

This ICCV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Generalized Real-World Super-Resolution through Adversarial Robustness

Angela Castillo^{*1}, María Escobar^{*1}, Juan C. Pérez^{1,2}, Andrés Romero³, Radu Timofte³, Luc Van Gool³, and Pablo Arbelaez¹

¹Center for Research and Formation in Artificial Intelligence, Universidad de los Andes, Colombia ²King Abdullah University of Science and Technology (KAUST), Saudi Arabia ³Computer Vision Lab, ETH Zürich, Switzerland

> ¹{a.castillo13, mc.escobar11, jc.perez13, pa.arbelaez}@uniandes.edu.co ³roandres@ethz.ch, {radu.timofte, vangool}@vision.ee.ethz.ch

Abstract

Real-world Super-Resolution (SR) has been traditionally tackled by first learning a specific degradation model that resembles the noise and corruption artifacts in lowresolution imagery. Thus, current methods lack generalization and lose their accuracy when tested on unseen types of corruption. In contrast to the traditional proposal, we present Robust Super-Resolution (RSR), a method that leverages the generalization capability of adversarial attacks to tackle real-world SR. Our novel framework poses a paradigm shift in the development of real-world SR methods. Instead of learning a dataset-specific degradation, we employ adversarial attacks to create difficult examples that target the model's weaknesses. Afterward, we use these adversarial examples during training to improve our model's capacity to process noisy inputs. We perform extensive experimentation on synthetic and real-world images and empirically demonstrate that our RSR method generalizes well across datasets without re-training for specific noise priors. By using a single robust model, we outperform state-of-theart specialized methods on real-world benchmarks.

1. Introduction

Super-Resolution (SR) is the task of increasing the resolution of a given image. The ever-growing use of deep learning has fostered the creation of SR models that obtain remarkable results with high fidelity on traditional SR benchmarks [25, 27, 31, 41, 48, 58]. These conventional models are trained in a supervised manner with a High-Resolution (HR) and a corresponding Low-Resolution (LR)



Figure 1. We present the comparison between our method and state-of-the-art methods: Impressionism [18] and ESRGAN-FS [12], for two different types of degradations: NTIRE_{*syn*} [29] and AIM_{*syn*} [30]. For reference, we show the bicubically up-sampled input, the result of a supervised SISR method (ESR-GAN [48]), and the ground-truth (GT). Blue frames denote training and validation on the same dataset. Red frames denote training and validation on different datasets. Green frames denote our method.

image pair. Since capturing the exact same scene in both HR and LR is complex and time-consuming, traditional SR datasets use clean HR images and the LR images are usually generated through a bicubic down-sampling operation [1, 45]. Using a bicubic kernel on clean HR images simplifies the ill-posed SR task because it ignores the fact that real-world images are subjected to sensor noise and artifacts. Thus, models trained with clean, paired datasets tend to underperform when evaluated on real-life scenes.

This limitation motivated the study of real-world SR on datasets with synthetic and natural corruptions [29, 30]. Several benchmarks design real-world artifacts and corrup-

^{*}equal contribution

tions under different assumptions or from varying sensors, and current solutions have to produce a specific model for each one. Consequently, state-of-the-art models for real-world SR [12, 18] generate photo-realistic results only when they are evaluated with the particular dataset for which they were trained, failing to generalize to new datasets with unseen corruptions. As an example, Fig. 1 shows two of the best methods in real-world SR, Impressionism [18] and ESRGAN-FS [12], evaluated on two benchmark datasets with different artificial corruptions. When evaluated on the image from the same training dataset (blue frames), both methods show excellent performance. However, they create significant artifacts in the image from the unseen dataset (red frames).

Even though the different noise in each of the input images does not modify human perceptual recognition (*i.e.*, we can easily discriminate the introduced artifacts in both cases), state-of-the-art models are not *robust* to unseen noise. To the best of our knowledge, there is no single realworld SR method that can generalize across different types of noise.

A way to assess a model's robustness to unseen noise is through Adversarial Attacks [32]. These attacks are based on adversarial examples, which are small intensity perturbations to the input image specifically designed to trick the model into failure cases. Adversarial examples have demonstrated improvement in noise generalization of models for various tasks such as classification [32, 44], semantic segmentation [3, 9, 49] and object detection [54]. Through robust training [32], models learn to be invariant to noise by creating more human-perception aligned filters [11].

Choi *et al.* [6] explored the effect of adversarial attacks on traditional SR methods, finding that these methods are brittle and easy to fool under adversarial examples. In more recent work, Choi *et al.* [7] proposed an adversarial defense by modifying the intermediate filters of the network, thus improving the performance of traditional SR models under their previous attacks. However, their work was focused on evaluating the methods based on pixel-wise metrics. To date, there are no works that report the study of robust training for the real-world SR problem.

In this work, we leverage adversarial attacks to create a real-world SR model robust to unseen noise types. As illustrated in Fig. 2b, we start with a standard SR model trained on a clean dataset and find adversarial examples of the LR input that damage the performance of the model. Simultaneously, we employ the created examples to teach the network that small perturbations in the LR image should result in the same HR ground-truth. Our approach presents a paradigm shift in the study of real-world SR because, instead of generating LR images with a type of noise that resembles that of a specific dataset, we focus on making the SR model more robust to any input. Fig. 1 shows the generalization capability of our method. Even though our model has never seen images from either of the artificially corrupt datasets, by applying robust training, we are able to obtain photo-realistic HR images for both types of input noise.

We perform extensive quantitative and qualitative evaluations on data from the NTIRE 2020 Challenge on Real-World Image Super-Resolution [29] and the AIM 2019 Real World Super-Resolution Challenge [30]. We show that our *single* model, trained robustly on clean images, achieves state-of-the-art performance compared to methods with specialized models for each specific dataset. We also provide qualitative evaluation and non-reference perceptual metrics for images retrieved from an iPhone 3 [17] and facial superresolution in the wild [51].

Our main contribution is twofold: first, we propose a novel use of adversarial attacks in real-world superresolution. Second, through the use of adversarial examples, we are able to create a generalized real-world SR model that achieves state-of-the-art results without training or fine-tuning on corrupt or real-world datasets. The code for reproducing our results is available at https://github.com/BCV-Uniandes/RSR.

2. Related work

2.1. Image Super-Resolution

2.1.1 Single Image Super-Resolution

Single Image Super-Resolution (SISR) is a problem that has been widely studied in computer vision [8, 37, 50, 16]. Early approaches incorporated simple linear interpolation methods [21, 56] estimating the correlation [2] or covariance of the low-resolution data [26]. Nevertheless, those methods fell short to fully capture the high frequencies of the images.

More recent methods have focused on improving the model's architecture to have better performance, either optimizing for pixel-wise scores (e.g., PSNR-oriented) or GAN-based [13] perceptual scores (e.g., LPIPSoriented [57]). SISR is very popular [25, 27, 31, 41, 48, 58] in its simplified form, where it is relatively easy to create paired datasets using a known downsampling strategy (e.g., bicubic interpolation). However, assuming prior knowledge of the downsampling kernel is unrealistic for images in the wild, as the kernel and the type of degradation are unknown, and therefore, SISR models fail to generalize. Moreover, images in the wild might contain corruptions and artifacts that should be removed in the super-resolution process. Therefore, blind SR approaches [4, 34, 15, 55, 38] tackle unknown degradation operations, whereas unsupervised real-world approaches [28, 33, 5, 53, 22, 39, 60] tackle severe degradation or compression artifacts in lowresolution images. Since our approach aims at performing super-resolution robust to noise and artifacts, we focus on the real-world literature.

2.1.2 Real-World Super-Resolution

Real-world images contain noise or corruptions inherent to the sensor acquisition, or data compression artifacts. Traditional SR techniques [48, 58, 25] fail at this task because they enhance the artifacts. Due to the problem formulation, it is hard to get paired real-world low-resolution and high-resolution alignment, so this problem is traditionally tackled in an unsupervised fashion by first learning a degradation procedure that can be used to produce an artificially paired dataset. As using paired datasets has been widely studied by SISR approaches, real-world methods focus on the degradation stage, that is, finding the best way to inject noise that resembles the one in the low-resolution dataset (Fig. 2a).

2.2. Adversarial Attacks

Adversarial attacks on Deep Neural Networks (DNNs) were firstly conducted by Szegedy et al. [43]. Their results showed that DNNs could be fooled by imperceptible perturbations that do not modify the image semantics. Particularly, Choi et al. [6, 7] demonstrated that SR methods are susceptible to adversarial attacks. Formally, an adversarial attack is a procedure that, given an image, produces a poisoned input image that the DNN misclassifies with high confidence [14]. Large efforts have been devoted to developing models that are both accurate *and* robust to attacks. Adversarial Training (AT) [32], a technique that has stood the test of time, casts the problem of training adversariallyrobust models as one of robust optimization of a saddlepoint problem. In practice, AT trains a DNN on adversarial examples generated on-the-fly through Projected Gradient Descent (PGD) by aiming at maximizing the DNN's loss.

While much of the initial concern regarding the brittleness of DNNs against attacks was related to security concerns, recent evidence provides new perspectives. Stutz et al. [42] explored how adversarial examples may be related to generalization in learning algorithms. In relation to perception, Engstrom et al. [11] demonstrated that robust features are aligned with human perception, and hence adversarial robustness can serve as a prior for learned representations. Further, Santurkar et al. [40] showed that adversarially-robust classifiers prove to be useful for image synthesis tasks. Inspired by these findings, in this paper we demonstrate how adversarial robustness can be used as a useful prior for learning DNNs for real-world SR. In particular, we show that by using adversarial attacks we can bypass the need of corrupted low-resolution data, and train the robust super-resolution architecture directly, thus producing a single model that generalizes to several datasets.

To avoid ambiguity with Generative Adversarial Net-

works (GANs), we refer to Adversarial Training as Robust Training.

3. Method

Traditionally, real-world SR approaches design handcrafted noise models for learning dataset-specific degradation and using it to create new LR images to train their SR models (Fig. 2a). However, these models tend not to generalize well to new datasets with an unknown degradation distribution. Unlike traditional approaches, our method creates LR images that are perceptually challenging for the SR model in order to improve its generalization capability (Fig. 2b).

Inspired by recent works that have drawn a link between adversarial robustness and both generalization and human perception [11, 40], we exploit adversarial robustness as useful prior knowledge for learning SR models. We underscore that, as with any ill-posed problem such as SR [52], useful priors are of utmost importance: what may prove to be a reasonable output from a SR model strongly relies on vast amounts of prior information humans possess about the real world. Thus, in search of SR models with improved performance on unseen types of degradation, we leverage adversarially-robust training to induce priors correlated with human perception. In particular, we encourage the training procedure to prefer robust solutions against adversarial attacks during the learning of SR models.

3.1. Mathematical Formulation

Super Resolution. Let I_{HR} be a HR image, and I_{LR} be a LR image downsampled \downarrow_s by a factor of s as

$$I_{LR} = (I_{HR} \otimes k) \downarrow_s + n, \tag{1}$$

where k is the downsampling kernel, and n represents additive noise. In common SR methods [48, 25], this formulation is simplified by using a bicubic kernel and small independent identically distributed (*i.i.d*) additive Gaussian noise [59]. In real-world approaches, n is not neglected, as the corruptions severely alter the low-resolution image. To circumvent this issue, we optimize n for each image using adversarial attacks. Without loss of generality, our model can be extended to blind approaches by considering unknown kernels k.

Our rationale to include adversarial examples that resemble real-world LR images is as follows: Employing GANbased methods that include a generator to inject noise [5, 28] introduces a domain shift that compromises the performance during inference. Hand-crafted methods [18, 12] do not introduce a domain shift because of their simplicity (Fig. 2a) but require extensive manual labor to simulate different types of corruption. By using adversarial attacks, we can modify the injected noise and thus optimize it to find the



Figure 2. **Traditional vs. Robust Super-Resolution.** In (*a*) *Traditional Real-world Training*, methods use hand-crafted noise models to learn dataset-specific degradations. In contrast, our (*b*) *Robust Training* strategy employs optimized adversarial examples to train a single SR model robustly.



Figure 3. Noise Optimization. We create adversarial examples by finding iteratively a noise that exploits the weaknesses of the generator. Once the t iterations of the optimization process are complete, the adversarial examples are used to improve our SR model during our *Robust training*.

hardest type of noise in a data-driven way (Fig. 2b) while maintaining the semantic information present in an image. With this approach, we avoid a domain shift and present an optimized method to faithfully remove noise and corruption robustly.

Robust Training. In a fashion similar to [32], we generate adversarial examples on-the-fly via the Projected Gradient Descent (PGD) method. PGD is an optimization procedure for finding ℓ_p -bounded adversarial examples. In this work, we focus in particular on ℓ_{∞} -bounded examples, computed as repeated iterations of

$$x_{t+1}^{adv} = \prod_{X} \left(x_t + \alpha \operatorname{sign} \left(\nabla_{x_t} L(x_t, y) \right) \right), \qquad (2)$$

where \prod is the projection operator, x_{t+1} is the input image of the attack iteration t+1, y is the ground-truth, X denotes the set of *valid* images (*i.e.*, the intersection of the ℓ_{∞} -ball of ϵ radius and the usual [0, 1] range for image pixels), α is the size of the gradient step and $L(\cdot)$ is the cost function with respect to which optimization is performed. For the first iteration, x_0 is perturbed with random noise of magnitude ϵ and projected back into X, as customary in previous works [14, 24]. In our method, we use PGD to maximize both the L_1 and perceptual (L_{percep}) losses of the generator. Furthermore, including adversarial attacks within the SR framework requires significantly more ingenuity than just replacing the loss function. Specifically, besides setting an objective function for the target module, designing adversarial attacks during training requires defining (*i*) how to maximize such objective, (*ii*) the constraints to maximize the optimization variables, and (*iii*) when/where to introduce adversarial examples.

3.2. Robust Training on Super-Resolution

The underlying idea of robust training is to continually include adversarial examples during the training stage of the model. Our Robust Super-Resolution (RSR) model extends this idea to SR through three main steps. Firstly, we implement a GAN-based model for SR that is previously trained on clean LR images. Secondly, we employ the trained generator to create adversarial examples through a noise optimization process (Fig. 3). Finally, we use the adversarial examples as LR inputs to further train the GAN robustly (Fig. 2b).

GAN-based model. We build upon ESRGAN [48, 47], a GAN-based approach to SR [13]. The Generator (G) is built upon the RRDBNet architecture [25] with Residualin-Residual Dense Blocks to improve the enhanced image's quality. G optimizes three losses: (i) the L_1 loss to evaluate the pixel distance of the super-resolved image and the HR, (ii) the L_{percep} [19] which is calculated as the distance of activation features of a pre-trained network between superresolved and HR images, (iii) and the term of the adversarial loss for the generator L_G^{GAN} . L_1 and L_{percep} are minimized to reduce the distance between the SR and HR images. The objective function for G is:

$$L_G = L_1 + L_{percep} + L_G^{GAN}.$$
 (3)

The Discriminator (D) is based on a VGG-128 architecture and follows the same principle as in Relativistic GAN [20], which estimates the probability that a real image is more realistic than a fake image. During the first stage of our method, we use an ESRGAN model pre-trained on clean images, *i.e.*, the model can super-resolve clean images at this stage but it is not robust to any type of noise.

Noise Optimization. We use the pre-trained network G as the starting point for our noise optimization process. Our goal is to find the optimal input noise that, when added to a LR image, results in a deficient super-resolved image. This process allows us to identify the generator's weaknesses in handling inputs with different types of corruption. Fig. 3 depicts our optimization process. We freeze the weights of G during this stage and create the adversarial examples based on the L_1 and L_{percep} losses. We add an ϵ -constraint random noise to the LR input image to retrieve an adversarial example. Then, we optimize the noise through PGD (Eq. 2). We iteratively update the input adversarial noise by considering the contribution of both losses and produce the adversarial example as follows:

$$I_{LR_{t+1}}^{adv} = I_{LR_t} + \alpha \operatorname{sign} \left(\nabla_{I_{LR_t}} \left(L_1(G(I_{LR_t}), I_{HR}) + L_{percep}(G(I_{LR_t}), I_{HR}) \right) \right).$$

$$(4)$$

The noise on real-world images tends to be grouped in multiple pixels, so we take inspiration from [10] to design structured noise by changing the input scale. To train our model, we start the optimization process with a synthetically grouped structure to better simulate real-world corruptions. Nevertheless, during the optimization iterations of the attack, we no longer control the structure of the noise as it changes according to the steps of the gradient. Please refer to the supplementary material for additional analysis regarding the noise structure.

RSR Training. We use the adversarial examples created during the optimization stage as LR input to train the GAN further. Since the new inputs are specifically optimized to exploit the weaknesses of G, using them during training ensures that the network learns to super-resolve noisy images. Furthermore, since the optimized noise does not follow a pre-defined distribution, the network must generalize to different types of noise. The noise optimization and training stages are performed interspersed throughout the training iterations. Thus, as the training process of the network progresses, the examples that the network considers as adversarial are more potent after each new iteration. This process,

in turn, makes the network more robust against increasing degradation in every new iteration.

4. Experimental setup

4.1. Datasets

Training. The DIV2K dataset [1, 45] is one of the widely used benchmarks for traditional SR. The training set contains 800 2K resolution images and their corresponding LR images, obtained using a bicubic downgrading operator. These images do not include any simulated or real noise, and we use them as the basis for our robust training. For memory-saving purposes during training, we crop the images into 480×480 sub-images.

Reference validation. We validate our framework over the validation datasets from the NTIRE 2020 Challenge on Real-World Image Super-Resolution [29] track 1, and the AIM 2019 Real World Super-Resolution Challenge [30] track 2. The LR images in each dataset include a different degradation obtained with an undisclosed artificial operator, resembling thus realistic corruptions and artifacts. The validation set is composed of the artificially degraded version of the 100 LR images from the DIV2K validation set and their corresponding HR ground-truth to calculate referencebased metrics. Since both datasets include synthetic corruptions, we refer to them as NTIRE_{syn} and AIM_{syn}.

Non-reference validation. We also validate the performance of our method on real-world datasets that do not have a ground-truth HR for reference. We use the validation set of the DPED [17] dataset that contains real-world images taken by an iPhone3 camera. This dataset represents a more challenging framework because the data contains noise, differences in lighting, and other low-quality artifacts. Furthermore, we evaluate our method on real-world faces from the WIDER FACE dataset [51], initially designed for facial detection in-the-wild. We randomly sample 100 cropped faces of medium size with no occlusions. We refer to these datasets as DPED_{*rw*} and FACES_{*rw*}.

4.2. Evaluation metrics

As it is well known that pixel-wise metrics (*e.g.*, PSNR) do not correlate well with human perception, the metric we focus on is the Learned Perceptual Image Patch Similarity (LPIPS) [57]. This metric is based on the comparison between features of a neural network, in this case, a pre-trained AlexNet [23] model, and it measures the distance in perceptual quality between the generated image and the ground truth. We also report the Peak Signal-to-Noise Ratio (PSNR) measured in decibels (dB) and Structural Similarity Index (SSIM) to evaluate the pixel-wise fidelity of the result with the ground-truth HR image. For the datasets that do not include ground-truth HR images, we use the non-reference image quality metrics: Perception-based Image

ε	Iters	Noise structure	Loss	PSNR↑			SSIM↑			LPIPS↓		
				NTIRE _{syn}	AIM _{syn}	Avg	NTIRE _{syn}	AIM _{syn}	Avg	NTIRE _{syn}	AIM _{syn}	Avg
14/255	2	1.5	$L_1 + L_{percep}$	24.31	21.99	23.15	0.65	0.60	0.62	0.23	0.37	0.30
14/255	2	1.5	L_1	25.62	22.47	24.04	0.68	0.63	0.65	0.30	0.40	0.35
14/255	2	1.5	$\overline{L_{percep}}$	24.20	22.07	23.13	0.64	0.61	0.63	0.24	0.38	0.31
10/255	2	1.5	$L_1 + L_{percep}$	23.16	21.84	22.50	0.56	0.59	0.58	0.31	0.37	0.34
<u>12/255</u>	2	1.5	$L_1 + L_{percep}$	23.68	21.74	22.71	0.62	0.59	0.60	0.25	0.37	0.31
<u>16/255</u>	2	1.5	$L_1 + L_{percep}$	23.84	22.06	22.95	0.63	0.62	0.62	0.25	0.38	0.31
<u>18/255</u>	2	1.5	$L_1 + L_{percep}$	24.02	21.89	22.96	0.64	0.61	0.62	0.25	0.38	0.31
14/255	4	1.5	$L_1 + L_{percep}$	24.13	22.04	23.09	0.65	0.61	0.63	0.26	0.38	0.32
14/255	<u>6</u>	1.5	$L_1 + L_{percep}$	24.00	21.73	22.86	0.63	0.60	0.62	0.27	0.37	0.32
14/255	<u>8</u>	1.5	$L_1 + L_{percep}$	23.94	21.56	22.75	0.63	0.58	0.61	0.27	0.36	0.32
14/255	2	1	$L_1 + L_{percep}$	25.58	22.20	23.89	0.70	0.63	0.67	0.24	0.38	0.31
14/255	2	<u>2</u>	$L_1 + L_{percep}$	19.99	21.40	20.69	0.32	0.54	0.43	0.54	0.41	0.47

Table 1. Ablation Study. Performance comparison of varying the hyper-parameters for the noise optimization stage of our method. ϵ denotes the maximum pixel intensity perturbation of the noise. *Iters* denotes the number iterations used to find the adversarial example. *Noise structure* denotes how grouped the initial noise values are. *Loss* indicates which loss function was set as objective for the noise optimization process. Red and blue colors highlight the best two scores. The hyper-parameters we use for our RSR model are in **bold**.

Quality Evaluator (PIQE) [46], Naturalness Image Quality Evaluator (NIQE) [36], and Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) [35]. We provide further validation in the supplementary material.

4.3. Implementation details

Robust training. We retrieve the adversarial examples using PGD with parameter values $\epsilon = 14/255$, iterations = 2, structured noise = $1.5 \times$. The ablation experiments for each of these parameters are presented in section 5.1. We initialize training with the pre-trained weights of ESRGAN on DIV2K. Then, we perform our robust training on two NVIDIA QUADRO RTX 800 GPUs for 18k iterations with an initial learning rate of $1e^{-4}$, 16 images per batch, and use an Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.99$.

Comparison with State-of-the-Art. We compare our method against ESRGAN [48] to establish a baseline against SISR models. We extend the ESRGAN framework for 18k iterations over the DIV2K dataset for a fair comparison. Furthermore, we compare our results against the best-performing method of the NTIRE20 and AIM19 challenges: Impressionism [18] and ESRGAN-FS [12], respectively. For Impressionism, we use the provided pre-trained weights for NTIRE_{syn} and DPED_{rw}, and for ESRGAN-FS, we use their weights for AIM_{syn} and DPED_{rw}. For a complete comparison in every dataset, we train the available code of Impressionism in AIM_{syn} and ESRGAN-FS in NTIRE_{syn}, using default parameters in their official public repositories. All our experiments are done considering a scaling factor of $4 \times$.

5. Results and Discussion

This section analyzes the quantitative and qualitative results obtained by our method, and we perform a thorough comparison with the state-of-the-art methods.

5.1. Ablation study

Robust training relies on the preservation of the perceptual and semantic information of the original image. Creating hard examples for the network that are unrealistic compared to the training dataset would confuse the network rather than help encourage adversarial robustness. Thus, we have to find a trade-off between how hard the example is and how realistic it is. In Table 1, we present an extensive validation of the effect of each hyper-parameter to demonstrate that our final method uses the optimal one. Since our method aims to improve the perceptual quality, we focus mainly on LPIPS and report PSNR and SSIM as complementary metrics. Please refer to the supplementary material for qualitative visualizations and a more in-depth analysis of the different hyper-parameters.

Choosing the right loss function to optimize is crucial because it directs the robust training towards creating hard examples for a part of the network. Using only the L_1 loss for robust adversarial optimization improves the pixel-wise metrics. This result is expected because we are specifically training the network to be more precise in a pixel-wise comparison. However, the perceptual quality is affected. In the same way, optimizing only the perceptual loss improves the LPIPS compared to the previous experiment. Nevertheless, when we use both losses, we are able to improve the average perceptual performance further.

 ϵ constrains the magnitude of the perturbation in the adversarial example. Larger values of ϵ result in larger perturbations as it allows the noise to be in a higher range, thus modifying the input more. Table 1 shows that performing smaller modifications to the input, *i.e.*, having a smaller ϵ , greatly affects the mean LPIPS. However, the perceptual quality remains virtually the same when ϵ is large enough.

Regarding the iterative nature of the attack, increasing the number of iterations leads to a worse result in every dimension of the evaluation. This result suggests that an



Figure 4. **Qualitative results on Synthetic Images.** Comparison between our method and state-of-the-art methods, for two synthetic corruption datasets: $NTIRE_{syn}$ and AIM_{syn} . For reference, we show the bicubically upsampled input, the result of a supervised SISR method (ESRGAN [48]), and the ground-truth (GT). Blue frames denote training and validation on the same dataset. Red frames denote training and validation on different datasets. Green frames denote our method.

adversarial attack with two iterations is strong enough to impact the SR model positively.

Finally, the scale of the structured noise determines how grouped the initial noise is. Using a value of 1 means that every noisy pixel of the LR image has a different value. Table 1 shows that a structured noise of 1 performs well for PSNR and SSIM but slightly worse for LPIPS. This phenomenon is the same that we find on the weakest ϵ modification. In contrast, the results suggest that using a value of 2 in the structure scale creates adversarial examples that are too different from the original input, so the network can not effectively use them for training. We find that the best trade-off is achieved with a structure scale of 1.5.

Effect of robust training. Table 2 shows that ESRGAN achieves the worst LPIPS for both NTIRE_{syn} and AIM_{syn}. This behavior could result from training the model on clean images, depriving the model of seeing any corruptions like those present on the datasets. In contrast, RSR outperforms ESRGAN for every metric and has an average LPIPS 0.35 points lower. These results are visually confirmed in Fig. 4. ESRGAN magnifies the input noise of every image while RSR removes it. Since we use ESRGAN as our baseline, the only difference between that model and ours is that we perform robust training. Thus, even though we train with a

clean dataset, the invariance to noise enforced during robust training significantly improves the generalization capacity of the model when evaluated on corrupt datasets.

5.2. Real-World Synthetic Datasets

The winners of the two recent real-world SR challenges employ noise-specific models, i.e., different models depending on the type of corruption. To compute their generalization capabilities, in Table 2, we evaluate each specialized model on both synthetic datasets NTIRE_{syn} and AIM_{sun} and rank them by their average performance. Stateof-the-art models perform well on the dataset they were trained on but have lower performance on unseen corruption statistics. Given that we use a *single* model and it is not statistically-dependent on the type of artifacts, our generalization capabilities are considerably better than those of previous works. Fig. 4 further reinforces this notion by showing the artifacts that state-of-the-art models create on unseen datasets (red frames on the Fig.). See the supplementary material for an in-depth study of the type of artifacts that each training dataset enforces. In comparison, we have a single robust model that achieves the best LPIPS for both unseen datasets and successfully removes input noise without creating artifacts.

Table 2. Quantitative comparison on Synthetic Images. We report the comparison of reference metrics between our method and the state-of-theart methods in different datasets. \uparrow and \downarrow indicate higher is better and lower is better, respectively. Red and blue colors highlight the best two scores. Bold represents the best method for LPIPS metric.

Method	Training	PSNR↑			S	SIM↑		LPIPS↓		
withiou	Dataset	NTIRE _{syn}	AIM _{syn}	Avg	NTIREsyn	AIM _{syn}	Avg	NTIRE _{syn}	AIM _{syn}	Avg
Bicubic	-	25.51	22.35	23.93	0.67	0.62	0.65	0.63	0.68	0.66
	NTIRE _{syn}	24.82	21.47	23.15	0.66	0.54	0.60	0.23	0.52	0.37
Impressionism [18]	AIM _{syn}	19.65	21.89	20.77	0.29	0.60	0.45	0.67	0.41	0.54
	DPED _{rw}	17.53	18.84	18.18	0.34	0.49	0.41	0.60	0.47	0.53
	NTIRE _{syn}	24.59	22.07	23.33	0.69	0.63	0.66	0.25	0.47	0.36
ESRGAN-FS [12]	AIM _{syn}	19.56	20.82	20.19	0.31	0.51	0.41	0.56	0.39	0.48
	$DPED_{rw}$	17.79	20.15	18.97	0.34	0.53	0.43	0.51	0.47	0.49
ESRGAN [48]	DIV2K	20.59	21.48	21.03	0.43	0.56	0.49	0.68	0.53	0.60
RSR (Ours)	DIV2K	24.31	21.99	23.15	0.65	0.60	0.62	0.23	0.37	0.30



Figure 5. Qualitative comparison on real-world images. Comparison between our method and state-of-the-art methods, for two real-world datasets: DPED_{rw} and FACES_{rw}. Note that these datasets do not have ground-truth. Blue frames denote training and validation on the same dataset. Red frames denote training and validation on different datasets. Green frames denote our method.

5.3. Real-World Datasets

Fig. 5 shows the output of different methods for two challenging in-the-wild datasets: DPED_{rw} [17] and FACES_{rw} [51]. A traditional SISR method [48] and specialized real-world models lack removing unseen real-world noise. Note that Impressionism trained on DPED_{rw} creates sharp super-resolved images on the same dataset (blue frames), but produces less pleasant images with stronger artifacts in FACES_{rw}. This result suggests that, even with a real-world noise training, there is no guarantee the model generalizes well towards images in the wild. Conversely, our model removes real-world noise from both datasets, having better performance on input images with clear edges (*e.g.*, street sign from DPED_{rw}). Thus, we show that our *single* robust model accurately super-resolves unseen synthetic and natural real-world images.

6. Conclusion

In this work, we explore the use of adversarial attacks to improve the robustness to unseen noise in the task of realworld SR. We present RSR, a novel SR method that leverages robust adversarial examples to create photo-realistic HR images regardless of the LR input noise. We evaluate our method on synthetic and natural real-world SR datasets. By using a *single* robust model trained only with a clean dataset, we outperform current state-of-the-art methods requiring specialized models for each type of corruption. Furthermore, we provide theoretical insights on how to adapt adversarial attacks for the particular needs of real-world SR models. We expect our work to catalyze further study of the fruitful relationship between adversarial robustness and real-world SR.

Acknowledgements: We thank Guillaume Jeanneret for insightful discussions on the subject.

References

- Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pages 126–135, 2017. 1, 5
- [2] Jan Allebach and Ping Wah Wong. Edge-directed interpolation. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 3, pages 707–710. IEEE, 1996.
 2
- [3] Anurag Arnab, Ondrej Miksik, and Philip HS Torr. On the robustness of semantic segmentation models to adversarial attacks. In *CVPR*, 2018. 2
- [4] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. Blind super-resolution kernel estimation using an internal-gan. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2
- [5] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *Proceedings of the European conference on computer vision (ECCV)*, pages 185–200, 2018. 2, 3
- [6] Jun-Ho Choi, Huan Zhang, Jun-Hyuk Kim, Cho-Jui Hsieh, and Jong-Seok Lee. Evaluating robustness of deep image super-resolution against adversarial attacks. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 303–311, 2019. 2, 3
- [7] Jun-Ho Choi, Huan Zhang, Jun-Hyuk Kim, Cho-Jui Hsieh, and Jong-Seok Lee. Adversarially robust deep image superresolution using entropy regularization. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 2, 3
- [8] Debabrata Chowdhuri, KS Sendhil Kumar, M Rajasekhara Babu, and Ch Pradeep Reddy. Very low resolution face recognition in parallel environment. *IJCSIT*) International Journal of Computer Science and Information Technologies, 3(3):4408–4410, 2012. 2
- [9] Moustapha Cisse, Yossi Adi, Natalia Neverova, and Joseph Keshet. Houdini: Fooling deep structured prediction models. arXiv preprint arXiv:1707.05373, 2017. 2
- [10] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pages 1310– 1320. PMLR, 2019. 5
- [11] Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Brandon Tran, and Aleksander Madry. Adversarial robustness as a prior for learned representations. arXiv preprint arXiv:1906.00945, 2019. 2, 3
- [12] Manuel Fritsche, Shuhang Gu, and Radu Timofte. Frequency separation for real-world super-resolution. In 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pages 3599–3608. IEEE, 2019. 1, 2, 3, 6, 8
- [13] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. arXiv preprint arXiv:1406.2661, 2014. 2, 4

- [14] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572, 2014. 3, 4
- [15] Jinjin Gu, Hannan Lu, Wangmeng Zuo, and Chao Dong. Blind super-resolution with iterative kernel correction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1604–1613, 2019. 2
- [16] Yong Guo, Jian Chen, Jingdong Wang, Qi Chen, Jiezhang Cao, Zeshuai Deng, Yanwu Xu, and Mingkui Tan. Closedloop matters: Dual regression networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [17] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings* of the IEEE International Conference on Computer Vision, pages 3277–3285, 2017. 2, 5, 8
- [18] Xiaozhong Ji, Yun Cao, Ying Tai, Chengjie Wang, Jilin Li, and Feiyue Huang. Real-world super-resolution via kernel estimation and noise injection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 466–467, 2020. 1, 2, 3, 6, 8
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 694–711. Springer, 2016. 4
- [20] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard gan. In *International Conference on Learning Representations*, 2018. 5
- [21] Robert Keys. Cubic convolution interpolation for digital image processing. *IEEE transactions on acoustics, speech, and signal processing*, 29(6):1153–1160, 1981. 2
- [22] Gwantae Kim, Jaihyun Park, Kanghyu Lee, Junyeop Lee, Jeongki Min, Bokyeung Lee, David K Han, and Hanseok Ko. Unsupervised real-world super resolution with cycle generative adversarial network and domain discriminator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 456–457, 2020. 2
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25:1097–1105, 2012. 5
- [24] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. arXiv preprint arXiv:1611.01236, 2016. 4
- [25] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photorealistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 1, 2, 3, 4
- [26] Xin Li and Michael T Orchard. New edge-directed interpolation. *IEEE transactions on image processing*, 10(10):1521– 1527, 2001. 2
- [27] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single

image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 1, 2

- [28] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Unsupervised learning for real-world super-resolution. In 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pages 3408–3416. IEEE, 2019. 2, 3
- [29] Andreas Lugmayr, Martin Danelljan, and Radu Timofte. Ntire 2020 challenge on real-world image super-resolution: Methods and results. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pages 494–495, 2020. 1, 2, 5
- [30] Andreas Lugmayr, Martin Danelljan, Radu Timofte, Manuel Fritsche, Shuhang Gu, Kuldeep Purohit, Praveen Kandula, Maitreya Suin, AN Rajagoapalan, Nam Hyung Joon, et al. Aim 2019 challenge on real-world image super-resolution: Methods and results. In 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pages 3575–3583. IEEE, 2019. 1, 2, 5
- [31] Andreas Lugmayr, Martin Danelljan, Luc Van Gool, and Radu Timofte. Srflow: Learning the super-resolution space with normalizing flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 715–732. Springer, 2020. 1, 2
- [32] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. 2, 3, 4
- [33] Shunta Maeda. Unpaired image super-resolution using pseudo-supervision. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 291–300, 2020. 2
- [34] Tomer Michaeli and Michal Irani. Nonparametric blind super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 945–952, 2013. 2
- [35] Anish Mittal, Anush K Moorthy, and Alan C Bovik. Blind/referenceless image spatial quality evaluator. In 2011 conference record of the forty fifth asilomar conference on signals, systems and computers (ASILOMAR), pages 723– 727. IEEE, 2011. 6
- [36] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. 6
- [37] Jinshan Pan, Sifei Liu, Deqing Sun, Jiawei Zhang, Yang Liu, Jimmy Ren, Zechao Li, Jinhui Tang, Huchuan Lu, Yu-Wing Tai, et al. Learning dual convolutional neural networks for low-level vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3070–3079, 2018. 2
- [38] Dongwei Ren, Kai Zhang, Qilong Wang, Qinghua Hu, and Wangmeng Zuo. Neural blind deconvolution using deep priors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2020. 2
- [39] Haoyu Ren, Amin Kheradmand, Mostafa El-Khamy, Shuangquan Wang, Dongwoon Bai, and Jungwon Lee. Realworld super-resolution using generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on*

Computer Vision and Pattern Recognition Workshops, pages 436–437, 2020. 2

- [40] Shibani Santurkar, Andrew Ilyas, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Image synthesis with a single (robust) classifier. In Advances in Neural Information Processing Systems (NeurIPS), 2019. 3
- [41] Assaf Shocher, Nadav Cohen, and Michal Irani. "zero-shot" super-resolution using deep internal learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3118–3126, 2018. 1, 2
- [42] David Stutz, Matthias Hein, and Bernt Schiele. Disentangling adversarial robustness and generalization. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2019. 3
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199, 2013. 3
- [44] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. 2
- [45] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, Lei Zhang, Bee Lim, et al. Ntire 2017 challenge on single image super-resolution: Methods and results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017. 1, 5
- [46] N Venkatanath, D Praneeth, Maruthi Chandrasekhar Bh, Sumohana S Channappayya, and Swarup S Medasani. Blind image quality evaluation using perception based features. In 2015 Twenty First National Conference on Communications (NCC), pages 1–6. IEEE, 2015. 6
- [47] Xintao Wang, Ke Yu, Kelvin C.K. Chan, Chao Dong, and Chen Change Loy. Basicsr. https://github.com/ xinntao/BasicSR, 2020. 4
- [48] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 1, 2, 3, 4, 6, 7, 8
- [49] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Yuyin Zhou, Lingxi Xie, and Alan Yuille. Adversarial examples for semantic segmentation and object detection. In *ICCV*, 2017. 2
- [50] Qingxiong Yang, Ruigang Yang, James Davis, and David Nistér. Spatial-depth super resolution for range images. In 2007 IEEE Conference on Computer Vision and Pattern Recognition, pages 1–8. IEEE, 2007. 2
- [51] Shuo Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 5, 8
- [52] Wenming Yang, Xuechen Zhang, Yapeng Tian, Wei Wang, Jing-Hao Xue, and Qingmin Liao. Deep learning for single image super-resolution: A brief review. *IEEE Transactions* on Multimedia, 2019. 3

- [53] Yuan Yuan, Siyuan Liu, Jiawei Zhang, Yongbing Zhang, Chao Dong, and Liang Lin. Unsupervised image superresolution using cycle-in-cycle generative adversarial networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pages 701–710, 2018. 2
- [54] Haichao Zhang and Jianyu Wang. Towards adversarially robust object detection. In *ICCV*, 2019. 2
- [55] Kai Zhang, Luc Van Gool, and Radu Timofte. Deep unfolding network for image super-resolution. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3217–3226, 2020. 2
- [56] Lei Zhang and Xiaolin Wu. An edge-guided image interpolation algorithm via directional filtering and data fusion. *IEEE transactions on Image Processing*, 15(8):2226–2238, 2006.
 2
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 2, 5
- [58] Wenlong Zhang, Yihao Liu, Chao Dong, and Yu Qiao. Ranksrgan: Generative adversarial networks with ranker for image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3096– 3105, 2019. 1, 2, 3
- [59] Ningning Zhao, Qi Wei, Adrian Basarab, Nicolas Dobigeon, Denis Kouamé, and Jean-Yves Tourneret. Fast single image super-resolution using a new analytical solution for $\ell_2 - \ell_2$ problems. *IEEE Transactions on Image Processing*, 25(8):3683–3697, 2016. 3
- [60] Yuanbo Zhou, Wei Deng, Tong Tong, and Qinquan Gao. Guided frequency separation network for real-world superresolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 428–429, 2020. 2