilize a 3D convolution to partition the input sequences and utilize MLP to embed the time step. The two operations can be expressed as:

$$X_p = F_{ts \times ss \times ss}(X_s), \quad (3)$$

$$T_e = \text{MLP}(T_s), \quad (4)$$

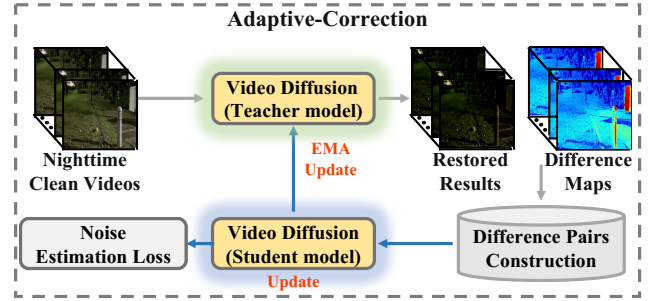


Figure 3: Overview of our adaptive-correction. We utilize our teacher model to generate predictions from nighttime clear videos (grey lines). The difference regions between clear videos and their corresponding predictions essentially represent errors produced by our model itself. We then use these difference pairs to train a student model, thus correcting our model's errors (blue lines). Finally, we utilize Exponential Moving Average (EMA) to update our teacher model.




Figure 4: Overview of our transformer-based noise estimation Network. Given a sequence and a time step as inputs, we convert the sequence into patches and encode the time step to the time embedding. Subsequently, several transformer blocks are applied to extract global spatiotemporal cues. Finally, a linear layer is proposed to convert feature representations to noise maps.

where $X_p \in \mathbb{R}^{P \times C_p}$ and $T_e \in \mathbb{R}^{C_p}$ are patches and the time embedding. $F_{ts \times ss \times ss}(\cdot)$ is a 3D convolution with the kernel size is $(ts \times ss \times ss)$.

The patches X_p and the time embedding T_e are then injected into our transformer blocks, consisting of two LN layers, a Multi-head Self-Attention (MSA) module and an MLP layer, to extract global temporal and spatial information. We follow (Peebles and Xie 2022) to inject the time embedding for our diffusion model. Our transformer block first linearly transforms the time embedding to six parts, *i.e.*, s_{msa} , g_{msa} , c_{msa} , s_{mlp} , g_{mlp} and c_{mlp} . The size of each part is \mathbb{R}^{C_p} . Then, a transformer block can be represented as:

$$\begin{aligned} Y_p^{msa} &= X_p + g_{msa} * \text{MSA}(\text{adaLN}(X_p, s_{msa}, c_{msa})), \\ Y_p &= Y_p^{msa} + g_{mlp} * \text{MLP}(\text{adaLN}(Y_p^{msa}, s_{mlp}, c_{mlp})), \end{aligned}$$

where $Y_p \in \mathbb{R}^{P \times C_p}$ and $T_e \in \mathbb{R}^{C_p}$ is the output of a transformer block. s_{msa} and s_{mlp} are learnable shift parameters of adaptive layer norm (adaLN), while c_{msa} and c_{mlp} are learnable scale parameters of adaLN. g_{msa} and g_{mlp} are learnable scaling parameters of our transformer block. After

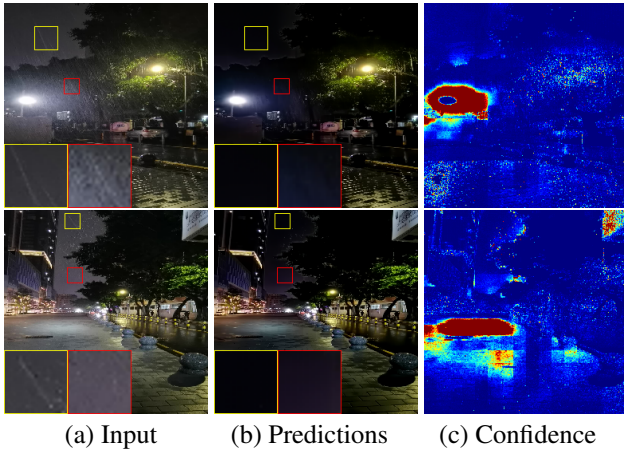


Figure 5: Visualization of predictions and confidence maps from our pretrained model. The red and blue regions are low-confidence and high-confidence predictions, respectively.

using transformer blocks and a linear layer, we can obtain noise maps $\epsilon \in \mathbb{R}^{\frac{C_s}{2} \times T_s \times H_s \times W_s}$.

Once the pretrained video diffusion model w is obtained, our adaptive-rain-removal and adaptive-correction train this model further using real-world nighttime rain and clear videos.

Adaptive-Rain-Removal

The core idea of adaptive-rain-removal is to improve our model’s real-world deraining capability by learning from high-confidence predictions of actual rain videos. As shown in Fig. 2, our adaptive-rain-removal is implemented through a teacher-student framework. Hence, this adaptive-rain-removal process is unsupervised and thus does not require ground-truths.

We use the pretrained video diffusion model w to initialize the teacher network w_T and the student network w_S . Given an unlabeled rain dataset $\mathbf{D}_{ur} = \{\mathbf{x}_i^{ur}\}_{i=1}^{N_{ur}}$, where \mathbf{x}_i^{ur} is the i -th rain video and N_{ur} is the number of rain videos, we utilize the teacher model w_T to generate predictions \mathbf{y}_i^{ur} and confidence maps \mathbf{u}_i^{ur} . Specifically, starting from random Gaussian noise ϵ , we can obtain a clear prediction for a rain video by gradually reversing the diffusion process (referring to Eqs (1) and (2)). Ideally, different initial Gaussian noise conditions with a rain video input should result in the same clear prediction. However, due to the complexity of rain streak distributions, the predictions from different noise maps might differ in challenging regions, such as those with complex light effects.

Motivated by this, we devise a resampling strategy to derive confidence maps for our rain removal predictions: We randomly generate N noise maps for each training video. Through a gradual reversal of the diffusion process, we obtain N predictions. The final prediction \mathbf{y}_i^{ur} is the mean of these predictions, while the confidence map \mathbf{u}_i^{ur} signifies the variance among them. These two processes can be defined as

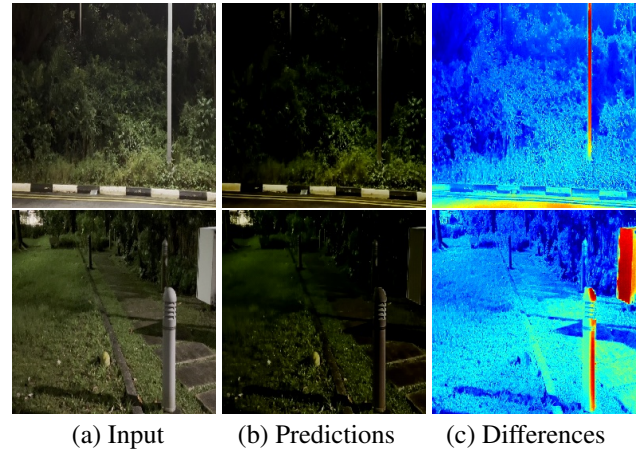


Figure 6: Visualization of predictions from our adaptive-rain-removal. (a) Input images. (b) Predictions. (c) Difference maps are calculated using the L1 distance between input images and predictions. The red and blue regions are high-difference and low-difference regions, respectively.

follows:

$$\mathbf{y}_i^{ur} = \frac{1}{N} \sum_{n=1}^N \mathbf{y}_{i,n}^{ur}, \quad (5)$$

$$\mathbf{u}_i^{ur} = \text{var}\{\mathbf{y}_{i,n}^{ur}\}_{n=1}^N. \quad (6)$$

Subsequently, we convert confidence maps to a binary mask by setting a threshold t_u . The values 0 and 1 of the binary mask represent low-confidence and high-confidence regions, respectively. Fig. 5 demonstrates that the high-confidence regions (blue regions) effectively remove most rain streaks. However, some high-confidence regions, like the sky, appear dark. This occurs because our confidence maps solely focus on the impact of rain streak removal. Consequently, the confidence scores remain high even when the rain-free background is dark. This issue can be addressed through our adaptive-correction process later.

Finally, we select high-confidence predictions and their corresponding inputs to construct region-based paired real data $\mathbf{D}_{ur} = \{\hat{\mathbf{x}}_i^{ur}, \hat{\mathbf{y}}_i^{ur}\}_{i=1}^{N_{ur}}$. These real-world pairs are then used to retrain our student model w_S . In each training step, we randomly augment the input videos by adding Gaussian noise and masking pixels. We then feed them into the student model w_S . Training our student model on these augmented inputs enables it to learn more challenging regions. After each optimization of the student model, we update the teacher model via Exponential Moving Average (EMA) that can be formulated as

$$w_T = 0.999 * w_T + 0.001 * w_S, \quad (7)$$

where w_T and w_S represent the weights of the teacher and student models, respectively. Through iterative learning, our adaptive-rain-removal enables our model to derain real-world rain videos, particularly in regions affected by complex light effects.

Adaptive-Correction

As our pretrained model relies on a synthetic dataset for training, it tends to produce over-saturated and color-shifted results, as shown in Fig. 5. Our adaptive-correction aims to address these issues, namely restoring the over-saturated and color-shifted regions. The core idea of our adaptive-correction is to rectify errors in our model’s predictions by utilizing nighttime clear videos.

Our adaptive-correction is also implemented based on a teacher-student framework. The parameters of both teacher and student networks are shared in our adaptive-rain-removal and adaptive-correction. Given a real-world clear dataset $\mathbf{D}_{uc} = \{\mathbf{y}_i^{uc}\}_{i=1}^{N_{uc}}$, where \mathbf{y}_i^{uc} is the i -th clear video and N_{uc} is the number of clear videos, we utilize the teacher model w_T to generate predictions \mathbf{x}_i^{uc} . Specifically, we first randomly sample a Gaussian noise ϵ . Then, we feed the noise and a clear video condition into our teacher model w_T . By gradually reversing the diffusion process (referring to Eqs (1) and (2)), our teacher model generate restored predictions. In theory, the prediction should be the same as the inputs since the input videos are already clear.

However, due to the lack of paired real-world data, our initial synthetic-trained diffusion model is difficult to model accurate clear distribution, leading to some over-suppressed regions in our predictions. Fig. 6 shows the input clear night videos and their corresponding predictions. We also visualize the L1 distance between them. The blue regions are low-difference regions, while the red regions represent high-difference regions. It can be observed that predictions obtained using our pretrained video diffusion model are different from their corresponding inputs. These differences essentially represent over-saturated and color-shifted problems and need to be addressed.

Motivated by this, we utilize these differences to construct paired real-world samples $\mathbf{D}_{uc} = \{\hat{\mathbf{x}}_i^{uc}, \hat{\mathbf{y}}_i^{uc}\}_{i=1}^{N_{uc}}$. We then retrain our student model w_S using these differing regions, thereby compelling our student model to rectify these errors. To be specific, the input and ground-truth videos are $\hat{\mathbf{x}}_i^{uc}$ and $\hat{\mathbf{y}}_i^{uc}$, respectively. In each training step, we randomly sample noise ϵ_t with a time step t and add it to the ground-truth $\hat{\mathbf{y}}_i^{uc}$. We subsequently feed the noise-affected ground truth and $\hat{\mathbf{x}}_i^{uc}$ into our noise estimation network for training. The objective function of the noise estimation network is to estimate accurate noise maps ϵ_t . This process can be formulated as follows:

$$\text{Loss} = \|\epsilon_t - \epsilon_{w_S}(\sqrt{\alpha_t}\hat{\mathbf{y}}_i^{uc} + \sqrt{1 - \alpha_t}\epsilon_t \hat{\mathbf{x}}_i^{uc}, t)\|^2, \quad (8)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ and $\epsilon_{w_S}(\cdot)$ is the noise estimation network of the student model. β_1, \dots, β_T are the variance changing schedule. By continually training on these differences, our student model can rectify errors from our teacher model. After each optimization of the student model, we update our teacher model via EMA (referring to Eq. (7)). This allows our teacher model to avoid producing erroneous results, leading to more precise restored results.

In the training stage, we jointly use adaptive-rain-removal and adaptive-correction to train our student model w_S . The two processes share the same teacher and student parameters. After each optimization of the student model, we use

EMA to update the teacher model w_T . In the inference stage, given a degraded video, we directly use the teacher model w_T to generate rain-free predictions. Note that, $w_T^{\text{SR}} = w_T^{\text{SC}}$ and $w_S^{\text{SR}} = w_S^{\text{SC}}$, where T represents the teacher model, S represents the student model, SR stands for adaptive-rain-removal and SC stands for adaptive-correction. The network parameters are always shared.

Experiments

SynNightRain (Patil et al. 2022b,a) is a synthetic nighttime video deraining dataset. It includes 30 nighttime videos, each of which includes 200 frames. Each video contains synthetic rain streaks that generated by (Ranftl et al. 2020) and (Li, Cheong, and Tan 2019). We follow the protocol (Patil et al. 2022b) to evaluate the effectiveness of our method, i.e., 10 videos are used as the training set and the rest 20 videos are taken as the test set.

RealNightRain To evaluate the generalization ability of our method, we collect 10 real-world nighttime rain videos, where 5 videos contain dense rain streaks and the rest 5 videos contain sparse rain streaks. We will publicly release our collected dataset.

Implementation Details

Our NightRain includes three parts: pretraining, adaptive-rain-removal and adaptive-correction, each of which needs to train a video diffusion model. In each training step, we randomly sample $P \times K$ videos, where P is the number of videos and K denotes the number of clips of a video. Each clip consists of 4 frames with the image size is 64×64 . The Adam is used to optimize our model and the learning rate is set to 0.0002.

Pretraining The total training steps and $P \times K$ are set to 2,000,000 and 16×4 , respectively. To evaluate our method on SynNightRain, we use 25 sampling steps to gradually restore a rain video.

Adaptive-Rain-Removal Fig. 2 shows the pipeline of our adaptive-rain-removal. The number of iterations in confidence generation N , the threshold t_u and the total training steps are set to 3, 0.5 and 10,000, respectively. We apply two data augmentation approaches (Gaussian noise and pixel masking) to augment the inputs of the student model. The variance range of the Gaussian noise is (0, 0.2) and the masking ratio of pixels is set to 25%. The total training steps and $P \times K$ are set to 10,000 and 10×2 , respectively.

Adaptive-Correction Fig. 3 shows the pipeline of our adaptive-rain-removal. The total training steps, the number of clear videos, P and K are set to 10K, 10, 10 and 2, respectively.

Transformer-based Video Diffusion Model As shown in Figure 4, our video diffusion model consists of three modules, i.e., Patch Partition, Time Embedding and Transformer blocks. For the patch partition part, we set the patch size ($ts \times ss \times ss$) and the output channel C_p to 2 and 768, respectively. For the time embedding, we utilize two fully-connected layers to encode the information, where the output channels are set to 768. In our method, we set the number of transformer blocks to 10.

Methods	HRIR	DLF	RMFD	DSTFM	MetaRain	WeatherDiff	FDM	Ours
Venue	CVPR'19	CVPR'19	TPAMI'21	PR'22	ECCV'22	TPAMI'23	NeurIPS'22	
PSNR	16.83	15.71	16.18	17.82	22.21	20.98	23.49	26.73
SSIM	0.6481	0.6307	0.6402	0.6486	0.6723	0.6697	0.7657	0.8647

Table 1: Quantitative results on the nighttime deraining dataset. DSTFM (Patil et al. 2022a) and MetaRain (Patil et al. 2022b) are nighttime video deraining methods. WeatherDiff (Özdenizci and Legenstein 2023) and FDM (Harvey et al. 2022) are diffusion-based methods, where FDM is a video diffusion model.




Figure 7: Qualitative results on the SynNightRain dataset. MetaRain (Patil et al. 2022b) is a nighttime video deraining method. WeatherDiff (Özdenizci and Legenstein 2023) and FDM (Harvey et al. 2022) are diffusion-based methods, where FDM is a video diffusion method. Zoom in for better visualisation.

Evaluation on the SynNightRain Dataset

We compare our NightRain on the SynNightRain dataset with state-of-the-art (SOTA) methods, including HRIR (Li, Cheong, and Tan 2019), DLF (Yang, Liu, and Feng 2019), RMFD (Yang et al. 2021), DSTFM (Patil et al. 2022a), MetaRain (Patil et al. 2022b), FDM (Harvey et al. 2022) and WeatherDiff (Özdenizci and Legenstein 2023). HRIR, DLF, RMFD, DSTFM and MetaRain are CNN-based methods, while the rest methods are diffusion-based methods. The quantitative results of HRIR, DLF, RMFD and DSTFM are collected from (Yang et al. 2021).

Quantitative Evaluation The experimental results are shown in Table 1. It can be observed that the PSNR and SSIM of MetaRain are 22.21 and 0.6723, respectively. Our method achieves 26.73 of PSNR and 0.8647 of SSIM, which outperforms MetaRain by 4.52 dB and 0.1924, respectively. We also compare our method with existing state-of-the-art diffusion methods, i.e., FDM (Harvey et al. 2022) and WeatherDiff (Özdenizci and Legenstein 2023). All experiments use the same sampling steps for a fair comparison. Our NightRain outperforms these methods by 3.24 of PSNR. This is because existing methods are implemented by convolution neural networks and thus cannot take full advantage of global spatiotemporal information. In contrast, we develop a transformer-based diffusion model that fully utilizes global spatiotemporal cues. As a result, our NightRain can estimate more accurate noise maps, leading to superior performance.

Qualitative Evaluation Fig. 7 shows the quantitative results, including MetaRain (Patil et al. 2022b), FDM (Harvey et al. 2022), WeatherDiff (Özdenizci and Legenstein 2023) and our NightRain. We can observe that existing methods struggle to restore video details in dense rain streaks since these methods do not effectively model clear distributions

and capture global information. In contrast, our NightRain introduces a transformer-based diffusion model that captures global relations and models more comprehensive clear distributions. As a result, it effectively recovers details from adjacent pixels, leading to superior performance.

Evaluation on Real-world Datasets

Fig. 8 shows the qualitative results on real-world nighttime rain videos. It can be observed that the restored results of existing nighttime video deraining methods still have a few rain streaks left. This is because these methods are trained on synthetic datasets and thus suffer from the domain gap problem. In contrast, our NightRain not only removes nighttime rain streaks but also restores the clear background. The main reason is that our adaptive-rain-removal utilizes high-confidence predictions from real-world rain videos to improve the deraining ability of our model. By training on actual paired data, our model reduces the domain gap problem between synthetic and real-world rain streak distributions.

Another problem of existing methods is that their restored predictions are over-saturated or suffer from a color shift (refer Figure 1). For example, MetaRain (Patil et al. 2022b) over-suppresses the white light. The results of FDM (Harvey et al. 2022) and WeatherDiff (Özdenizci and Legenstein 2023) are too dark. This problem is caused by the domain gap between synthetic and real-world rain streak distributions. Our NightRain also relies on synthetic datasets to initialize our video diffusion model. However, we propose adaptive-correction to address these problems. Specifically, adaptive-correction forces our model to continually rectifies prediction errors during the training stage. As a result, our NightRain effectively removes rain streaks without introducing over-saturation or color-shift issues, resulting in



Figure 8: Qualitative comparison on real-world data. MetaRain (Patil et al. 2022b) is nighttime video deraining methods. FDM (Harvey et al. 2022) and WeatherDiff (Özdenizci and Legenstein 2023) are diffusion-based methods, where FDM is a video diffusion method. Zoom in for better visualisation.

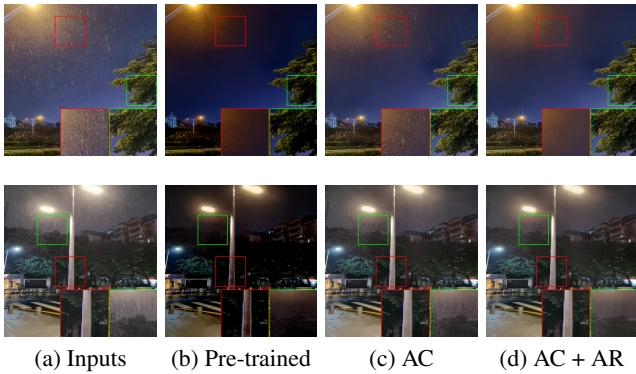


Figure 9: Ablation studies on real-world videos. (a) Input (b) Results obtained using the pre-trained model. (c) Results obtained using our adaptive-correction. (d) Results obtained using our adaptive-correction and adaptive-rain-removal. Zoom in for better visualisation.

more accurate detail recovery.

Ablation Studies

In this section, we conduct several ablation studies on real-world datasets to verify the effectiveness of our proposed processes: adaptive-rain-removal and adaptive-correction. Figure 9 displays the results of the pre-trained model, our adaptive-rain-removal, and adaptive-correction. Moreover, we use Hyper-IQA (Su et al. 2020), a non-reference metric, to quantitatively evaluate the performance. The Hyper-IQA scores (higher is better) for (a) Inputs, (b) Pre-trained, (c) AC, and (d) AC + AR are 0.6341, 0.6401, 0.6540, and 0.6791, respectively. It can be observed that predic-

tions obtained using the pre-trained model suffer from over-saturation and color shifts. Specifically, the results are excessively dark, hindering the visibility of object details. This is due to the significant domain gap between synthetic and real-world rain streak distributions.

As illustrated in Figure 9 (c), our adaptive-correction addresses over-saturation and color-shift issues by continually rectifying errors in our predictions. However, despite these improvements, some rain streaks remain in the predictions of our adaptive-correction. To tackle this, our adaptive-rain-removal utilizes unlabeled rain videos, enabling our model to remove rain streaks from real-world rain videos, especially in regions affected by complex light effects. Consequently, the final outputs of our NightRain not only eliminate rain streaks but also retain crucial video details.

Conclusion

We have presented NightRain, a novel nighttime video deraining method with adaptive-rain-removal and adaptive-correction processes. Our adaptive-rain-removal utilizes unlabeled rain videos to enable our model to enhance the deraining capability in real-world rain videos, particularly in regions affected by complex light effects. By learning high-confidence predictions, this process compels our model to progressively address more challenging regions, leading to effective real-world rain streak removal. Our adaptive-correction aims to address over-saturation and color shift problems. By learning from differences between clear night videos and their corresponding predictions, our adaptive-correction continually rectifies errors in our models' predictions, resulting in more precise restoration. Extensive experiments show that our NightRain achieves a significant performance improvement both quantitatively and qualitatively.

References

- Chen, X.; Li, H.; Li, M.; and Pan, J. 2023. Learning A Sparse Transformer Network for Effective Image Deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5896–5905.
- Fang, G.; Ma, X.; Song, M.; Mi, M. B.; and Wang, X. 2023. DepGraph: Towards Any Structural Pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16091–16101.
- Harvey, W.; Naderiparizi, S.; Masrani, V.; Weillbach, C.; and Wood, F. 2022. Flexible diffusion modeling of long videos. *NeurIPS*.
- Islam, M. R.; and Paul, M. 2021. Video rain-streaks removal by combining data-driven and feature-based models. *Sensors*, 21(20): 6856.
- Jiang, K.; Wang, Z.; Yi, P.; Chen, C.; Huang, B.; Luo, Y.; Ma, J.; and Jiang, J. 2020. Multi-scale progressive fusion network for single image deraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8346–8355.
- Kulkarni, A.; Patil, P. W.; and Murala, S. 2021. Progressive subtractive recurrent lightweight network for video deraining. *IEEE Signal Processing Letters*.
- Li, B.; Yan, T.; He, W.; and Zhu, X. 2023. STANet: a Spatial-Temporal Aggregation Network for Video Deraining. In *2023 3rd International Conference on Neural Networks, Information and Communication Engineering (NNICE)*, 470–479. IEEE.
- Li, R.; Cheong, L.-F.; and Tan, R. T. 2019. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1633–1642.
- Ma, X.; Fang, G.; and Wang, X. 2023. DeepCache: Accelerating Diffusion Models for Free. *arXiv:2312.00858*.
- Özdenizci, O.; and Legenstein, R. 2023. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Patil, P. W.; Gupta, S.; Rana, S.; and Venkatesh, S. 2022a. Dual-frame spatio-temporal feature modulation for video enhancement. *Pattern Recognition*, 130: 108822.
- Patil, P. W.; Gupta, S.; Rana, S.; and Venkatesh, S. 2022b. Video Restoration Framework and Its Meta-adaptations to Data-Poor Conditions. In *ECCV*, 143–160. Springer.
- Peebles, W.; and Xie, S. 2022. Scalable Diffusion Models with Transformers. *arXiv preprint arXiv:2212.09748*.
- Ranftl, R.; Lasinger, K.; Hafner, D.; Schindler, K.; and Koltun, V. 2020. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3): 1623–1637.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Shi, Z.; Feng, Y.; Zhao, M.; and He, L. 2020. A joint deep neural networks-based method for single nighttime rainy image enhancement. *Neural Computing and Applications*, 32: 1913–1926.
- Su, R.; Zhang, L.; Zhang, Y.; Xu, F.; Lu, K.; Tong, N.; and Li, F. 2023. Complex scene video frames alignment and multi-frame fusion deraining with deep neural network. *Neural Computing and Applications*, 35(7): 5369–5380.
- Su, S.; Yan, Q.; Zhu, Y.; Zhang, C.; Ge, X.; Sun, J.; and Zhang, Y. 2020. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3667–3676.
- Tang, Q.; Yang, J.; Liu, H.; and Guo, Z. 2020. A Modified Syn2Real Network for Nighttime Rainy Image Restoration. In *Advances in Visual Computing: 15th International Symposium, ISVC 2020, San Diego, CA, USA, October 5–7, 2020, Proceedings, Part II 15*, 344–356. Springer.
- Wang, H.; Wu, Y.; Li, M.; Zhao, Q.; and Meng, D. 2019. A survey on rain removal from video and single image. *arXiv preprint arXiv:1909.08326*.
- Wang, S.; Zhu, L.; Fu, H.; Qin, J.; Schönlieb, C.-B.; Feng, W.; and Wang, S. 2022a. Rethinking Video Rain Streak Removal: A New Synthesis Model and a Deraining Network with Video Rain Prior. In *ECCV*.
- Wang, X.; Chen, J.; Han, Z.; Zhu, Q.; and Ruan, W. 2022b. Real-Time Video Deraining via Global Motion Compensation and Hybrid Multi-Scale Temporal Correlations. *IEEE Signal Processing Letters*, 29: 672–676.
- Xue, X.; Ding, Y.; Mu, P.; Ma, L.; Liu, R.; and Fan, X. 2020. Sequential deep unrolling with flow priors for robust video deraining. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1813–1817. IEEE.
- Yan, W.; Tan, R. T.; Yang, W.; and Dai, D. 2021. Self-aligned video deraining with transmission-depth consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11966–11976.
- Yang, W.; Liu, J.; and Feng, J. 2019. Frame-consistent recurrent video deraining with dual-level flow. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1661–1670.
- Yang, W.; Tan, R. T.; Feng, J.; Wang, S.; Cheng, B.; and Liu, J. 2021. Recurrent multi-frame deraining: Combining physics guidance and adversarial learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 8569–8586.
- Yang, W.; Tan, R. T.; Wang, S.; and Liu, J. 2020. Self-learning video rain streak removal: When cyclic consistency meets temporal correspondence. In *CVPR*, 1720–1729.
- Zhang, K.; Li, D.; Luo, W.; Ren, W.; and Liu, W. 2022. Enhanced Spatio-Temporal Interaction Learning for Video Deraining: Faster and Better. *IEEE TPAMI*.
- Zhuang, J.; Luo, Y.; Zhao, X.; Jiang, T.; and Guo, B. 2022. UConNet: Unsupervised Controllable Network for Image and Video Deraining. In *Proceedings of the 30th ACM International Conference on Multimedia*, 5436–5445.