

Unified Removal of Raindrops and Reflections: A New Benchmark and A Novel Pipeline

Xingyu Liu^{1,2}, Zewei He^{1,2}[✉], Yu Chen^{1,2}, Chunyu Zhu³, Zixuan Chen⁴,
Xing Luo², and Zhe-Ming Lu^{1,2}

¹ Huanjiang Laboratory, Zhuji, China

² School of Aeronautics and Astronautics, Zhejiang University, Hangzhou, China

³ Hangzhou Institute of Technology, Xidian University, Hangzhou, China

⁴ The Chinese University of Hong Kong, Hong Kong, China

{xingyu_liu, zeweihe}@zju.edu.cn

Project page: <https://xingyuliu00.github.io/diffur3/>

Abstract. When capturing images through glass surfaces or windshields on rainy days, raindrops and reflections frequently co-occur to significantly reduce the visibility of captured images. This practical problem lacks attention and needs to be resolved urgently. Prior de-raindrop, de-reflection, and all-in-one models have failed to address this composite degradation. To this end, we first formally define the unified removal of raindrops and reflections (UR³) task for the first time and construct a real-shot dataset, namely RainDrop and ReFlection (RDRE), which provides a new benchmark with substantial, high-quality, diverse image pairs. Then, we propose an effective diffusion-based framework (i.e., DiffUR³) with several target designs to address this challenging task. By leveraging the powerful generative prior, DiffUR³ successfully removes both types of degradations. Extensive experiments demonstrate that our method achieves state-of-the-art performance on our benchmark and on challenging in-the-wild images.

Keywords: Raindrop removal · Reflection removal · New benchmark

1 Introduction

On rainy days, raindrops and reflections frequently co-occur in autonomous driving scenarios, posing challenges for onboard visual recognition systems or vehicle cameras during recording. Adherent raindrops and the reflections from camera side significantly reduce the visibility of captured images [44], and may lead to severe driving safety hazards. Therefore, Unified Removal of Raindrops and Reflections (UR³) is a practical problem that requires urgent solution.

However, this task has not received significant attention and enough research. We notice that the physical models of raindrop (i.e., $\mathbf{I} = (1 - \mathbf{A}) \odot \mathbf{B} + \mathbf{A} \odot \mathbf{R}_d$ from [33]) and reflection (i.e., $\mathbf{I} = (1 - \mathbf{W}) \odot \mathbf{B} + \mathbf{W} \odot \mathbf{R}_f$ from [49]) share a

[✉] Corresponding author (zeweihe@zju.edu.cn).




Fig. 1: We compare our DiffUR³ pipeline with other methods on low-quality images with raindrops and reflections from our newly collected real-world benchmark. Specifically, (c) DAI [16] is designed for reflection removal, (d) A cascaded method, and (e) A re-trained all-in-one method (i.e., Histoformer [38] and † indicates re-trained on our dataset), (f) Our DiffUR³ pipeline jointly removes both degradations in a single pass.

similar multiplicative compositing structure. $\mathbf{I}$ and $\mathbf{B}$ denote the observed image and the background scene, respectively. $\mathbf{R}_d$ and $\mathbf{R}_r$ indicate the raindrop layer and the reflection layer, respectively. $\mathbf{A}$ is the transparent matrix and $\mathbf{W}$ is the reflective amplitude coefficient map. When both degradations co-occur, the two processes are applied sequentially. This yields a nested compositing model: $\mathbf{I} = (1 - \mathbf{A}) \odot [(1 - \mathbf{W}) \odot \mathbf{B} + \mathbf{W} \odot \mathbf{R}_r] + \mathbf{A} \odot \mathbf{R}_d$. By expanding this expression, the cross-term $\mathbf{A} \odot \mathbf{W} \odot \mathbf{R}_r$ couples both degradations. This means the raindrop-related coefficient $\mathbf{A}$ does not merely occlude the clean background $\mathbf{B}$, but also modulates the reflection component $\mathbf{W} \odot \mathbf{R}_r$. Therefore, the raindrop degradation and the reflection degradation are mutually coupled and interacted with each other. They do not simply superimpose upon each other. The model designed for reflection removal can also partially eliminate raindrops (see in Fig. 1 (c)), and vice versa. Previously, researchers treated raindrop removal and reflection removal as two separate tasks [16, 19, 31, 33, 35, 47]. Though these methods can achieve relatively good performance in removing the target type of degradation (i.e., raindrop or reflection) from a single image, they often fail to remove both types at the same time (see in Fig. 1 (c)). In addition, arranging raindrop removal and reflection removal in a cascaded manner essentially overlooks the correlation (see in Fig. 1 (d)). In this work, we aim to finding a practical solution capable of simultaneously eliminating this raindrop-reflection composite degradation, thereby enhancing the clarity of captured images. We hope this endeavor can provide useful support for applications such as autonomous driving, photography, and video surveillance [8, 9, 51, 52].

UR³ is a fundamental but complex task. Its key challenges lie in the following three aspects. (1) Lack of data: substantial image pairs are required for training. Data synthesis represents a potential solution, yet the gap between synthetic and real-world data cannot be ignored. Existing real-captured datasets [31, 53] normally contain only a single degradation. Currently, there is no publicly available dataset with the raindrop-reflection composite degradation. (2) Task interac-

tion: treating raindrop removal and reflection removal as two isolated tasks often yields sub-optimal performance. Recent all-in-one methods also demonstrate poor generalization performance (see in Fig. 1 (e)). A new pipeline is required for this task. (3) Void information: background scene is completely lost in regions with exceptionally large/dense raindrops or intense reflections. These occluded regions are extremely challenging, somewhat analogous to the inpainting task.

Trying to address these challenges, we first set up an image acquisition platform to collect corresponding data for addressing the UR³ task. We collect a substantial number of image pairs to constitute our **RainDrop** and **ReFlection** (RDRF) dataset. Within each pair, one features a clean image and the other is a degraded image containing both raindrops and reflections. To the best of our knowledge, RDRF is the first real-shot dataset with raindrop-reflection composite degradation. We hope RDRF dataset can contribute to the advancement and development of the UR³ task, and benefit the entire community.

Then, we focus on proposing a qualified baseline for this new UR³ task. By leveraging the powerful generative prior, we design an effective multi-condition-controlled diffusion framework (i.e., DiffUR³) to jointly remove both kinds of degradations. To the best of our knowledge, DiffUR³ is the first pipeline for the UR³ task, which contains several target designs. (1) A Modulate&Gate module is proposed to step-wisely align each condition with the noisy latent and adaptively select the effective components in the latent space. This simple yet effective module can enhance the control signals by modulating and mining the condition latents. (2) We train an additional Fidelity Encoder to correct the distortions caused by the compression operation in the VAE encoder.

2 Related work

2.1 Raindrop removal

In the realm of raindrop removal, recent studies have explored diverse methodologies. Eigen *et al.* [11] pioneered single-image raindrop removal using CNNs. Qian *et al.* [31] introduced a generative adversarial network (GAN) to enhance raindrop removal. Transformer-based approaches like IDT [43], UDR-S²Former [4] and Histoformer [38] have achieved superior performance. Meanwhile, the CCN [32] adopts a unique approach by employing neural architecture search. More recently, diffusion-based methods like WeatherDiff [29] and T³-DiffWeather [5] have emerged, leveraging the generative capabilities of diffusion models to enhance raindrop removal.

2.2 Reflection removal

In the field of single image reflection removal, various advanced techniques have been proposed to address the ill-posed nature of separating superimposed transmission and reflection layers. Early methods such as CEILNet [12] leverage edge information and deep learning. IBCLN [23] introduces a cascaded refinement




Fig. 2: (a) Sketch diagram and actual equipment of our image acquisition platform. To suppress shutter-induced micro-vibrations which may potentially induce image misalignment, we implement a wireless triggering mechanism. It comprises a remote controller and a camera-mounted signal receiver, enabling contact-free shutter operation. (b) The data collection pipeline for our RDRF dataset. ✕ denotes light occlusion

strategy to iteratively enhance the estimations of transmission and reflection layers. More recent advancements, YTMT [17], DSRNet [18], DSIT [19], RDNet [47], and GFRRN [7] employ dual-stream networks to enhance feature interaction. Further more, diffusion-based models like L-Differ [15] and DAI [16] also show their capabilities across a wide range of real-world scenarios.

2.3 Datasets

Existing real-world datasets for raindrop removal include AGAN [31], RainDS [32], RobotCar-Rainy [30], and Raindrop Clarity [20], these datasets provide low-quality images with raindrops and their corresponding ground truth images. Differently, Windshield [36] contains degraded images along with their corresponding binary masks that indicate the raindrop-affected areas. For reflection removal task, it is noteworthy that synthetic data is commonly employed for training. Recently, some real-world datasets have been proposed, such as RRW [53] and DRR [16]. However, there are no existing datasets that specifically address the unified removal of raindrops and reflections, which is the focus of our work.

3 RDRF dataset

Similar to most deep learning based methods, our task (i.e., unified removal of raindrops and reflections) requires a large number of degraded images with corresponding clean labels for training. There are no existing training or testing datasets for this new task. As shown in Fig. 2 (a), we set up an image acquisition platform to collect our own **R**ain**D**rop and **R**e**F**lection (RDRF) dataset. In our case, a substantial volume of image pairs are required, where each pair comprises



Fig. 3: Our RDRF dataset comprises a diverse collection of scenes, each contains a ground truth and multiple low-quality images. As illustrated in this figure, the clean ground truths are highlighted in red boxes, while corresponding low-quality images are arranged around. We divide it into the training and testing subsets, ensuring no overlapping samples between them. Please zoom in on screen for a better view.

two images with the identical background scene, yet one has a clean foreground and the other is corrupted by raindrops and reflections.

3.1 Hardware

Drawing inspirations from previous works of Zhu *et al.* [53] and Li *et al.* [24], our RDRF dataset is captured under real scenarios deliberately constructed in controlled environments. For the hardware configuration, the camera is mounted on a tripod using an adjustable base, with the glass slab positioned in front of the lens. We connect a signal receiver onto the camera, thereby enabling remote control of the shutter. This wireless triggering mechanism can effectively avoid image misalignment caused by camera vibrations resulting from manual operation. To ensure diversity, neither the camera nor the glass is fixed. They can be adjusted to simulate different shooting situations (e.g., camera-to-glass distances/angles). In addition, we utilize two cameras (Sony ILCE-7RM4A and Nikon D7100) with zoom lens and choose different glass thicknesses (3 mm, 5 mm, and 8 mm) to further enhance diversity.

3.2 Data collection pipeline

Fig. 2 (b) exhibits our data collection pipeline. For step 1, we utilize a light-blocking box to suppress the reflections from the camera side (✕ denotes light occlusion). The obtained image is regarded as the ground-truth. For step 2, we keep the background scenario and camera unchanged. The light-blocking box is removed and the raindrops are created by spraying water onto the glass surface. By randomly rotating the glass at different angles, we create varying reflections

with different scenes and intensities. For each scene, multiple images are captured as the low-quality ones.

Some samples are illustrated in Fig. 3. As demonstrated in the dataset, our RDRF dataset comprises a comprehensive collection of scenes. The raindrops are captured under diverse shapes and sizes (circular, elliptical, and irregular), ranging from sparse to dense. Raindrop flow traces are also included. In addition, the reflections are also captured with diverse reflection scenes, ranging from weak to strong. All the images are captured in 4752×3168 resolution to ensure high-quality. In total, our RDRF dataset consists of 307 unique scenes. **The category distribution diagram can be found in the supplementary material.** It is divided into a training set (216 scenes with 9003 image pairs) and a testing set (91 scenes with 277 image pairs). Note that, we capture 1 to 5 image pairs for each scene in the testing set.

To further address the spatial misalignment caused by our hardware, we follow the procedures proposed in [40]. It starts by extracting SIFT [27] key-points and descriptors, which are matched with L2 distance. Using the matched key-points, a homography matrix is estimated via RANSAC [13] to handle outliers and find a robust geometric transformation. The low-quality image is aligned to the ground truth by applying a perspective warp using the computed homography.

Previously, some valuable datasets are collected to address the raindrop removal task (e.g., Qian’s dataset [31], RainDS [32], Raindrop Clarity [20]). Their brilliant datasets has different focus with ours. Some of their authors also noticed the reflection artifacts [20, 31] and tried to avoid them during data collection stage. Even so, there are still some reflection artifacts present in Qian’s dataset [31]. This also demonstrates that reflections and raindrops tend to occur simultaneously. UR³ is a practical task which required urgent solution. Our RDRF dataset represents the first-of-its-kind contribution to the UR³ task.

4 Methodology

Our RDRF dataset provides sufficient and diverse training data for UR³ task. Formally, given a low-quality (LQ) image $I_{lq} \in \mathbb{R}^{3 \times H \times W}$ with both raindrops and reflections on it, a straightforward idea is to employ conventional restoration methods [6, 38] to directly learn the mapping function from the low-quality to the ground-truth. However, their results are perceptually unsatisfying, because of the complexity of UR³ task. Instead, we try to utilize the powerful generative priors of the diffusion model as an effective solution. In this way, UR³ is regarded as a conditional image generation problem.

Specifically, we design a two-stage network (i.e., DiffUR³) shown in Fig. 4 (a), which consists of **(I) a restoration stage**, and **(II) a multi-condition generation stage**. The restoration stage outputs an initial result I_s , which is then used as the condition image in the following generation stage.

Compared with DiffBIR framework, our DiffUR³ has two major different designs: (1) Considering the fact that raindrops and reflections merely affect




Fig. 4: (a) Overall pipeline of our DiffUR³ framework. Given a low-quality image I_{lq} , the restoration stage removes the undesired degradation to obtain the initial result I_s . Both I_{lq} and I_s are fed into the next stage as the condition images. We inject the **effective** condition information through a control branch, which outputs control signals for the noise prediction U-Net. (b) Details of the Modulate&Gate module within the control branch. (c) The generation of noisy latent z_t during the training phase. Note that the noisy latent starts from random Gaussian noise during the inference.

certain regions, some parts in the LQ image I_{lq} can be regarded as clean. We also integrate I_{lq} as one of the conditions during stage II. A Modulate&Gate module is designed to adaptively fuse the effective information from I_{lq} and I_s . (2) We design an additional Fidelity Encoder to correct the distortions caused by the compression in VAE. Details of our DiffUR³ pipeline are described below.

4.1 Restoration stage

In the first stage, our aim is to remove some easy yet undesired degradations from LQ input I_{lq} . The output image I_s provides a reliable condition image for training the generation stage.

$$I_s = \mathcal{RM}(I_{lq}), \quad (1)$$

where $\mathcal{RM}(\cdot)$ denotes the restoration model. In our implementation, we select the DRSformer [6] as the restoration model in stage I due to its superior performance and generalization capability.

4.2 Multi-condition generation stage

Our multi-condition generation stage is based on the Stable Diffusion Model [34], because its powerful generative prior can facilitate the restoration of regions that are challenging to be recovered in stage I through a conditional generation approach. To achieve better efficiency and stabilized training, the pretrained VAE [22] encoder $\mathcal{E}$ is employed to encode the condition images into the latent

space. Both diffusion and denoising processes are performed in this space instead of the pixel space. The main denoising network is a pretrained U-Net. The denoising output is then converted back to the pixel space using the pretrained VAE decoder $\mathcal{D}$.

As mentioned above, for UR³ task we argue LQ image I_{lq} contains some clean information within certain regions¹. Therefore, both I_{lq} and I_s are encoded by the VAE encoder $\mathcal{E}$:

$$c_{lq}, c_s = \mathcal{E}(I_{lq}, I_s), \quad (2)$$

where $c_{lq} \in \mathbb{R}^{4 \times \frac{H}{8} \times \frac{W}{8}}$ and $c_s \in \mathbb{R}^{4 \times \frac{H}{8} \times \frac{W}{8}}$ denote the obtained condition latent from I_{lq} and I_s , respectively. Besides, the noisy latent z_t is also embedded, since it has been proven to enhance image quality [25]. The generation of z_t is shown in Fig. 4 (c).

Similar to previous work [25], we also inject the condition information via a control branch. We make a trainable copy of the pretrained U-Net encoder and middle block (i.e., $\mathbf{F}_{con}$ in Fig. 4 (a)), which receives condition information and then outputs control signals. A normal solution is to add or concatenate c_{lq} , c_s and z_t before sending to $\mathbf{F}_{con}$ [3, 29]. However, we observe that the noisy latent z_t varies at different time steps, yet the condition latent (i.e., c_{lq} or c_s) remains unchanged. Instead of direct addition or concatenation, we propose a more reasonable solution to modulate c_{lq} and c_s through z_t . Since there are more than one condition latent, a gate mechanism is introduced to adaptively assign different spatial weights to c_{lq} and c_s .

To this end, before entering the $\mathbf{F}_{con}$, we design a Modulate&Gate module which consists of a Modulate block and a Gate block. Fig. 4 (b) shows the details of our Modulate&Gate module. We describe them as below.

Modulate block Take c_s as an example, c_{lq} can be similarly derived. First, both c_s and z_t individually pass through a convolutional layer to extract their features $f_c \in \mathbb{R}^{C \times \frac{H}{8} \times \frac{W}{8}}$ and $f_z \in \mathbb{R}^{C \times \frac{H}{8} \times \frac{W}{8}}$. C denotes the channel number of the extracted feature. In our implementation, we set $C = 32$. Then, their concatenation result is fed into two consecutive transformer layers [39] to perform the cross attention operation, which can facilitate the information interaction between f_c and f_z . Our cross attention operation aligns the dimensions of the output f_{cross} and f_z at the end. Finally, we add $f_{cross} \in \mathbb{R}^{C \times \frac{H}{8} \times \frac{W}{8}}$ with f_z , and employ another convolutional layer to reduce the channel number back to 4. The formulations are as follows:

$$\begin{aligned} f_c, f_z &= \text{Conv}(c_s, z_t), \\ f_{cross} &= \text{CrAttn}([f_c, f_z]), \\ \hat{c}_s &= \text{Conv}(f_{cross} + f_z), \end{aligned} \quad (3)$$

¹ Unlike blind super-resolution and blind image denoising in [25], where the entire LQ image is degraded.

where $Conv(\cdot)$ denotes the convolutional layer, $CrAttn(\cdot)$ denotes the cross attention operation, $[\cdot, \cdot]$ denotes the concatenation, $\hat{c}_s$ is the modulated condition latent, and $\hat{c}_{lq}$ can be derived by replacing c_s with c_{lq} in Eqn. 3.

Gate block After obtaining $\hat{c}_s$ and $\hat{c}_{lq}$, we need to selectively extract the components that are beneficial to our DiffUR³. We concatenate them together, and then send to a spatial attention to generate a spatial weight $\alpha \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s}}$. The spatial attention operation consists of two convolutional layers, one activation layer, and one sigmoid layer. The formulations are as follow:

$$\begin{aligned}\alpha &= SpAttn([\hat{c}_s, \hat{c}_{lq}]), \\ \bar{c}_s &= \alpha \cdot \hat{c}_s, \\ \bar{c}_{lq} &= (1 - \alpha) \cdot \hat{c}_{lq},\end{aligned}\tag{4}$$

where $SpAttn(\cdot)$ denotes the spatial attention operation, $\bar{c}_s$ and $\bar{c}_{lq}$ are the output condition latent variables. Note that our Modulate&Gate module is simple yet effective. More sophisticated designs can be considered for better performance, which is not the focus of this work. We concatenate $\bar{c}_s$, $\bar{c}_{lq}$, and the noisy latent z_t together, and send them to $\mathbf{F}_{con}$ for generating the control signals, which are added to the denoising U-Net via zero convolutions [45]. At each time step, the noise prediction U-Net estimates the noise component and performs denoising on the noisy latent z_t . During the inference phase, the noisy latent starts from random Gaussian noise and iteratively passes through the pretrained U-Net to estimate the clean latent $\hat{z}_0$.

Fidelity encoder By decoding $\hat{z}_0$ back to the pixel space, we can obtain the final output image. As shown in Fig. 5, we observe that while stable diffusion demonstrates superior performance in recovering the degradations caused by raindrops or reflections (green box of Fig. 5 (b)), some unwanted distortions are introduced by the compression operation in the VAE encoder (red box of Fig. 5 (b)). The stable diffusion model lacks the ability to correct these kinds of distortions. To deal with this issue and improve the fidelity of the generated results, we train an additional fidelity encoder (FE) inspired by [1]. The FE is proposed to extract multi-scale features from the initial result I_s and the LQ image I_{lq} , which are not affected by the down-sampling, for preserving local structural semantics.

Fig. 6 shows the training pipeline of our fidelity encoder, which shares the same architecture with VAE encoder [22] (besides the first convolutional layer). To keep consistent with the control branch, both LQ image I_{lq} and initial result I_s are fed into the fidelity encoder through a Gate block to extract faithful features. The extracted fidelity features are added to corresponding positions in VAE Decoder via zero convolutions. Then, we encode the ground truth I_{gt} via the pretrained VAE encoder to latent space, simulating the denoised latent i.e., z_0 . Finally, guided by the fidelity features, pretrained VAE decoder [22] converts the compressed latent z_0 to a reconstructed image $\hat{I}_{gt}$. The whole pipeline is trained by minimizing a mean absolute error (i.e., L_1 loss) between $\hat{I}_{gt}$ and I_{gt} .




Fig. 5: The motivation of employing an additional fidelity encoder.




Fig. 6: Training pipeline of our fidelity encoder.

Color correction Diffusion models can occasionally exhibit color shifts [10]. To tackle this issue, some simple techniques (e.g., color normalization, wavelet color correction) have been explored before [42]. We experimentally choose color normalization to align the mean and variance of the output with color reference I_s , and it is conducted patch-wisely. Note that, I_{lq} is not suitable here since the composite degradation may cause false correction.

5 Experimental results

5.1 Implementation details and metrics

For the training of stage I, we adopt DRSformer [6] as the restoration model. Compared with the original, we make certain simplifications to accelerate the computation in Stage I. $\{N_0, N_1, N_2, N_3, N_4\}$ are set to $\{0, 2, 4, 4, 8\}$, and the number of attention heads for sparse transformer blocks (i.e., STBs) in level 1 to level 4 are set to $\{1, 2, 4, 8\}$. The initial channel C is set to 16. Patches of size 256×256 are randomly cropped from the RDRF-training dataset, and horizontal and vertical flips are applied as the data augmentation techniques. This model is trained with an initial learning rate of $3e^{-4}$ for the first 100K iterations, which will gradually reduced to $1e^{-6}$ using cosine annealing schedule [14] during the remaining 200K iterations.

For stage II, the images in our RDRF-training dataset are randomly cropped into 640×640 patches. Horizontal and vertical flips, resizing, rotation are applied as the data augmentation techniques. The control branch is trained with a batch size of 40 and with a fixed learning rate of $1e^{-4}$ for entire 50K iterations. To accelerate the sampling process, we adopt a spaced DDPM sampling schedule [28] which requires 50 sampling steps. For the training of the additional fidelity encoder, we follow the training settings of stage II, except the batch size and number of iterations.

Note that the images in our RDRF-training dataset are firstly resized to a fixed resolution of 1080×720 , and we use AdamW [26] optimizer with default settings for all the training procedures. We train the restoration model in stage I for 300k iterations (batch size = 4) on a single A6000 GPU. Then we adopt the Stable Diffusion 2.1-base [34] as the generative prior, and train the control

branch in stage II for 50k iterations (batch size = 40) on two A6000 GPUs. The fidelity encoder is trained for 300k iterations (batch size = 6) on two A6000 GPUs. We adopt three traditional metrics (PSNR, SSIM, LPIPS [46]) and three no-reference image quality assessment metrics (MUSIQ [21], CLIPQA+ [41], HyperIQA [37]) to comprehensively evaluate our performance.

5.2 Ablation study and discussions

Modulate&Gate module First, we perform ablation study to validate the effectiveness of Modulate&Gate module. We employ the naive diffusion-based method which solely adopts I_{lq} as the condition image (similar to [29]) and denote it as **Baseline 1**. In addition, **Baseline 2** means only the I_s is regarded as the condition image (similar to [25]). Table 1 summarizes the quantitative results. **w/o Modulate&Gate** Module means both I_{lq} and I_s are embedded as the condition images and fused by channel-wise concatenation in latent space. We observe that our Modulate&Gate module is critically important for our DiffUR³, as omitting this component leads to an obvious performance drop. Note that, both Modulate and Gate sub-modules are helpful for improving the performance.

Table 1: Ablation study on Modulate&Gate module.

Model	Condition	PSNR↑	SSIM↑	LPIPS↓
Baseline 1	I_{lq}	28.33	0.9271	0.0944
Baseline 2	I_s	28.54	0.9273	0.0943
w/o Modulate&Gate	I_{lq}, I_s	28.60	0.9294	0.0899
w/o Modulate	I_{lq}, I_s	29.11	0.9355	0.0832
w/o Gate	I_{lq}, I_s	29.12	0.9344	0.0841
DiffUR ³	I_{lq}, I_s	29.41	0.9372	0.0813

Table 2: Ablation study on Fidelity Encoder.

Model	PSNR↑	SSIM↑	LPIPS↓
Baseline	26.95	0.8042	0.1082
+CFW($w = 0.5$)	28.78	0.9120	0.0909
+CFW($w = 1.0$)	29.18	0.9320	0.0832
+FE (DiffUR ³)	29.41	0.9372	0.0813

Further, an in-depth analysis is provided in Fig. 7. **Baseline 1** outputs occasionally exhibit generation errors due to the inherent characteristics of diffusion model. In contrast, I_s is a better condition to avoid errors. However, some restoration artifacts in I_s may influence the **Baseline 2** outputs. They demonstrate distinct advantages across different regions. By considering I_{lq} and I_s together, simple channel-wise concatenation fails to systematically integrate their complementary strengths. The introduced Modulate&Gate module enables adaptive integration of information from dual condition images, thereby enhancing the model performance.

Fidelity Encoder Then, we perform ablation study to validate the effectiveness of the Fidelity Encoder (FE). **Baseline** means directly decoding $\hat{z}_0$ to pixel space without fidelity features. We observe a significant decline in fidelity, since the compression operation within the VAE encoder introduces considerable distortions (Fig. 5 (b)). In addition, the Controllable Feature Wrapping (CFW) [42,50],




Fig. 7: In-depth analysis on the function of our Modulate&Gate (M&G) module.

which leverages the encoder features to modulate corresponding decoder features for fidelity improvement, is also employed for comparison. In CFW, there is a adjustable hyper-parameter w to control the fidelity. We choose $w = 0.5$ and $w = 1.0$. The experimental results are listed in Table 2, and both CFW and FE can boost the fidelity. Our FE ranks first in terms of PSNR, SSIM, and LPIPS. Some visual results are shown in Fig. 8. The texts recovered by our FE exhibit less distortions. Note that, the training of our fidelity encoder is independent of both restoration and control branch, enabling its flexible application to various stable diffusion architectures.



Fig. 8: Visual results with CFW and FE.

Color correction We experimentally test the color normalization and wavelet color correction in Fig. 9. It can be observed that the wavelet method mistakenly introduces some artifacts from color reference I_s in the roof region. The color normalization corrects the color shift to make the result closer to the ground truth. Our experimental results are inconsistent with [42]. Since wavelet-based method introduces the low-frequency part from I_s , where may contain low-frequency artifacts caused by raindrops or reflections. Therefore, we adopt the color normalization instead of the wavelet color correction in our case.

5.3 Comparisons with state-of-the-art methods

Since this work is the first exploration for unified removal of raindrops and reflections (UR³) task. There are no prior methods. We employ two classical




Fig. 9: Color correction results with different methods.




Fig. 10: User study on RDRF-testing.

raindrop removal methods (i.e., AGAN [31], UMAN [35]), three classical reflection removal methods (i.e., RDNet [47], DSIT [19], DAI [16]), four cascaded methods, and four all-in-one methods (i.e., Histoformer [38], WeatherDiff₆₄ [29], DiffUIR [48], UniRestore [2]) as the competitors. We adopt their public models for calculating metrics. In addition, we re-train the all-in-one methods by using their published codes.




Fig. 11: Visual results of various methods on our RDRF-testing. Superscript [†] means this method is re-trained on our RDRF-training dataset. Please check and zoom in on screen for a better view.

We conduct a user study to evaluate our DiffUR³ subjectively against five methods. Specifically, we randomly select 50 images from our RDRF-testing dataset and invite 20 experts with image restoration background as volunteers.

Table 3: Benchmark results on our RDRF-testing dataset. We report PSNR, SSIM, LPIPS and three no-reference image quality assessment metrics (i.e., MUSIQ, CLIP-IQA+, HyperIQA) to perform comprehensive comparisons. The **bold** and underline indicate the best and second best. Superscript † means re-trained on RDRF.

Type	Method	Venue	RDRF-testing					
			PSNR↑	SSIM↑	LPIPS↓	MUSIQ↑	CLIP-IQA+↑	HyperIQA↑
Raindrop removal	AGAN [31]	CVPR2018	25.18	0.9209	0.1136	73.39	0.6241	0.6650
	UMAN [35]	TIP2021	25.61	0.9222	0.1111	73.23	0.6262	0.6525
Reflection removal	RDNet [47]	CVPR2025	26.79	0.9169	0.1269	70.68	0.5983	0.6260
	DSIT [19]	NeurIPS2024	26.41	0.9128	0.1339	71.29	0.6058	0.6248
	DAI [16]	AAAI2026	27.48	0.9251	0.1058	73.25	0.6335	0.6518
Cascaded	AGAN+RDNet	-	26.68	0.9291	0.0973	73.87	0.6550	0.6787
	RDNet+AGAN	-	26.64	0.9268	0.0991	73.87	0.6453	0.6737
	AGAN+DSIT	-	26.68	0.9283	0.0984	<u>74.05</u>	<u>0.6661</u>	0.6760
	DSIT+AGAN	-	26.49	0.9243	0.1046	73.89	0.6536	0.6725
	UMAN+DAI	-	27.50	0.9288	0.0973	74.01	0.6655	<u>0.6789</u>
	DAI+UMAN	-	27.44	0.9284	0.0966	73.95	0.6653	0.6766
All-in-One	Histoformer [38]	ECCV2024	25.74	0.9217	0.1214	72.34	0.5955	0.6450
	WeatherDiff ₆₄ [29]	TPAMI2024	24.45	0.8936	0.1313	72.39	0.6383	0.6412
	DiffUIR [48]	CVPR2024	23.32	0.8777	0.2190	68.39	0.5268	0.5901
	UniRestore [2]	CVPR2025	22.77	0.8291	0.2952	66.71	0.5459	0.5123
Re-trained All-in-One	Histoformer [†]	-	28.31	0.9315	0.0955	72.19	0.6221	0.6507
	WeatherDiff ₆₄ [†]	-	26.11	0.9130	0.1140	72.33	0.6471	0.6225
	DiffUIR [†]	-	28.39	0.9300	0.0991	71.83	0.6216	0.6311
	UniRestore [†]	-	25.66	0.8858	0.1631	73.01	0.6350	0.5994
	Stage I	-	<u>28.82</u>	<u>0.9354</u>	<u>0.0925</u>	73.48	0.6342	0.6648
Ours	DiffUR ³	-	29.41	0.9372	0.0813	74.72	0.6705	0.7046

For every image, each expert is asked to compare the result of our DiffUR³ with the alternatives one by one. For each comparison, the observers are demanded to choose the favored one after at least 10 seconds of observation. Afterward, we statistic the percentage of certain method to be selected. The statistical results in Fig. 10 indicates that our DiffUR³ is more favored by the invited experts.

Table 3 shows the quantitative results on our RDRF-testing dataset. Note that, our DiffUR³ ranks top among all metrics. Though the output from stage I exhibits relatively higher PSNR and SSIM values, it still suffers from certain distortions and residual degradations (Fig. 11 (e)), resulting in unsatisfactory visual quality and lower no-reference scores. In contrast, stage II and Fidelity Encoder effectively refines the output and produces significantly higher-quality results.

In addition, some visual comparisons of our DiffUR³ and the competitors are provided in Fig. 11. It is worth mentioning that the results of our DiffUR³ are closer to the ground truth with less degradation residuals and artifacts than the alternatives. We further evaluate the generalization capability of our DiffUR³ on a raindrop-only dataset [31] (single degradation situation). **More details are provided in the supplementary material.** Since real-world driving scenarios in rainy weather are very challenging, we capture some testing images to form




Fig. 12: Visual results of various methods on our RDRF-wild dataset. Please check and zoom in on screen for a better view.

a RDRF-wild dataset. Fig. 12 shows the visual comparisons on RDRF-wild. For this situation, the whole scene is more complicated than the controlled one. Our DiffUR³ can still obtain appealing results. This indicates the robustness and generalization capability of our approach in real-world applications. **More visual results can be found in the supplementary material.**

6 Conclusion

This work introduces a pioneering approach to the challenging task of UR³. By establishing the first dedicated RDRF dataset and proposing an effective diffusion-based framework (DiffUR³), we successfully address the limitations of previous methods that treat raindrop and reflection removal as separate tasks. Our two-stage pipeline, incorporating a restoration stage and a multi-condition generation stage, effectively leverages generative priors to remove both types of degradations simultaneously. Extensive experiments demonstrate the superiority of our approach over SOTA methods. The RDRF dataset and DiffUR³ framework contribute significantly to the advancement of the UR³ task, offering valuable resources for future research.

Acknowledgment

This work was supported in part by National Natural Science Foundation of China under Grant No. 52305590, Zhejiang Provincial Natural Science Foundation of China under Grant No. LQ24F010004, and Tianshan Talent Cultivation Plan - Science and Technology Innovation Team Project under Grant No. 2024TSYCTD0011.

References

1. Chang, Z., Weng, S., Zhang, P., Li, Y., Li, S., Shi, B.: L-CAD: Language-based Colorization with Any-level Descriptions using Diffusion Priors. In: NeurIPS. vol. 36, pp. 1–13 (2023)

2. Chen, I., Chen, W.T., Liu, Y.W., Chiang, Y.C., Kuo, S.Y., Yang, M.H., et al.: Unirestore: Unified perceptual and task-oriented image restoration model using diffusion prior. In: CVPR. pp. 17969–17979 (2025)
3. Chen, J., Pan, J., Dong, J.: FaithDiff: Unleashing Diffusion Priors for Faithful Image Super-resolution. In: CVPR. pp. 28188–28197 (2025)
4. Chen, S., Ye, T., Bai, J., Chen, E., Shi, J., Zhu, L.: Sparse Sampling Transformer with Uncertainty-Driven Ranking for Unified Removal of Raindrops and Rain Streaks. In: ICCV. pp. 13060–13071 (2023)
5. Chen, S., Ye, T., Zhang, K., Xing, Z., Lin, Y., Zhu, L.: Teaching Tailored to Talent: Adverse Weather Restoration via Prompt Pool and Depth-Anything Constraint. In: ECCV. pp. 95–115 (2024)
6. Chen, X., Li, H., Li, M., Pan, J.: Learning A Sparse Transformer Network for Effective Image Deraining. In: CVPR. pp. 5896–5905 (2023)
7. Chen, Y., He, Z., Liu, X., Chen, Z., Lu, Z.: Gfrn: Explore the gaps in single image reflection removal. In: CVPR. pp. 5690–5699 (2026)
8. Chen, Z., He, Z., Lu, Z.M.: DEA-Net: Single Image Dehazing Based on Detail-Enhanced Convolution and Content-Guided Attention. *IEEE Transactions on Image Processing* **33**, 1002–1015 (2024)
9. Chen, Z., He, Z., Lu, Z., Sun, X., Lu, Z.M.: Prompt-based test-time real image dehazing: a novel pipeline. In: ECCV. pp. 432–449 (2024)
10. Choi, J., Lee, J., Shin, C., Kim, S., Kim, H., Yoon, S.: Perception Prioritized Training of Diffusion Models. In: CVPR. pp. 11462–11471 (2022)
11. Eigen, D., Krishnan, D., Fergus, R.: Restoring an image taken through a window covered with dirt or rain. In: ICCV. pp. 633–640 (2013)
12. Fan, Q., Yang, J., Hua, G., Chen, B., Wipf, D.: A Generic Deep Architecture for Single Image Reflection Removal and Image Smoothing. In: ICCV. pp. 3238–3247 (2017)
13. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6), 381–395 (1981)
14. He, T., Zhang, Z., Zhang, H., Zhang, Z., Xie, J., Li, M.: Bag of tricks for image classification with convolutional neural networks. In: CVPR. pp. 558–567 (2019)
15. Hong, Y., Zhong, H., Weng, S., Liang, J., Shi, B.: L-Differ: Single Image Reflection Removal with Language-Based Diffusion Model. In: ECCV. pp. 58–76 (2024)
16. Hu, J., Yang, C., Zhou, Z., Fang, J., Yang, X., Tian, Q., Shen, W.: Dereflexion Any Image with Diffusion Priors and Diversified Data. In: AAAI. pp. 4860–4868 (2026)
17. Hu, Q., Guo, X.: Trash or Treasure? An Interactive Dual-Stream Strategy for Single Image Reflection Separation. In: NeurIPS. vol. 30, pp. 24683–24694 (2021)
18. Hu, Q., Guo, X.: Single Image Reflection Separation via Component Synergy. In: ICCV. pp. 13092–13101 (2023)
19. Hu, Q., Wang, H., Guo, X.: Single image reflection separation via dual-stream interactive transformers. In: NeurIPS. vol. 37, pp. 55228–55248 (2024)
20. Jin, Y., Li, X., Wang, J., Zhang, Y., Zhang, M.: Raindrop Clarity: A Dual-Focused Dataset for Day and Night Raindrop Removal. In: ECCV. pp. 1–17 (2024)
21. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: Musiq: Multi-scale image quality transformer. In: ICCV. pp. 5148–5157 (2021)
22. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014)
23. Li, C., Yang, Y., He, K., Lin, S., Hopcroft, J.E.: Single image reflection removal through cascaded refinement. In: CVPR. pp. 3562–3571 (2020)

24. Li, Y., Monno, Y., Okutomi, M.: Dual-Pixel Raindrop Removal. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(12), 10748–10762 (2024)
25. Lin, X., He, J., Chen, Z., Lyu, Z., Dai, B., Yu, F., Qiao, Y., Ouyang, W., Dong, C.: DiffBIR: Toward Blind Image Restoration with Generative Diffusion Prior. In: *ECCV*. pp. 430–448 (2024)
26. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *ICLR* (2019)
27. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International journal of computer vision* **60**(2), 91–110 (2004)
28. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *ICML*. pp. 8162–8171 (2021)
29. Özdenizci, O., Legenstein, R.: Restoring Vision in Adverse Weather Conditions With Patch-Based Denoising Diffusion Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(8), 10346–10357 (2023)
30. Porav, H., Bruls, T., Newman, P.: I can see clearly now: Image restoration via de-raining. In: *ICRA*. pp. 7087–7093 (2019)
31. Qian, R., Tan, R.T., Yang, W., Su, J., Liu, J.: Attentive Generative Adversarial Network for Raindrop Removal from a Single Image. In: *CVPR*. pp. 2482–2491 (2018)
32. Quan, R., Yu, X., Liang, Y., Yang, Y.: Removing Raindrops and Rain Streaks in One Go. In: *CVPR*. pp. 9143–9152 (2021)
33. Quan, Y., Deng, S., Chen, Y., Ji, H.: Deep learning for seeing through window with raindrops. In: *ICCV*. pp. 2463–2471 (2019)
34. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR*. pp. 10684–10695 (2022)
35. Shao, M.W., Li, L., Meng, D.Y., Zuo, W.M.: Uncertainty Guided Multi-Scale Attention Network for Raindrop Removal from a Single Image. *IEEE Transactions on Image Processing* **30**, 4828–4839 (2021)
36. Soboleva, V., Shipitko, O.: Raindrops on Windshield: Dataset and Lightweight Gradient-Based Detection Algorithm. In: *IEEE Symposium Series on Computational Intelligence*. pp. 1–7 (2021)
37. Su, S., Yan, Q., Zhu, Y., Zhang, C., Ge, X., Sun, J., Zhang, Y.: Blindly Assess Image Quality in the Wild Guided by a Self-Adaptive Hyper Network. In: *CVPR*. pp. 3664–3673 (2020)
38. Sun, S., Ren, W., Gao, X., Wang, R., Cao, X.: Restoring Images in Adverse Weather Conditions via Histogram Transformer. In: *ECCV*. pp. 111–129 (2024)
39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: *NeurIPS*. vol. 30, pp. 5998–6008 (2017)
40. Wan, R., Shi, B., Duan, L.Y., Tan, A.H., Kot, A.C.: Benchmarking Single-Image Reflection Removal Algorithms. In: *ICCV*. pp. 3922–3930 (2017)
41. Wang, J., Chan, K.C.K., Loy, C.C.: Exploring CLIP for assessing the look and feel of images. In: *AAAI*. pp. 2555–2563 (2023)
42. Wang, J., Yue, Z., Zhou, S., Chan, K.C., Loy, C.C.: Exploiting Diffusion Prior for Real-World Image Super-Resolution. *International Journal of Computer Vision* **132**(12), 5929–5949 (2024)
43. Xiao, J., Fu, X., Liu, A., Wu, F., Zha, Z.J.: Image De-Raining Transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(11), 12978–12995 (2023)
44. You, S., Tan, R.T., Kawakami, R., Ikeuchi, K.: Adherent Raindrop Detection and Removal in Video. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(9), 1721–1733 (2016)

45. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: ICCV. pp. 3836–3847 (2023)
46. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)
47. Zhao, H., Li, M., Hu, Q., Guo, X.: Reversible decoupling network for single image reflection removal. In: CVPR. pp. 26430–26439 (2025)
48. Zheng, D., Wu, X.M., Yang, S., Zhang, J., Hu, J.F., Zheng, W.s.: Selective hour-glass mapping for universal image restoration based on diffusion model. In: CVPR. pp. 25445–25455 (2024)
49. Zheng, Q., Chen, J., Lu, Z., Shi, B., Jiang, X., Yap, K.H., Duan, L.Y., Kot, A.C.: What does plate glass reveal about camera calibration? In: CVPR. pp. 3022–3032 (2020)
50. Zhou, S., Chan, K.C., Li, C., Loy, C.C.: Towards Robust Blind Face Restoration with Codebook Lookup Transformer. In: NeurIPS. vol. 35, pp. 30599–30611 (2022)
51. Zhu, C., Deng, S., Song, X., Li, Y., Wang, Q.: Mamba Collaborative Implicit Neural Representation for Hyperspectral and Multispectral Remote Sensing Image Fusion. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–15 (2025)
52. Zhu, C., Song, X., Li, Y., Deng, S., Zhang, T.: A spatial-frequency dual-domain implicit guidance method for hyperspectral and multispectral remote sensing image fusion based on Kolmogorov–Arnold Network. *Information Fusion* **123**, 103261 (2025)
53. Zhu, Y., Fu, X., Jiang, P.T., Zhang, H., Sun, Q., Chen, J., Zha, Z.J., Li, B.: Revisiting Single Image Reflection Removal In the Wild. In: CVPR. pp. 25468–25478 (2024)