

Towards Robust Rain Removal Against Adversarial Attacks: A Comprehensive Benchmark Analysis and Beyond

Yi Yu^{1,2} Wenhan Yang^{1*} Yap-Peng Tan¹ Alex C. Kot¹

¹School of Electrical and Electronic Engineering, Nanyang Technological University

²ROSE Lab, Interdisciplinary Graduate Programme, Nanyang Technological University

yuyi0010@e.ntu.edu.sg

{wenhan.yang, eyptan, eackot}@ntu.edu.sg

Abstract

Rain removal aims to remove rain streaks from images/videos and reduce the disruptive effects caused by rain. It not only enhances image/video visibility but also allows many computer vision algorithms to function properly. This paper makes the first attempt to conduct a comprehensive study on the robustness of deep learning-based rain removal methods against adversarial attacks. Our study shows that, when the image/video is highly degraded, rain removal methods are more vulnerable to the adversarial attacks as small distortions/perturbations become less noticeable or detectable. In this paper, we first present a comprehensive empirical evaluation of various methods at different levels of attacks and with various losses/targets to generate the perturbations from the perspective of human perception and machine analysis tasks. A systematic evaluation of key modules in existing methods is performed in terms of their robustness against adversarial attacks. From the insights of our analysis, we construct a more robust deraining method by integrating these effective modules. Finally, we examine various types of adversarial attacks that are specific to deraining problems and their effects on both human and machine vision tasks, including 1) rain region attacks, adding perturbations only in the rain regions to make the perturbations in the attacked rain images less visible; 2) object-sensitive attacks, adding perturbations only in regions near the given objects. Code is available at https://github.com/yuyi-sd/Robust_Rain_Removal.

1. Introduction

Rain removal methods aim to remove the disruptive effects caused by rain streaks to restore a clean version of the image. It not only largely improves the visibility of the rainy image but can also improve the performance of many

*Corresponding author.




Figure 1. Left is from corresponding patch of clean input/output, and right is from adversarial input/output with perturbation bound $\epsilon = 4/255$. The semantic segmentation overlap of output is shown in the last column. The testing image is from RainCityscape [13].

subsequent (downstream) computer vision applications.

Early approaches are mainly model-driven and address the deraining problem based on the statistical properties of rain streaks and background scenes, e.g. image decomposition [17], sparse coding [27], and Gaussian mixture model [26]. These methods can well handle light rain. However, they fall short in cases of handle heavy rain and often blur the background scenes. Recently, deep-learning based deraining methods [9, 41] have become the mainstream. These methods have the capacity to model more complicated mappings from rain images to clean images and offer a better performance in terms of less remaining rain streaks and better preserved background scene. They also enhanced the performance of practical applications such as video surveillance.

While many deep learning-based deraining methods have been introduced, it isn't a thorough study on the robustness of these methods against adversarial attacks. It is also a real concern of deraining methods that some unnoticeable perturbations can alter the predicted results of deep networks [12, 31]. As the rainy images can be usually highly degraded by intensive rain streaks, perturbations can be easily and noticeably hidden in such images. These adversarially generated outputs can also compromise the reliability and stability of the subsequent applications such as video surveillance and autonomous driving, that adopt rain removal methods as a pre-processing module.

In this paper, we make the first attempt to investigate, improve, and evaluate the robustness of deep learning-based rain removal methods against adversarial attacks. Our main contributions are summarized below.

- A thorough analysis of existing deraining methods at different levels of robustness against adversarial attacks and with various losses/targets to generate the perturbations, from both the perspectives of human and machine vision tasks.
- A systematic evaluation of the modules in existing methods at both the model side (recurrence, attention, receptive field, and side information) and loss side (adversarial loss) in terms of their impact on robustness.
- A more robust deraining method constructed by integrating the effective modules which are identified from our analysis.
- Investigation of various adversarial attack types and scenarios specific to rain removal degradation in terms of both human and machine vision tasks.

2. Related Work

2.1. Adversarial Attacks

Deep neural networks [32] are vulnerable to adversarial examples that are visually indistinguishable from natural images but can compromise the performance of models. From the attack method aspect, Szegedy *et al.* [32] developed an optimization-based approach to find adversarial examples within specific amounts of input perturbation. Goodfellow *et al.* [12] proposed the Fast Gradient Sign Method (FGSM) and Kurakin *et al.* [19] developed multiple variations. Madry *et al.* [28] proposed a more powerful multi-step variant, which is essentially Projected Gradient Descent (PGD). There are also studies on the effects of adversarial attacks on various computer vision applications and scenarios, *e.g.* human pose estimation [5], object detection [39], and semantic segmentation [39].

There are also works on adversarial attacks and robustness analysis in low-level processing tasks. In [29], super-resolution is adopted as the tool to mitigate the adversarial attacks for deep image classification models. In [47], adversarial attacks are exploited on super-resolution model to attack the down-stream computer vision tasks, *e.g.* image classification, style transfer, and image captioning. In [3], the robustness of deep learning-based super-resolution methods against adversarial attacks is examined. In [11, 49], rain streak/haze is considered as one type of adversarial noise. In [4], the robustness of deep image-to-image models against adversarial attacks is evaluated from the perspective of human vision. However, there is no work investigating adversarial attacks and robustness analysis in

rain removal scenarios. It is also absent to consider the adversarial attacks from the perspective of human and machine vision tasks.

2.2. Single-image rain removal

It is non-trivial to separate rain streaks and background scenes from their mixed versions. Methods based on signal separation and texture classification have been proposed to address the challenge. Early non-deep learning-based works [2, 17, 18, 26, 27] are constructed based on statistical models. In [17], an attempt is made to perform the single-image deraining based on image decomposition using morphological component analysis. In a successive work [27], a mutual exclusivity property is incorporated into a sparsity model to form a discriminative sparse coding that can better separate the rain/background layers from their non-linear composite.

Recently, deep-learning based deraining methods [9, 22, 41, 50] have been proposed to obtain superior rain removal results. More recent works [10, 20, 23, 34, 48] also focus on advanced deep networks. Others [7, 8, 16, 21, 33, 40, 43–45] make use of effective priors for deraining. Some studies focus on the generalization capability of deraining models. In [37], Wei *et al.* make use of real rainy images in network training by formulating the residual between an input rainy image and its expected deraining result as a mixture of Gaussians. In [46], a Gaussian process-based semi-supervised learning is proposed to allow the network learn with synthetic data and generalize deraining performance with unlabeled real rainy images. These methods can provide better deraining results in terms of less remaining rain streaks and rich background textures. However, there is little work on the vulnerability of these deep networks against adversarial attacks. This work aims to fill this gap.

3. Benchmarking Adversarial Robustness of Deraining Models

3.1. Attack Framework

Adversarial attacks aim to deteriorate the output of the deraining methods by adding a small amount of visually unperceivable perturbations to the input rainy images. To generate the adversarial perturbations, we develop an attack method based on one optimization-based method PGD [28], which has been extensively used and considered as a powerful attack method so far for evaluating the robustness of classification models [6].

We consider a deraining model $f(\cdot|\theta)$ parametrized by θ . We denote X as the input image, Y as the groundtruth image, δ as the perturbations, and D as the metric to measure the degradation. The objective of adversarial attacks is to maximize the deviation of the output from the target, *e.g.*

the original output:

$$\delta = \arg \max_{\delta, \|\delta\|_p \leq \epsilon} D(f(X|\theta), f(X + \delta|\theta)). \quad (1)$$

To solve the maximization problem with ℓ_p -norm bound constraint (usually ℓ_∞ -norm), we use the PGD approach to obtain the perturbations iteratively:

$$\begin{aligned} \omega^{t+1} &= \delta^t + \alpha \text{sgn}(\nabla_{\delta} D(f(X|\theta), f(X + \delta|\theta))), \quad (2) \\ \delta^{t+1} &= \text{clip}_{[-\epsilon, \epsilon] \cap [-X, 1-X]}(\omega^{t+1}), \quad (3) \end{aligned}$$

where ∇ represents the gradient operation, sgn extracts the sign of gradients, and the clip operation guarantees that the perturbations are within $[-\epsilon, \epsilon]$ and the perturbed input is within $[0, 1]$. The term α controls the step length each iteration, and ϵ represents the maximum perturbation allowed for each pixel value. The initial δ^0 is sampled from the uniform distribution $U(-\epsilon, \epsilon)$, and the final adversarial perturbations δ^T is obtained after T iterations.

Based on different attack objectives, we define two types of corresponding metrics for D :

- *Restoration (Human Vision)*: per-pixel measures, e.g. ℓ_2 Euclidean distance:

$$\delta = \arg \max_{\delta, \|\delta\|_\infty \leq \epsilon} \|f(X + \delta|\theta) - f(X|\theta)\|_2. \quad (4)$$

- *Downstream CV Tasks (Machine Vision)*: the distance of feature extracted from pretrained models, e.g. Learned Perceptual Image Patch Similarity (LPIPS) [52]:

$$\delta = \arg \max_{\delta, \|\delta\|_\infty \leq \epsilon} \ell_{\text{lips}}(f(X + \delta|\theta), f(X|\theta)). \quad (5)$$

We denote the first one as LMSE attack and the second as LPIPS attack. For each attack type and perturbation bound, $\delta_{D, \epsilon}(X)$ is obtained as Eqs. (2-3), and the general form of adversarial robustness for deraining model f is given by

$$R_f^{D, \epsilon} = \mathbb{E}_{X \sim \mathbb{P}_{\text{data}}} [P(Y, f(X + \delta_{D, \epsilon}(X)))], \quad (6)$$

where the samples follow the distribution \mathbb{P}_{data} , and P is the evaluation metric. While it is impossible to evaluate the robustness by considering all values of ϵ , we estimate it using finite ϵ and limited test samples. With the estimation, we use mean-Adversarial-Performance (mAP) to evaluate the adversarial robustness against each type of attack D :

$$mAP_D = \frac{1}{n(E)} \sum_{\epsilon \in E} \frac{1}{n(\mathcal{D}_t)} \sum_{X \in \mathcal{D}_t} P(Y, f(X + \delta_{D, \epsilon}(X))), \quad (7)$$

where E is the set of ϵ to be evaluated, \mathcal{D}_t is the test dataset, and $n(\cdot)$ counts the numbers. Note that P is usually measured in terms of PSNR and SSIM.

Table 1. PSNR, parameter and module comparisons on two datasets. AT denotes attention module, RB recurrent blocks, SI side information and DD diverse dilations, and RC RainCityscape.

Model	Para.	AT	RB	SI	DD	Rain100H	RC
JORDER-E [41]	4.17M		✓	✓	✓	29.75	32.51
RCDNet [33]	3.17M		✓			29.65	31.44
MPRNet [48]	3.64M	✓				30.56	36.31
PReNet [30]	1.50M		✓			29.58	33.09
UMRL [45]	984K					26.05	30.15
RESCAN [25]	1.04M	✓	✓		✓	28.90	34.90

3.2. Evaluation Metrics

In our benchmark, we consider two types of performance evaluation metrics:

- The commonly used quality measures including peak signal to noise ratio (PSNR) and structural index similarity (SSIM) [36], for evaluating signal fidelity.
- Task-driven metrics that evaluate the performance of down-stream machine analysis tasks. Specifically, we evaluate the performance of object detection and semantic segmentation, *i.e.* mean Intersection over Union (mIoU) for semantic segmentation and Average Precision (AP) for pedestrian detection, with the approaches SSeg [53] and Pedestron [14] on the RainCityscape dataset [13].

3.3. Implementation Details

1) Datasets. We employ two synthetic deraining datasets (Rain100H [42] and RainCityscape [13]) and some real images [37]. Rain100H dataset is one of the most commonly used datasets consisting of 1800 paired rain/non-rain images of size 480×320 for training and 100 paired images of the same size for testing. PSNR and SSIM of input testing set are 12.05 and 0.3623. RainCityscape dataset consists of 2875 paired images of size 512×256 for training and 100 paired images for testing. PSNR and SSIM of input testing set are 19.71 and 0.7087. Besides paired images, RainCityscape dataset also provides additional labels for object detection and semantic segmentation, which allow us to evaluate the performance of down-stream tasks. In our experiment, the model evaluated on the testing set of Rain100H [42] and RainCityscape [13] is trained on their respective training sets.

2) Deraining Methods. We consider six state-of-the-art deep learning-based deraining models as shown in Table 1.

3) Perturbation Levels. The adversarial attacks are implemented based on PGD [28]. We also set $\epsilon \in E = \{1/255, 2/255, 4/255, 8/255\}$ regarding to ℓ_∞ -norm, $\alpha = \epsilon/4$, and $T = 20$.

4) Results and Remarks. Figure 2 and Table 2 show the robustness results of six deraining models against adversarial attacks. For the RainCityscape dataset, we also consider




Figure 2. Adversarial Robustness of six deraining models and our proposed method evaluated by PSNR, SSIM, mIoU of Semantic Segmentation and AP of Pedestrian Detection on the Rain100H [42] and RainCityscape [13] datasets. Subfigures in the first row denote the performance against LMSE attack, and subfigures in the second row denote the LPIPS attack. RC denotes RainCityscape. Note that $\epsilon = 0/255$ denotes the performance of output to input without perturbations.

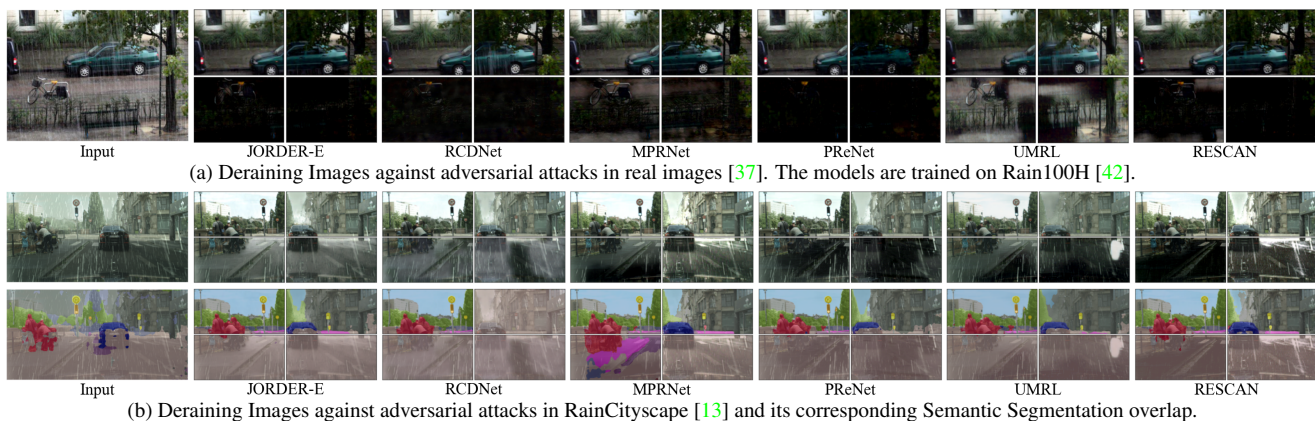


Figure 3. Visual comparison of the deraining outputs for input with perturbations based on **LMSE** attack. In each image other than input, patches on the top-left, top-right, bottom-left and bottom-right are corresponding patches from the output to input with perturbation bound $\epsilon = \{0/255, 1/255, 2/255, 4/255\}$. Best view by zooming in. More visualization results can be found in the supplementary material.

the performance of pedestrian detection and semantic segmentation. From the benchmark results, we can have the following:

- Even very small perturbations ($\epsilon = 1$) can heavily degrade the deraining performance of existing methods. Larger perturbations ϵ lead to significantly performance drops in image quality.
- On the Rain100H dataset, MPRNet [48] is the most robust model in terms of PSNR and SSIM, which might be attributed to its use of attention modules to suppress the effects of adversarial noise at the feature level.
- On the RainCityscape dataset, the gaps of PSNR and SSIM values against adversarial attacks with identical ϵ are closer for all the models, while MPRNet is most robust for the down-stream tasks in terms of mIoU and AP.
- When attacking with LPIPS, it is observed that perturbations from 0 to 1 have the largest performance drop in terms of PSNR and SSIM while larger perturbations lead to smaller performance gaps. Especially on RainCityscape, perturbations from 1 to 8 bring about the same degradation in terms of PSNR, while SSIM, which considers more structural information, contin-

Table 2. Adversarial Robustness (mean-Adversarial-Performance) of six deraining models on two datasets.

Datasets	Rain100H		RainCityscape			
	PSNR	SSIM	PSNR	SSIM	mIoU	AP
Restoration (Human Vision): MSE Loss						
JORDER-E	10.91	0.3383	12.89	0.5920	0.2799	0.6511
RCDNet	12.29	0.3744	13.21	0.5386	0.1924	0.6307
MPRNet	13.47	0.5253	12.76	0.5693	0.3945	0.7217
PReNet	8.82	0.2107	13.12	0.4864	0.1151	0.6056
UMRL	12.94	0.4297	12.02	0.5512	0.1886	0.6086
RESCAN	11.40	0.3528	11.13	0.5088	0.2168	0.6065
Downstream CV Tasks (Machine Vision): LPIPS Loss						
JORDER-E	16.22	0.5423	20.18	0.6692	0.2354	0.6014
RCDNet	18.54	0.6086	20.39	0.6498	0.1001	0.5427
MPRNet	22.70	0.6878	20.10	0.6815	0.2451	0.6753
PReNet	11.22	0.3628	20.48	0.6598	0.2291	0.6398
UMRL	21.44	0.6594	19.03	0.6420	0.1140	0.4548
RESCAN	17.54	0.5640	20.44	0.6681	0.1652	0.6376

ues to drop. And it’s not surprising to find that LMSE attack brings more drop on PSNR and SSIM, while LPIPS attack brings more on mIoU and AP.

- It is interesting to note that the deraining results are vulnerable in terms of semantic segmentation performance. The attack with the perturbation ($\epsilon = 2$) has already crashed the segmentation performance.
- Existing methods are more robust for object detection than semantic segmentation.

4. Towards Robust Deraining Model

In this section, we perform a comprehensive analysis on the modules and losses to evaluate their adversarial robustness.

4.1. Module Analysis

The state-of-the-art deraining methods usually consist of recurrent blocks, several attention modules, side information injection and larger receptive field. We perform an ablation study of these modules based on the baseline of RCDNet [33] and JOEDER-E [42].

1) Recurrent Blocks. As there might be different rain streak layers overlapping with each other and it is not easy to remove all streaks in one stage, many deraining models incorporate a recurrent structure to perform rain removal into multiple stages. The process can be formulated as:

$$\begin{aligned}
 O_0 &= O, \\
 [R_t, B_t] &= f^t(O_{t-1}, R_{t-1}), 1 \leq t \leq T, \\
 O_t &= O - R_t = B_t, 1 \leq t \leq T,
 \end{aligned} \tag{8}$$

where O is the input, T is the number of recurrent blocks, R_t is the output rain streaks at the t -th stage, B_t is the output background at the t -th stage, and f^t is the t -th CNN block. In each stage, R_t and B_t are updated, and the final output is B_T .

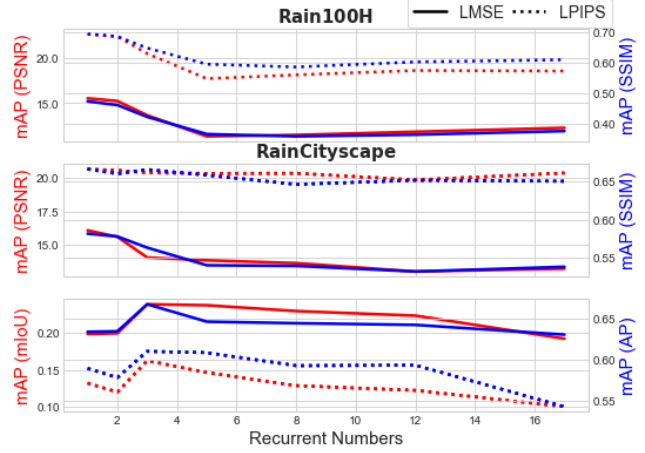


Figure 4. Recurrent Blocks: Adversarial Robustness for RCDNet with various stage numbers. Red curves correspond to the left y-axis, and blue curves correspond to the right y-axis.

Table 3. Attention Module: Adversarial Robustness of RCDNet with various attention module. Clean represents the performance for input without perturbations in terms of SSIM.

Datasets	Rain100H		RainCityscape				
	Clean	PSNR SSIM	Clean	PSNR	SSIM	mIoU	AP
Restoration (Human Vision): MSE Loss							
RCDNet	0.8867	12.29 0.3744	0.9517	13.21	0.5386	0.1924	0.6307
+ SE (MUL)	0.8834	16.10 0.5325	0.9626	13.43	0.5768	0.2874	0.6425
+ SE (ADD)	0.8859	15.63 0.5246	0.9697	13.83	0.6100	0.3502	0.6672
+ CBAM	0.8777	14.47 0.4930	0.9553	13.55	0.6039	0.2699	0.6465
+ GC (MUL)	0.8797	15.83 0.5371	0.9617	13.75	0.5959	0.2869	0.6666
+ GC (ADD)	0.8694	15.81 0.5595	0.9644	10.99	0.5648	0.2947	0.6347
+ SK	0.8813	15.29 0.5260	0.9553	12.77	0.5574	0.3478	0.6341
Downstream CV Tasks (Machine Vision): LPIPS Loss							
RCDNet	0.8867	18.54 0.6086	0.9517	20.39	0.6498	0.1001	0.5427
+ SE (MUL)	0.8834	24.00 0.7108	0.9626	20.23	0.6405	0.1976	0.6838
+ SE (ADD)	0.8859	23.68 0.7011	0.9697	20.44	0.6867	0.2461	0.6730
+ CBAM	0.8777	22.33 0.6552	0.9553	19.25	0.6783	0.1373	0.6603
+ GC (MUL)	0.8797	24.50 0.7176	0.9617	20.02	0.6574	0.2242	0.6725
+ GC (ADD)	0.8694	23.81 0.7159	0.9644	19.43	0.6114	0.1710	0.6516
+ SK	0.8813	23.58 0.7048	0.9553	20.48	0.7109	0.2389	0.5603

We examine the adversarial robustness of RCDNet with various stage numbers $\{1, 2, 3, 5, 8, 12, 17\}$. For a fair comparison, the models with fewer stages have wider feature maps to keep the number of parameters around $3M$.

Insight: As shown in Fig. 4, the models with one or two recurrent blocks are more robust in terms of PSNR and SSIM, while the model with three stages is the most robust considering the mIoU and AP for down-stream tasks and the second most robust in terms of PSNR and SSIM.

2) Attention Module. RESCAN [25] and MPRNet [48] both consist of attention modules to enhance its deraining performance. We look into their effects in terms of robustness. We consider several attention modules based on feature re-calibration in CNN:

- *SENet* [15]: Features are first passed through an avg-pooling *squeeze* operator across spatial dimension and

Table 4. Side Information: Effect of side information on the adversarial robustness of JORDER-E. Clean represents the performance for input without perturbations in terms of SSIM.

Datasets	Rain100H			RainCityscape				
	Clean	PSNR	SSIM	Clean	PSNR	SSIM	mIoU	AP
Restoration (Human Vision): MSE Loss								
JORDER-E	0.8921	10.91	0.3383	0.9600	12.89	0.5920	0.2799	0.6511
NO Mask	0.8809	15.37	0.4991	0.9626	14.89	0.6143	0.2955	0.6202
Downstream CV Tasks (Machine Vision): LPIPS Loss								
JORDER-E	0.8921	16.22	0.5423	0.9600	20.18	0.6692	0.2354	0.6014
NO Mask	0.8809	22.71	0.6947	0.9626	20.01	0.6707	0.2574	0.6176

then produce the channel-wise activations by the subsequent *excitation* operation. Note that the activations can be added or multiplied to the feature maps.

- *GCNet* [1]: GCNet differs from SENet in the *squeeze* operation, and adopts the structures of Non-Local Networks [35] for spatial modeling.
- *CBAM* [38]: CBAM has two sequential sub-modules: channel and spatial. The subsequent module concatenates the output of Max and Avg pooling (across channel), and uses a Conv layer to produce the spatial attention map.
- *SKNet* [24]: "Selective Kernel" (SK) block fuses the features in multiple branches with different kernel sizes via softmax attention. Different channel attentions on these branches are performed as well.

Insight: By inserting the attention module into the same location of RCDNet (after each conv block), we find that *all attention modules indeed improve the robustness and SE with the addition is the most robust considering the downstream tasks; it also performs well in terms of PSNR and SSIM, as can be seen from Table 3.*

3) Side Information. A few deraining models require side information during training, *e.g.*, JORDER-E [42] adopts rain mask and density as supervision. We look into the effect of this side information on adversarial robustness.

Insight: From Table 4, JORDER-E without rain mask performs better against adversarial attacks, suggesting that *the model without side information tends to be more robust.*

4) Receptive Field. It is well known that, increasing the receptive field can boost the model's capacity in handling more complex signal mapping. However, it is unclear how different receptive fields affect the robustness of a model. Some deraining models adopt multi-path dilated convolution structure to increase its receptive field, which provides a means to analyze the robustness of a model while keeping others unchanged. JORDER-E [42]'s basic block consists of three paralleled convolution paths with a dilation factor of 1, 2, and 3, respectively. By removing paths with larger receptive fields, the effect of receptive fields on adversarial robustness is discussed.

Table 5. Receptive field: Effect of receptive field on the adversarial robustness of JORDER-E. Clean represents the performance for input without perturbations in terms of SSIM.

Datasets	Rain100H			RainCityscape				
	Clean	PSNR	SSIM	Clean	PSNR	SSIM	mIoU	AP
Restoration (Human Vision): MSE Loss								
D = {1, 2, 3}	0.8921	10.91	0.3383	0.9600	12.89	0.5920	0.2799	0.6511
D = {1, 2}	0.8969	9.32	0.2716	0.9622	12.60	0.5060	0.2192	0.6154
D = {1}	0.8956	9.19	0.2650	0.9608	12.56	0.5193	0.2268	0.6215
Downstream CV Tasks (Machine Vision): LPIPS Loss								
D = {1, 2, 3}	0.8921	16.22	0.5423	0.9600	20.18	0.6692	0.2354	0.6014
D = {1, 2}	0.8969	12.65	0.4507	0.9622	19.61	0.6522	0.1405	0.5573
D = {1}	0.8956	12.50	0.4426	0.9608	20.04	0.6607	0.1500	0.5576

Table 6. Effect of adversarial loss on the adversarial robustness of RCDNet. Clean represents the performance for input without perturbations in terms of SSIM.

Datasets	Rain100H			RainCityscape				
	Clean	PSNR	SSIM	Clean	PSNR	SSIM	mIoU	AP
Restoration (Human Vision): MSE Loss								
RCDNet	0.8867	12.29	0.3744	0.9517	13.21	0.5386	0.1924	0.6307
+ Adv Loss	0.8584	21.85	0.7619	0.9308	17.78	0.6503	0.2855	0.6944
Downstream CV Tasks (Machine Vision): LPIPS Loss								
RCDNet	0.8867	18.54	0.6086	0.9517	20.39	0.6498	0.1001	0.5427
+ Adv Loss	0.8584	25.93	0.8050	0.9308	22.68	0.7230	0.2365	0.6289

Insight: Superior adversarial performance can be observed from Table 5 when a *larger receptive field* is adopted.

4.2. Fidelity Loss vs. Adversarial Loss

Zhang *et al.* [51] explored the trade-off between accuracy and robustness for classification problems and proposed a new formulation of adversarial loss to defense the adversarial attacks:

$$L_{\text{adv}} = \lambda \max_{X' \in \mathbb{B}(X, \epsilon)} \mathcal{L}(f(X|\theta), f(X'|\theta)), \quad (9)$$

where $f(X|\theta)$ is the output vector of learning model, $\mathcal{L}(\cdot, \cdot)$ is a multi-class calibrated loss such as cross-entropy loss, and $\mathbb{B}(X, \epsilon)$ represents the neighborhood of X : $\{X' : \|X' - X\|_{\infty} \leq \epsilon\}$. The trade-off can be balanced by adjusting the loss weight λ in the final combined loss and the perturbation bound ϵ .

For rain removal model, the adversarial loss can be obtained as:

$$L_{\text{adv}} = \lambda \max_{\delta, \|\delta\|_{\infty} \leq \epsilon} \|f(X + \delta|\theta) - f(X|\theta)\|_2. \quad (10)$$

Insight: Table 6 shows that incorporating *adversarial loss can improve the adversarial robustness, but it also compromises the performance on input without perturbations.*

4.3. Ensemble Modules into a Robust Deraining Model

Based on the above insights, we build a robust deraining model by integrating MPRNet, SE with addition, convolution layer with diverse dilations, three recurrent blocks, and

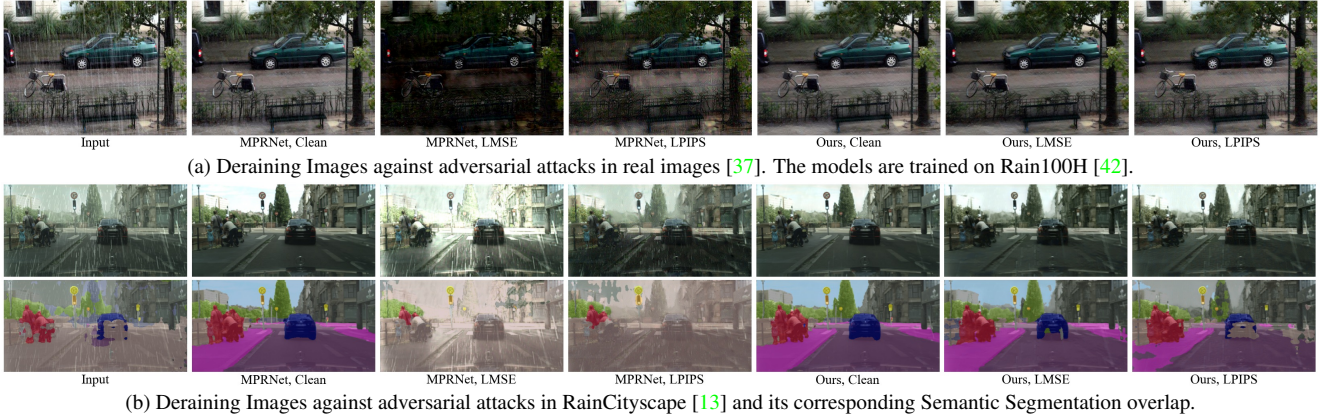


Figure 5. Visual comparison of the derain outputs of RainCityscape [13] test data with adversarial perturbation bound $\epsilon = 4/255$. Images annotated with Clean are the output of input without perturbations, and images annotated with LMSE/LPIPS are the output of input with adversarial perturbations. Best view by zooming in. More visualization results can be found in the supplementary material.

Table 7. Ablation Study: Robustness of enhanced model compared with MPRNet. Clean represents the performance for input without perturbations in terms of SSIM. AL denotes adversarial loss, AT denotes attention module, RB denotes recurrent blocks and DD denotes dilation dilations.

Datasets	Rain100H			RainCityscape			
	Clean	PSNR	SSIM	Clean	PSNR	SSIM	mIoU AP
Restoration (Human Vision): MSE Loss							
MPRNet	0.8974	13.47	0.5253	0.9811	12.76	0.5693	0.3945 0.7217
Ours	0.8670	24.46	0.8099	0.9401	22.82	0.8264	0.7592 0.7487
Ours w/o AL	0.8841	15.35	0.5997	0.9794	12.67	0.5707	0.4745 0.7328
Ours w/o RB	0.8546	22.09	0.7710	0.9389	21.24	0.7992	0.6512 0.7242
Ours w/o DD	0.8685	22.08	0.7782	0.9466	22.05	0.8023	0.7011 0.7230
Ours w/o AT	0.8474	21.41	0.7388	0.9411	18.75	0.7184	0.3062 0.6552
Downstream CV Tasks (Machine Vision): LPIPS Loss							
MPRNet	0.8974	22.70	0.6878	0.9811	20.10	0.6815	0.2451 0.6753
Ours	0.8670	26.81	0.8138	0.9401	25.75	0.8484	0.7212 0.7033
Ours w/o AL	0.8841	24.49	0.7110	0.9794	21.00	0.6811	0.3107 0.6955
Ours w/o RB	0.8546	26.58	0.8011	0.9389	24.84	0.8355	0.5904 0.6814
Ours w/o DD	0.8685	26.55	0.7993	0.9466	25.06	0.8342	0.6352 0.6844
Ours w/o AT	0.8474	25.75	0.7689	0.9411	22.15	0.7530	0.2194 0.5917

adversarial loss. Figure 2 and Table 7 shows that our model is much more robust than MPRNet. Figure 5 presents the visual results of different methods. While MPRNet is already the most robust among all state-of-the-art methods, its restoration results in Figure 5 (a) and (b) as well as the down-stream segmentation results in Figure 5 (c) are still heavily degraded. Benefiting from the robust modules and loss derived from our module analysis, our enhanced model is more robust and obtains visually similar results compared to the ones without adversarial perturbations.

We also conduct an ablation study by subtracting one of the robust modules from our model. We can make the following conclusions from Table 7:

- All aforementioned modules contribute to the robustness in our model.
- Adversarial loss is the key component to improve robustness at the cost of reducing the performance for

input without perturbations a bit.

- The attention module can help maintain robustness when the results are evaluated by down-stream tasks.

5. Discussion: Advanced Attack Scenarios

The two basic attacks described in Sec. 3.1 aim to maximize the deviation of the output with the perturbations embedded over the entire input image. In this section, we investigate other advanced attack scenarios and show that our model can handle them well, especially for the case where human and machine vision applications are considered jointly.

5.1. Object-sensitive Attack

In this scenario, the attack aims to compromise the downstream applications on a specific object, such as pedestrians and cars. With such perturbations added to the input, the adversarial attacks minimize the perceptual loss between the source patch S and target patch T . Thus, the object at source patch S is more similar to the target patch T from the machine vision perspective. The attack can be formulated as follows:

$$\hat{\delta} = \arg \max_{\delta, \|\delta\|_{\infty} \leq \epsilon} - \ell_{pips}(\text{Crop}_S f(X + \delta|\theta), \text{Crop}_T f(X|\theta)),$$

$$\delta = \hat{\delta} \circ M_S, \quad (11)$$

where $\text{Crop}_S(\cdot)$ denotes the image-crop operation at the patch S , M_S is the binary mask specifying the source patch S , and \circ represents the element-wise multiplication. We perform the object-sensitive attack for both MPRNet and our proposed method, and the results are shown in Figure 6 (a). In the output of MPRNet, one person riding a bicycles in the source patch (red box) was mis-classified as the ground in the target patch (blue box) by the semantic segmentation approach SSeg [53], while the output of our model survived the attack.



Figure 6. Derain Images of **MPRNet** and **Ours** against several attacking scenarios and its corresponding semantic segmentation overlap in RainCityscape [13]. Note that the adversarial perturbation bound $\epsilon = 4/255$. In (a), the patches with red boxes are the source objects to attack, and patches with blue boxes are the target objects. More visualization results can be found in the supplementary material.

5.2. Partial Attack (Rain Region Attack)

We also consider the case that perturbations are added to part of the input image. The attack can be formulated as follows:

$$\hat{\delta} = \arg \max_{\delta, \|\delta\|_{\infty} \leq \epsilon} D(f(X|\theta), f(X + \delta \circ M|\theta)), \quad (12)$$

$$\delta = \hat{\delta} \circ M, \quad (13)$$

where M is the binary mask specifying the attack region. Figure 6 (b) shows the results of MPRNet and our method attacked by adding the perturbations to the rain region. The results show that our method is more robust than MPRNet under such an attack.

5.3. Unnoticeable Attack for Down-stream Tasks

We also examine the attack of which the output degradation is unnoticeable but can much affect the performance of downstream tasks, *i.e.* benefiting human vision but degenerating machine vision. Our objective is to maximize the perceptual loss deviation and minimize the MSE loss deviation in the output, which can be described as follows:

$$\delta = \arg \max_{\delta, \|\delta\|_{\infty} \leq \epsilon} \ell_{pips}(f(X + \delta|\theta), f(X|\theta)) - \lambda \|f(X + \delta|\theta) - f(X|\theta)\|_2. \quad (14)$$

where λ balances the trade-off between human vision and machine vision deterioration. From the example in Figure 6 (c), both results of MPRNet and the proposed method

are visually pleasing, but our model is much more robust against this attack in the downstream segmentation task.

5.4. Input-close Attack

The goal of input-close attack is to minimize the MSE loss between the input and output. That is to say, we hope the attacked model totally fails to remove the rain streaks of the input:

$$\delta = \arg \max_{\delta, \|\delta\|_{\infty} \leq \epsilon} - \|f(X + \delta|\theta) - X\|_2. \quad (15)$$

Figure 6 (d) shows that our model can handle this attack.

6. Conclusion

In this paper, we conducted the first comprehensive study on the robustness of deep learning-based rain streak removal methods against adversarial attacks. Based on the analysis, we proposed a more robust deraining method by integrating effective modules. We also evaluated several types of adversarial attacks that are specific to the properties of rain/deraining problems and exploited the properties or requirements of both human and machine vision tasks.

Acknowledgement. This work was done at Rapid-Rich Object Search (ROSE) Lab, Nanyang Technological University. This research is supported in part by the NTU-PKU Joint Research Institute (a collaboration between the Nanyang Technological University and Peking University that is sponsored by a donation from the Ng Teng Fong Charitable Foundation).

References

- [1] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Global context networks. *CoRR*, abs/2012.13375, 2020. [6](#)
- [2] Yi-Lei Chen and Chiou-Ting Hsu. A generalized low-rank appearance model for spatio-temporally correlated rain streaks. In *Proc. IEEE Int'l Conf. Computer Vision*, pages 1968–1975, 2013. [2](#)
- [3] J. Choi, H. Zhang, J. Kim, C. Hsieh, and J. Lee. Evaluating robustness of deep image super-resolution against adversarial attacks. In *Proc. IEEE Int'l Conf. Computer Vision*, pages 303–311, 2019. [2](#)
- [4] Jun-Ho Choi, Huan Zhang, Jun-Hyuk Kim, Cho-Jui Hsieh, and Jong-Seok Lee. Deep Image Destruction: A Comprehensive Study on Vulnerability of Deep Image-to-Image Models against Adversarial Attacks. *arXiv e-prints*, page arXiv:2104.15022, April 2021. [2](#)
- [5] Moustapha M Cisse, Yossi Adi, Natalia Neverova, and Joseph Keshet. Houdini: Fooling deep structured visual and speech recognition models with adversarial examples. In *Proc. Annual Conf. Neural Information Processing Systems*, volume 30, 2017. [2](#)
- [6] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *Proc. Int'l Conf. Machine Learning*, volume 119, pages 2206–2216, 2020. [2](#)
- [7] S. Deng, M. Wei, J. Wang, Y. Feng, L. Liang, H. Xie, F. L. Wang, and M. Wang. Detail-recovery image deraining via context aggregation networks. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 14548–14557, 2020. [2](#)
- [8] Y. Du, J. Xu, X. Zhen, M. Cheng, and L. Shao. Conditional variational image deraining. *IEEE Trans. on Image Processing*, 29:6288–6301, 2020. [2](#)
- [9] X. Fu, J. Huang, D. Zeng, Y. Huang, X. Ding, and J. Paisley. Removing rain from single images via a deep detail network. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 1715–1723, 2017. [1](#), [2](#)
- [10] X. Fu, B. Liang, Y. Huang, X. Ding, and J. Paisley. Lightweight pyramid networks for image deraining. *IEEE Trans. on Neural Networks and Learning Systems*, 31(6):1794–1807, 2020. [2](#)
- [11] Ruijun Gao, Qing Guo, Felix Juefei-Xu, Hongkai Yu, and Wei Feng. AdvHaze: Adversarial Haze Attack. *arXiv e-prints*, page arXiv:2104.13673, April 2021. [2](#)
- [12] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proc. Int'l Conf. Learning Representations*, 2015. [1](#), [2](#)
- [13] Shirsendu Sukanta Halder, Jean-François Lalonde, and Raoul de Charette. Physics-based rendering for improving robustness to rain. In *Proc. IEEE Int'l Conf. Computer Vision*, pages 10202–10211, 2019. [1](#), [3](#), [4](#), [7](#), [8](#)
- [14] Irtiza Hasan, Shengcai Liao, Jinpeng Li, Saad Ullah Akram, and Ling Shao. Generalizable pedestrian detection: The elephant in the room. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 11328–11337, June 2021. [3](#)
- [15] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 7132–7141, 2018. [5](#)
- [16] K. Jiang, Z. Wang, P. Yi, C. Chen, B. Huang, Y. Luo, J. Ma, and J. Jiang. Multi-scale progressive fusion network for single image deraining. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 8343–8352, 2020. [2](#)
- [17] L. W. Kang, C. W. Lin, and Y. H. Fu. Automatic single-image-based rain streaks removal via image decomposition. *IEEE Trans. on Image Processing*, 21(4):1742–1755, April 2012. [1](#), [2](#)
- [18] J. H. Kim, C. Lee, J. Y. Sim, and C. S. Kim. Single-image deraining using an adaptive nonlocal means filter. In *Proc. IEEE Int'l Conf. Image Processing*, pages 914–917, Sept 2013. [2](#)
- [19] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *Proc. Int'l Conf. Learning Representations*, 2017. [2](#)
- [20] Guanbin Li, Xiang He, Wei Zhang, Huiyou Chang, Le Dong, and Liang Lin. Non-locally enhanced encoder-decoder network for single image de-raining. In *ACM Trans. Multimedia*, pages 1056–1064, 2018. [2](#)
- [21] R. Li, L. Cheong, and R. T. Tan. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 1633–1642, 2019. [2](#)
- [22] Ruoteng Li, Loong-Fah Cheong, and Robby T. Tan. Single Image Deraining using Scale-Aware Multi-Stage Recurrent Network. *arXiv e-prints*, page arXiv:1712.06830, December 2017. [2](#)
- [23] R. Li, R. T. Tan, and L. F. Cheong. All in one bad weather removal using architectural search. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 3172–3182, 2020. [2](#)
- [24] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 510–519, 2019. [6](#)
- [25] Xia Li, Jianlong Wu, Zhouchen Lin, Hong Liu, and Hongbin Zha. Recurrent squeeze-and-excitation context aggregation net for single image deraining. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Proc. IEEE European Conf. Computer Vision*, volume 11211, pages 262–277, 2018. [3](#), [5](#)
- [26] Y. Li, R. T. Tan, X. Guo, J. Lu, and M. S. Brown. Rain streak removal using layer priors. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 2736–2744, 2016. [1](#), [2](#)
- [27] Y. Luo, Y. Xu, and H. Ji. Removing rain from a single image via discriminative sparse coding. In *Proc. IEEE Int'l Conf. Computer Vision*, pages 3397–3405, 2015. [1](#), [2](#)
- [28] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *Proc. Int'l Conf. Learning Representations*, 2018. [2](#), [3](#)
- [29] A. Mustafa, S. H. Khan, M. Hayat, J. Shen, and L. Shao. Image super-resolution as a defense against adversarial attacks. *IEEE Trans. on Image Processing*, 29:1711–1724, 2020. [2](#)

- [30] Dongwei Ren, Wangmeng Zuo, Qinghua Hu, Pengfei Zhu, and Deyu Meng. Progressive image deraining networks: A better and simpler baseline. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 3937–3946, 2019. [3](#)
- [31] Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. Is robustness the cost of accuracy? – a comprehensive study on the robustness of 18 deep image classification models. In *Proc. IEEE European Conf. Computer Vision*, pages 644–661, 2018. [1](#)
- [32] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proc. Int'l Conf. Learning Representations*, 2014. [2](#)
- [33] Hong Wang, Qi Xie, Qian Zhao, and Deyu Meng. A model-driven deep neural network for single image rain removal. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 3100–3109, 2020. [2](#), [3](#), [5](#)
- [34] Tianyu Wang, Xin Yang, Ke Xu, Shaozhe Chen, Qiang Zhang, and Rynson W.H. Lau. Spatial attentive single-image deraining with a high quality real rain dataset. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, June 2019. [2](#)
- [35] Xiaolong Wang, Ross B. Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. [6](#)
- [36] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. on Image Processing*, 13(4):600–612, 2004. [3](#)
- [37] W. Wei, D. Meng, Q. Zhao, Z. Xu, and Y. Wu. Semi-supervised transfer learning for image rain removal. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 3872–3881, 2019. [2](#), [3](#), [4](#), [7](#)
- [38] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. CBAM: convolutional block attention module. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Proc. IEEE European Conf. Computer Vision*, volume 11211, pages 3–19, 2018. [6](#)
- [39] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Yuyin Zhou, Lingxi Xie, and Alan Yuille. Adversarial examples for semantic segmentation and object detection. In *Proc. IEEE Int'l Conf. Computer Vision*, Oct 2017. [2](#)
- [40] Wenhan Yang, Jiaying Liu, Shuai Yang, and Zongming Guo. Scale-free single image deraining via visibility-enhanced recurrent wavelet learning. *IEEE Trans. on Image Processing*, 28(6):2948–2961, 2019. [2](#)
- [41] Wenhan Yang, Robby T. Tan, Jiashi Feng, Zongming Guo, Shuicheng Yan, and Jiaying Liu. Joint rain detection and removal from a single image with contextualized deep networks. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 42(6):1377–1393, 2020. [1](#), [2](#), [3](#)
- [42] Wenhan Yang, Robby T. Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 1685–1694, 2017. [3](#), [4](#), [5](#), [6](#), [7](#)
- [43] Wenhan Yang, Shiqi Wang, and Jiaying Liu. Removing arbitrary-scale rain streaks via fractal band learning with self-supervision. *IEEE Trans. on Image Processing*, 29:6759–6772, 2020. [2](#)
- [44] Wenhan Yang, Shiqi Wang, Dejia Xu, Xiaodong Wang, and Jiaying Liu. Towards scale-free rain streak removal via self-supervised fractal band learning. *Proc. AAAI Conf. on Artificial Intelligence*, pages 12629–12636, 2020. [2](#)
- [45] Rajeev Yasarla and Vishal M. Patel. Uncertainty guided multi-scale residual learning-using a cycle spinning CNN for single image de-raining. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 8405–8414, 2019. [2](#), [3](#)
- [46] R. Yasarla, V. A. Sindagi, and V. M. Patel. Syn2real transfer learning for image deraining using gaussian processes. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 2723–2733, 2020. [2](#)
- [47] Minghao Yin, Yongbing Zhang, Xiu Li, and Shiqi Wang. When deep fool meets deep prior: Adversarial attack on super-resolution network. In *ACM Trans. Multimedia*, page 1930–1938, 2018. [2](#)
- [48] Syed Waqas Zamir, Aditya Arora, Salman H. Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 14821–14831, 2021. [2](#), [3](#), [4](#), [5](#)
- [49] Liming Zhai, Felix Juefei-Xu, Qing Guo, Xiaofei Xie, Lei Ma, Wei Feng, Shengchao Qin, and Yang Liu. It's Raining Cats or Dogs? Adversarial Rain Attack on DNN Perception. *arXiv e-prints*, page arXiv:2009.09205, September 2020. [2](#)
- [50] He Zhang and Vishal M. Patel. Density-aware single image de-raining using a multi-stream dense network. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, June 2018. [2](#)
- [51] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *Proc. Int'l Conf. Machine Learning*, volume 97, pages 7472–7482, 09–15 Jun 2019. [6](#)
- [52] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, 2018. [3](#)
- [53] Yi Zhu, Karan Sapra, Fitsum A. Reda, Kevin J. Shih, Shawn D. Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pages 8856–8865, 2019. [3](#), [7](#)