

Towards Adversarially Robust Deep Image Denoising

Hanshu Yan¹, Jingfeng Zhang², Jiashi Feng³, Masashi Sugiyama^{2,4} and Vincent Y. F. Tan^{5,1}

¹ECE, NUS

²RIKEN-AIP

³ByteDance Inc.

⁴GSFS, UTokyo

⁵Math, NUS

hanshu.yan@u.nus.edu

Abstract

This work systematically investigates the adversarial robustness of deep image denoisers (DIDs), i.e, how well DIDs can recover the ground truth from noisy observations degraded by adversarial perturbations. Firstly, to evaluate DIDs’ robustness, we propose a novel adversarial attack, namely Observation-based Zero-mean Attack (OBSATK), to craft adversarial zero-mean perturbations on given noisy images. We find that existing DIDs are vulnerable to the adversarial noise generated by OBSATK. Secondly, to robustify DIDs, we propose an adversarial training strategy, hybrid adversarial training (HAT), that jointly trains DIDs with adversarial and non-adversarial noisy data to ensure that the reconstruction quality is high and the denoisers around non-adversarial data are locally smooth. The resultant DIDs can effectively remove various types of synthetic and adversarial noise. We also uncover that the robustness of DIDs benefits their generalization capability on unseen real-world noise. Indeed, HAT-trained DIDs can recover high-quality clean images from real-world noise even without training on real noisy data. Extensive experiments on benchmark datasets, including Set68, PolyU, and SIDD, corroborate the effectiveness of OBSATK and HAT.

1 Introduction

Image denoising, which aims to reconstruct clean images from their noisy observations, is a vital part of the image processing systems. The noisy observations are usually modeled as the addition between ground-truth images and zero-mean noise maps [Dabov *et al.*, 2007; Zhang *et al.*, 2017]. Recently, deep learning-based methods have made significant advancements in denoising tasks [Zhang *et al.*, 2017; Anwar and Barnes, 2019] and have been applied in many areas including medical imaging [Gondara, 2016] and photography [Abdelhamed *et al.*, 2018]. Despite the success of deep denoisers in recovering high-quality images from a

certain type of noisy images, we still lack knowledge about their robustness against adversarial perturbations, which may cause severe safety hazards in high-stake applications like medical diagnosis. To address this problem, the first step should be developing attack methods dedicated for denoising to evaluate the robustness of denoisers. In contrast to the attacks for classification [Goodfellow *et al.*, 2015; Madry *et al.*, 2018], attacks for denoising should consider not only the adversarial budget but also some assumptions of natural noise, such as zero-mean, because certain perturbations, such as adding a constant value, do not necessarily result in visual artifacts. Although Choi *et al.* [2021; 2019] studied the vulnerability for various deep image processing models, they directly applied the attack from classification. To the best of our knowledge, no attacks are truly dedicated for the denoising task till now.

To this end, we propose the observation-based zero-mean attack (OBSATK), which crafts a worst-case zero-mean perturbation for a noisy observation by maximizing the distance between the output and the ground-truth. To ensure that the perturbation satisfies the adversarial budget and the zero-mean constraints, we utilize the classical projected-gradient-descent (PGD) [Madry *et al.*, 2018] method for optimization, and develop a two-step operation to project the perturbation back into the feasible region. Specifically, in each iteration, we first project the perturbation onto the zero-mean hyperplane. Then, we linearly rescale the perturbation to adjust its norm to be less or equal to the adversarial budget. We examine the effectiveness of OBSATK on several benchmark datasets and find that deep image denoisers are indeed susceptible to OBSATK: the denoisers cannot remove adversarial noise completely and even yield atypical artifacts, as shown in Figure 2g.

To robustify deep denoisers against adversarial perturbations, we propose an effective adversarial training strategy, namely hybrid adversarial training (HAT), to train denoisers by using adversarially noisy images and non-adversarial noisy images together. The loss function of HAT consists of two terms. The first term ensures the reconstruction performance from common non-adversarial noisy images, and the second term ensures the reconstructions between non-adversarial and adversarial images to be close to each other. Thus, we can obtain denoisers that perform well on both non-adversarial noisy images and their adversarial perturbed ver-

Please refer to the full-length paper [Yan *et al.*, 2022] for the appendices, proofs, and codes.

sions. Extensive experiments on benchmark datasets verify the effectiveness of HAT.

Moreover, we reveal that adversarial robustness benefits the generalization capability to unseen types of noise, i.e., HAT can train denoisers for real-world noise removal only with synthetic noise sampled from common distributions like Gaussians. That is because OBSATK searches for the worst-case perturbations around different levels of noisy images, and training with adversarial data ensures the denoising performance on various types of noise. In contrast, other reasonable methods for real-world denoising [Guo *et al.*, 2019; Lehtinen *et al.*, 2018] mostly require a large number of real-world noisy data for the training, which are unfortunately not available in some applications like medical radiology. We conduct experiments on several real-world datasets. Numerical and visual results demonstrate the effectiveness of HAT for real-world noise removal.

In summary, there are three main contributions in this work: 1) We propose a novel attack, OBSATK, to generate adversarial examples for noisy observations, which facilitates the evaluation of the robustness of deep image denoisers. 2) We propose an effective adversarial training strategy, HAT, for robustifying deep image denoisers. 3) We build a connection between adversarial robustness and the generalization to unseen noise, and show that HAT serves as a promising framework for training generalizable deep image denoisers.

2 Notation and Background

Adversarial robustness and adversarial training Consider a deep neural network (DNN) $\{f_\theta : \theta \in \Theta\}$ mapping an input \mathbf{y} to a target \mathbf{x} , the model is trained to minimize a certain loss function that is measured by particular distance $d(\cdot, \cdot)$ between output $f_\theta(\mathbf{y})$ and the target \mathbf{x} . In high stake applications, the DNN should resist small perturbations on the input data and map the perturbed input to a result close to the target. The notion of *robustness* has been proposed to measure the resistance of DNNs against the slight changes of the input [Szegedy *et al.*, 2014; Goodfellow *et al.*, 2015]. The robustness is characterized by the distance $d(f_\theta(\mathbf{y}'), \mathbf{x})$ between $f_\theta(\mathbf{y}')$ and target \mathbf{x} , where the worst-case perturbed input \mathbf{y}' is located within a small neighborhood of the original input \mathbf{y} and maximizes the distance between its output and target \mathbf{x} .

$$\mathbf{y}' = \arg \max_{\mathbf{y}' : \|\mathbf{y}' - \mathbf{y}\| \leq \rho} d(f_\theta(\mathbf{y}'), \mathbf{x}). \quad (1)$$

The worst-case perturbation \mathbf{y}' can be approximated via many adversarial attack methods, such as FGSM [Goodfellow *et al.*, 2015], I-FGSM [Kurakin *et al.*, 2017], and PGD [Madry *et al.*, 2018], which solve (1) via gradient descent methods. The distance $d(f_\theta(\mathbf{y}'), \mathbf{x})$ is an indication of the robustness of f_θ around \mathbf{y} : a small distance implies strong robustness and vice versa. In terms of image classification, the ρ -neighborhood is usually defined by the ℓ_∞ -norm and the distance $d(\cdot, \cdot)$ is measured by the cross-entropy loss [Madry *et al.*, 2018] or a margin loss [Carlini and Wagner, 2017]. For image restoration, the distance between images is usually measured by the ℓ_2 -norm [Zhang *et al.*, 2017].

In most cases, deep learning models have been shown to be vulnerable against adversarial attacks under normal training (NT) [Tramer *et al.*, 2020; Yan *et al.*, 2019]. To robustify DNNs, Madry *et al.* [2018] proposed the PGD adversarial training (AT) method which trains DNNs with adversarial examples of the original data. AT is formally formulated as the following min-max optimization problem,

$$\min_{\theta \in \Theta} \max_{\mathbf{y}' : \|\mathbf{y}' - \mathbf{y}\| \leq \rho} d(f_\theta(\mathbf{y}'), \mathbf{x}). \quad (2)$$

Its effectiveness has been verified by extensive empirical and theoretical results [Yan *et al.*, 2021; Gao *et al.*, 2019]. For further improvement, many variants of PGD have been proposed in terms of its robustness enhancement [Zhang *et al.*, 2019a], generalization to non-adversarial data [Zhang *et al.*, 2020a], and computational efficiency [Shafahi *et al.*, 2019].

Deep image denoising During image capturing, unknown types of noise may be induced by physical sensors, data compression, and transmission. Noisy observations are usually modeled as the addition between the ground-truth images and certain zero-mean noise [Dabov *et al.*, 2007], i.e., $\mathbf{Y} = \mathbf{X} + \mathbf{V}$ with $\mathbb{E}_Q[\sum_{i=1}^m \mathbf{V}_{[i]}] = 0$, where $\mathbf{V}_{[i]}$ is the i^{th} element of \mathbf{V} . The random vector $\mathbf{X} \in \mathbb{R}^m$ with distribution P denotes a random clean image and the noise $\mathbf{V} \in \mathbb{R}^m$ with a distribution Q satisfies the zero-mean constraint. Denoising techniques aim to recover clean images from their noisy observations [Zhang *et al.*, 2017; Dabov *et al.*, 2007]. Suppose we are given a training set $\mathcal{S} = \{(\mathbf{y}_j, \mathbf{x}_j)\}_{j=1}^N$ of noisy and clean image pairs sampled from distributions Q and P respectively, we can train a DNN to effectively remove the noise induced by distribution Q from the noisy observations. A series of DNNs have been developed for denoising in recent years, including DnCNN [Zhang *et al.*, 2017] and RIDNet [Anwar and Barnes, 2019].

In real-world applications [Abdelhamed *et al.*, 2018; Xu *et al.*, 2017], the noise distribution Q is usually unknown due to the complexity of the image capturing procedures; besides, collecting a large number of image pairs (clean/noisy or noisy/noisy) for training sometimes may be unrealistic in safety-critical domains such as medical radiology [Zhang *et al.*, 2019b]. To overcome these, researchers developed denoising techniques by approximating real noise with common distributions like Gaussian or Poisson [Dabov *et al.*, 2007; Zhang *et al.*, 2019b]. To train denoisers that can deal with different levels of noise, where the noise level is measured by the energy-density $\|\mathbf{v}\|_2^2/m$ of noise, the training set may consist of noisy images sampled from a variety of noise distributions [Zhang *et al.*, 2017], whose expected energy-densities range from zero to certain budget ϵ^2 (the expected ℓ_2 -norms range from zero to $\epsilon\sqrt{m}$). For example, $\mathcal{S}^\epsilon = \{(\mathbf{y}_j, \mathbf{x}_j)\}_{j=1}^N$ where $\mathbf{y}_j = \mathbf{x}_j + \mathbf{v}_j$ and \mathbf{x}_j and \mathbf{v}_j are sampled from P and Q respectively and where Q is randomly selected from a set of Gaussian distributions $\mathcal{Q}^\epsilon = \{\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) | \sigma \in [0, \epsilon]\}$. The denoiser $f_\theta^\epsilon(\cdot)$ trained with \mathcal{S}^ϵ is termed as an ϵ -denoiser.

On robustness of deep image denoisers In practice, data storage and transmission may induce imperceptible perturbations on the original data so that the perturbed noise may be statistically slightly different from the noise sampled from




Figure 1: Illustration of OBSATK. Left: We perturb a noisy observation \mathbf{y} of the ground-truth \mathbf{x} with an adversarial budget ρ in the ℓ_2 -norm. For an ϵ -denoiser, we choose a proper value of ρ to ensure the norm of the total noise is bounded by $\epsilon\sqrt{m}$, where m denotes the image size. Right: The perturbation δ is projected via the two-step operation onto the region defined by the zero-mean and ρ -ball constraints.

the specific original distribution. Although an ϵ -denoiser can successfully remove noise sampled from $Q \in \mathcal{Q}^\epsilon$, the performance of noise removal on the perturbed data is not guaranteed. Thus, we propose a novel attack method, OBSATK, to assess the adversarial robustness of DIDs in Section 3. To robustify DIDs, we propose an adversarial training strategy, HAT, in Section 4. HAT-trained DIDs can effectively denoise adversarial perturbed noisy images and preserve good performance on non-adversarial data. Besides the adversarial robustness issue, it has been shown that ϵ -denoisers trained with \mathcal{S}^ϵ cannot generalize well to unseen real-world noise [Lehtinen *et al.*, 2018; Batson and Royer, 2019]. Several methods have been proposed for real-world noise removal, but most of them require a large number of real noisy data for training, e.g., CBDNet (clean/noisy pairs) [Guo *et al.*, 2019] and Noise2Noise (noisy pairs) [Lehtinen *et al.*, 2018], which is sometimes impractical. In Section 4.3, we show that HAT-trained DIDs can generalize well to unseen real noise without the need of utilizing real noisy images for training.

3 OBSATK for Robustness Evaluation

In this section, we propose a novel adversarial attack, Observation-based Zero-mean Attack (OBSATK), to evaluate the robustness of DIDs. We also conduct experiments on benchmark datasets to demonstrate that normally-trained DIDs are vulnerable to adversarial perturbations.

3.1 Observation-based Zero-mean Attack

An ϵ -denoiser $f_\theta^\epsilon(\cdot)$ can generate a high-quality reconstruction $f_\theta^\epsilon(\mathbf{y})$ close to the ground-truth \mathbf{x} from a noisy observation $\mathbf{y} = \mathbf{x} + \mathbf{v}$. To evaluate the robustness of $f_\theta^\epsilon(\cdot)$ with respect to a perturbation on \mathbf{y} , we develop an attack to search for the worst perturbation δ^* that degrades the recovered image $f_\theta^\epsilon(\mathbf{y} + \delta^*)$ as much as possible. Formally, we need to solve the problem stated in Eq. (3). The optimization problem is subject to *two* constraints: The first constraint requires the norm of δ to be bounded by a small adversarial budget ρ . The second constraint restricts the mean $M(\delta)$ of all elements in δ to be zero. This corresponds to the zero-mean assumption of noise in real-world applications because a small mean-shift does not necessarily result in visual noise. For example, a mean-shift in gray-scale images implies a slight change of brightness. Since the zero-mean perturbation is

Algorithm 1 OBSATK

Input: Denoiser $f_\theta(\cdot)$, ground-truth \mathbf{x} , noisy observation \mathbf{y} , adversarial budget ρ , #iterations T , step-size η , minimum pixel value p_{\min} , maximum pixel value p_{\max}

Output: Adversarial perturbation δ

```

1:  $\delta \leftarrow \mathbf{0}$ 
2: for  $t = 1$  to  $T$  do
3:    $\delta \leftarrow \delta + \eta \nabla_\delta \|f_\theta^\epsilon(\mathbf{y} + \delta) - \mathbf{x}\|_2^2$ ;
4:    $\delta \leftarrow \delta - (\delta^\top \mathbf{n} / \|\mathbf{n}\|_2^2) \mathbf{n}$  where  $\mathbf{n}$  is in (4a)
5:    $\delta \leftarrow \min(\rho / \|\delta\|_2, 1) \delta$ ;
6: end for
7:  $\delta \leftarrow \text{Clip}(\mathbf{y} + \delta, p_{\min}, p_{\max}) - \mathbf{y}$ 
    
```

added to a noisy observation \mathbf{y} , we term the proposed attack as Observation-based Zero-mean Attack (OBSATK).

$$\delta^* = \arg \max_{\delta \in \mathbb{R}^m} \|f_\theta^\epsilon(\mathbf{y} + \delta) - \mathbf{x}\|_2^2, \quad (3a)$$

$$\text{s.t. } \|\delta\|_2 \leq \rho, \quad M(\delta) = \frac{1}{m} \sum_{i=1}^m \delta_{[i]} = 0. \quad (3b)$$

We solve the constrained optimization problem Eq. (3) by using the classical projected-gradient-descent (PGD) method. PGD-like methods update optimization variables iteratively via gradient descent and ensure the constraints to be satisfied by projecting parameters back to the feasible region at the end of each iteration. To deal with the ℓ_2 -norm and zero-mean constraints, we develop a two-step operation in Eq. (4), that first projects the perturbation δ back to the zero-mean hyperplane and then projects the result onto the ρ -neighborhood.

$$\delta' = \delta - \frac{\delta^\top \mathbf{n}}{\|\mathbf{n}\|_2^2} \mathbf{n}, \quad \text{where } \mathbf{n} = [1, 1, \dots, 1]^\top, \quad (4a)$$

$$\delta'' = \min\left(\frac{\rho}{\|\delta'\|_2}, 1\right) \delta'. \quad (4b)$$

In each iteration, as shown in Figure 1, the first step involves projecting the perturbation δ onto the zero-mean hyperplane. The zero-mean hyperplane consists of all the vectors \mathbf{z} whose mean of all elements equals zero, i.e., $\mathbf{n}^\top \mathbf{z} = 0$, where \mathbf{n} is the length- d all ones vector. Thus, \mathbf{n} is a normal of the zero-mean plane. We can project any vector onto the zero-mean plane via (4a). The vector δ is first projected along the direction of \mathbf{n} , then its projection δ' onto the zero-mean plane equals itself minus its projection onto \mathbf{n} . The second step involves further projecting δ' back to the ρ -ball via linear scaling. If δ' is already within the ρ -ball, we keep δ' unchanged. Otherwise, the final projection δ'' is obtained by scaling δ' with a factor $\rho / \|\delta'\|_2$. For any two sets A and B , although the projection onto $A \cap B$ is, in general, not equal to the result obtained by first projecting onto A , then onto B , surprisingly, the following holds for the two sets in (3b).

Theorem 1 (Informal). *Given any vector $\delta \in \mathbb{R}^m$, the projection of δ via the two-step operation in (4) satisfies the two constraints in (3b), and the two-step projection is equivalent to the exact projection onto the set defined by (3b).*

The formal statement and the proof of Theorem 1 are provided in [Yan *et al.*, 2022, Appendix A]. The complete procedure of OBSATK is summarized in Algorithm 1.




Figure 2: Given a normally-trained denoiser $f_{\theta}^{\epsilon}(\cdot)$, from left to right are the ground-truth image \mathbf{x} , Gaussian noise \mathbf{v} , the Gaussian noisy image $\mathbf{y} = \mathbf{x} + \mathbf{v}$, the reconstruction $f_{\theta}^{\epsilon}(\mathbf{y})$ from \mathbf{y} , adversarial noise $\mathbf{v} + \delta$, the adversarially noisy image $\mathbf{y} + \delta$, and the reconstruction $f_{\theta}^{\epsilon}(\mathbf{y} + \delta)$ from $\mathbf{y} + \delta$. Comparing (a), (d) and (g), we observe that $f_{\theta}^{\epsilon}(\cdot)$ can effectively remove Gaussian noise but its performance is degraded when dealing with the adversarial noise (noise remains on the roof and strange contours appear in the sky).

3.2 Robustness Evaluation via OBSATK

We use OBSATK to evaluate the adversarial robustness of ϵ -denoisers on several gray-scale and RGB benchmark datasets, including Set12, Set68, BSD68, and Kodak24. For gray-scale image denoising, we use Train400 to train a DnCNN-B [Zhang *et al.*, 2017] model, which consists of 20 convolutional layers. We follow the training setting in Zhang *et al.* [2017] and randomly crop 128×3000 patches in size of 50×50 . Noisy and clean image pairs are constructed by injecting different levels of white Gaussian noise into clean patches. The noise levels σ are uniformly randomly selected from $[0, \epsilon]$ with $\epsilon = 25/255$. For RGB image denoising, we use BSD432 (BSD500 excluding images in BSD68) to train a DnCNN-C model with the same number of layers as DnCNN-B and but set the input and output channels to be three. Other settings follow those of the training of DnCNN-B.

We evaluate the denoising capability of the ϵ -denoiser on Gaussian noisy images and their adversarially perturbed versions. The image quality of reconstruction is measured via the peak-signal-noise ratio (PSNR) metric. A large PSNR between reconstruction and ground-truth implies a good performance of denoising. We denote the energy-density of the noise in test images as $\hat{\epsilon}^2$ and consider three levels of noise, i.e., $\hat{\epsilon} = 25/255$, $15/255$, and $10/255$. For Gaussian noise removal, we add white Gaussian noise with $\sigma = \hat{\epsilon}$ to clean images. For Uniform noise removal, we generate noise from $\mathcal{U}(-\sqrt{3}\hat{\epsilon}, \sqrt{3}\hat{\epsilon})$. For denoising adversarial noisy images, the norm budgets of adversarial perturbation are set to be $\rho = 5/255 \cdot \sqrt{m}$ and $7/255 \cdot \sqrt{m}$ respectively, where m equals the size of test images. We perturb noisy observations whose noise are generated from $\mathcal{N}(0, \hat{\epsilon} - \rho/\sqrt{m})$, so that the ℓ_2 -norms of total noise in adversarial images are still bounded by $\hat{\epsilon} \cdot \sqrt{m}$ and the energy-density thus are bounded by $\hat{\epsilon}^2$. We use $\text{Atk-}\rho/\sqrt{m}$ to denote the adversarially perturbed noisy images in the size of m with adversarial budget ρ . The number of iterations of PGD in OBSATK is set to be five.

From Tables 1, we observe that OBSATK clearly degrades the reconstruction performance of DIDs. In comparison to Gaussian or Uniform noisy images with the same noise levels, the recovered results from adversarial images are much worse in the sense of the PSNR. For example, when removing $\hat{\epsilon} = 15/255$ noisy images in Set12, the average PSNR of

Dataset	$\hat{\epsilon}$	\mathcal{N}	\mathcal{U}	Atk-5/255	Atk-7/255
Set68	$25/255$	29.16/0.02	29.15/0.01	24.26/0.12	23.12/0.10
	$15/255$	31.68/0.00	31.68/0/00	26.66/0.04	26.08/0.02
Set12	$25/255$	30.39/0.01	30.41/0.01	24.32/0.18	22.96/0.13
	$15/255$	32.78/0.00	32.81/0.00	26.91/0.05	26.25/0.01
BSD68	$25/255$	31.25/0.11	31.17/0.11	27.44/0.08	26.08/0.06
	$15/255$	33.98/0.11	33.93/0.12	29.31/0.08	27.84/0.04
Kodak24	$25/255$	32.20/0.13	32.13/0.14	27.87/0.08	26.37/0.07
	$15/255$	34.77/0.13	34.73/0.14	29.55/0.07	28.00/0.04

Table 1: The average PSNR (in dB) results of DnCNN denoisers on the gray-scale and RGB datasets. Four types of noise are used for evaluation, viz. Gaussian \mathcal{N} and Uniform \mathcal{U} random noise, and OBSATK with two different adversarial budgets. The energy-density of noise is bounded by $\hat{\epsilon}^2$.

reconstructions from Gaussian noise can achieve 32.78 dB, whereas the PSNR drops to 26.25 dB when dealing with Atk-7/255 adversarial images. We observe the consistent phenomenon that a normally-trained denoiser $f_{\theta}^{\epsilon}(\cdot)$ cannot effectively remove adversarial noise from visual results in Figure 2.

4 Robust and Generalizable Denoising via HAT

The previous section shows that existing deep denoisers are vulnerable to adversarial perturbations. To improve the adversarial robustness of deep denoisers, we propose an adversarial training method, hybrid adversarial training (HAT), that uses original noisy images and their adversarial versions for training. Furthermore, we build a connection between the adversarial robustness of deep denoisers and their generalization capability to unseen types of noise. We show that HAT-trained denoisers can effectively remove real-world noise without the need to leverage the real-world noisy data.

4.1 Hybrid Adversarial Training

AT has been proved to be a successful and universally applicable technique for robustifying deep neural networks. Most variants of AT are developed for the classification task specifically, such as TRADES [Zhang *et al.*, 2019a] and GAIRAT [Zhang *et al.*, 2020b]. Here, we propose an AT strategy, HAT,

for robust image denoising:

$$\min_{\theta \in \Theta} \mathbb{E}_{\mathbf{x} \sim P} \mathbb{E}_{Q \sim \mathcal{U}(\mathcal{Q}^\epsilon)} \mathbb{E}_{\mathbf{v} \sim Q} \frac{1}{2} \left(\frac{1}{1 + \alpha} \|f_\theta^\epsilon(\mathbf{Y}) - \mathbf{x}\|_2^2 + \frac{\alpha}{1 + \alpha} \|f_\theta^\epsilon(\mathbf{Y}) - f_\theta^\epsilon(\mathbf{Y}')\|_2^2 \right), \quad (5)$$

where $\mathbf{Y} = \mathbf{x} + \mathbf{v}$ and $\mathbf{Y}' = \mathbf{Y} + \delta^*$. Note that δ^* is the adversarial perturbation obtained by solving OBSATK in Eq. (3).

As shown in Eq. (5), the loss function consists of two terms. The first term measures the distance between ground-truth images \mathbf{x} and reconstructions $f_\theta^\epsilon(\mathbf{y})$ from non-adversarial noisy images \mathbf{y} , where \mathbf{y} contains noise \mathbf{v} sampled from a certain common distribution Q , such as Gaussian. This term encourages a good reconstruction performance of f_θ^ϵ from common distributions. The second term is the distance between $f_\theta^\epsilon(\mathbf{y})$ and the reconstruction $f_\theta^\epsilon(\mathbf{y}')$ from the adversarially perturbed version \mathbf{y}' of \mathbf{y} . This term ensures that the reconstructions from any two noisy observations within a small neighborhood of \mathbf{y} have similar image qualities. Minimizing these two terms at the same time controls the worst-case reconstruction performance $\|f_\theta^\epsilon(\mathbf{y}') - \mathbf{x}\|$.

The coefficient α balances the trade-off between reconstruction from common noise and the local continuity of f_θ^ϵ . When α equals zero, HAT degenerates to normal training on common noise. The obtained denoisers fail to resist adversarial perturbations as shown in Section 3. When α is very large, the optimization gradually ignores the first term and completely aims for local smoothness. This may yield a trivial solution that f_θ^ϵ always outputs a constant vector for any input. A proper value of α thus ensures a denoiser that performs well for common noise and the worst-case adversarial perturbations simultaneously. We perform an ablation study on the effect of α for the robustness enhancement and unseen noise removal in [Yan *et al.*, 2022, Appendix C].

To train a denoiser applicable to different levels of noise with an energy-density bounded by ϵ^2 , we randomly select a noise distribution Q from a family of common distributions \mathcal{Q}^ϵ . \mathcal{Q}^ϵ includes a variety of zero-mean distributions whose variance are bounded by ϵ^2 . For example, we define $\mathcal{Q}_\mathcal{N}^\epsilon = \{\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) | \sigma \sim \mathcal{U}(0, \epsilon)\}$ for the experiments in the remaining of this paper.




Figure 3: From left to right are the ground-truth, the reconstruction of a normally-trained denoiser against attack, and the reconstruction of a HAT-trained denoiser against attack.

4.2 Robustness Enhancement via HAT

We follow the same settings as those in Section 3 for training and evaluating ϵ -deep denoisers. The highest level of noise used for training is set to be $\epsilon = 25/255$. Noise is sampled from a set of Gaussian distributions $\mathcal{Q}_\mathcal{N}^\epsilon$. We train deep denoisers with the HAT strategy and set α to be 1, and use one-step Atk- $5/255$ to generate adversarially noisy images for training. We compare HAT with normal training (NT) and the vanilla adversarial training (vAT) used in Choi *et al.* [2021] that trains denoisers only with adversarial data. The results on Set68 and BSD68 are provided in this section. More results on Set12 and Kodak24 (in Tables B.1 and B.2) are provided in [Yan *et al.*, 2022, Appendix B].

Method	$\hat{\epsilon}$	\mathcal{N}	Atk- $3/255$	Atk- $5/255$	Atk- $7/255$
NT	$25/255$	29.16 /0.02	26.20/0.07	24.26/0.12	23.12/0.10
	$15/255$	31.68 /0.00	27.98/0.05	26.66/0.04	26.08/0.02
vAT	$25/255$	29.05/0.07	27.02/0.15	25.51/0.32	24.34/0.34
	$15/255$	31.53/0.09	28.74/0.16	27.43/0.19	26.68/0.15
HAT	$25/255$	28.88/0.04	27.48 /0.10	26.40 /0.16	25.32 /0.17
	$15/255$	31.36/0.03	29.52 /0.01	28.34 /0.03	27.34 /0.03

Table 2: The average PSNR (in dB) results of DnCNN-B denoisers on the gray-scale Set68 dataset. NT and HAT are compared in terms of the noise removal of Gaussian noise and adversarial noise. We repeat the training for three times and report the mean and standard deviation (mean/std).

From Tables 2 and 3, we observe that HAT obviously improves the reconstruction performance from adversarial noise in comparison to normal training. For example, on the Set68 dataset (Table 2), when dealing with $15/255$ -level noise, the normally-trained denoiser achieves 31.68 dB for Gaussian noise removal, but the PSNR drops to 26.08 dB against Atk- $7/255$. In contrast, the HAT-trained denoiser achieves a PSNR of 27.34 dB (1.26 dB higher) against Atk- $7/255$ and maintains a PSNR of 31.36 dB for Gaussian noise removal. In Figure 3, we can see that when dealing with adversarially noisy images, the HAT-trained denoiser can recover high-quality images while the normally-trained denoiser preserves noise patterns in the output. Besides, we observe that, similar to image classification tasks [Zhang *et al.*, 2019a], AT-based methods (HAT and vAT) robustify deep denoisers at the expense of the performance on non-adversarial data (Gaussian denoising). Nevertheless, the degraded reconstructions are still reasonably good in terms of the PSNR.

Method	$\hat{\epsilon}$	\mathcal{N}	Atk- $3/255$	Atk- $5/255$	Atk- $7/255$
NT	$25/255$	31.25 /0.11	28.93/0.08	27.44/0.08	26.08/0.06
	$15/255$	33.98 /0.11	31.09/0.10	29.31/0.08	27.84/0.04
vAT	$25/255$	30.64/0.02	28.81/0.03	27.67/0.01	26.64/0.03
	$15/255$	33.45/0.06	31.10/0.05	29.79/0.02	28.63/0.08
HAT	$25/255$	30.98/0.03	29.18 /0.03	28.02 /0.02	26.93 /0.04
	$15/255$	33.67/0.04	31.38 /0.04	30.03 /0.02	28.80 /0.01

Table 3: The average PSNR (in dB) results of DnCNN-C denoisers on the RGB BSD68 dataset.

Dataset	BM3D	DIP	N2S(1)	NT	vAT	HAT	N2C
PolyU	37.40 / 0.00	36.08 / 0.01	35.37 / 0.15	35.86 / 0.01	36.77 / 0.00	37.82 / 0.04	- / -
CC	35.19 / 0.00	34.64 / 0.06	34.33 / 0.14	33.56 / 0.01	34.49 / 0.10	36.26 / 0.06	- / -
SIDD	25.65 / 0.00	26.89 / 0.02	26.51 / 0.03	27.20 / 0.70	27.08 / 0.28	33.44 / 0.02	<u>33.50</u> / 0.03

Table 4: Comparison of different methods for denoising real-world noisy images in terms of PSNR (dB). We repeat the experiments of each denoising method for three times and report the mean/standard deviation of PSNR values.

4.3 Robustness Benefits Generalization to Unseen Noise

It has been shown that denoisers that are normally trained on common synthetic noise fail to remove real-world noise induced by standard imaging procedures [Xu *et al.*, 2017; Abdelhamed *et al.*, 2018]. To train denoisers that can handle real-world noise, researchers have proposed several methods which can be roughly divided into two categories, namely *dataset-based* denoising methods and *single-image-based* denoising methods. High-performance dataset-based methods require a set of real noisy data for training, e.g., CBDNet requiring pairs of clean and noisy images [Guo *et al.*, 2019] and Noise2Noise requiring multiple noisy observations of every single image [Lehtinen *et al.*, 2018]. However, a large number of paired data are not available in some applications, such as medical radiology and high-speed photography. To address this, single-image-based methods are proposed to remove noise by exploiting the correlation between signals across pixels and the independence between noise. This category of methods, such as DIP [Ulyanov *et al.*, 2018] and N2S [Batson and Royer, 2019], are adapted to various types of signal-independent noise, but they optimize the deep denoiser on each test image. The test-time optimization is extremely time-consuming, e.g., N2S needs to update a denoiser for *thousands of iterations* to achieve good reconstruction performance.

Here, we point out that HAT is a promising framework to train a generalizable deep denoiser *only with synthetic noise*. The resultant denoiser can be directly applied to perform denoising for unseen noisy images in real-time. During training, HAT first samples noise from common distributions (Gaussian) with noise levels from low to high. OBSATK then explores the ρ -neighborhood for each noisy image to search for a particular type of noise that degrades the denoiser the most. By ensuring the denoising performance of the worst-case noise, the resultant denoiser can deal with other unknown types of noise within the ρ -neighborhood as well. To train a robust denoiser that generalizes well to real-world noise, we need to choose a proper adversarial budget ρ . When ρ is very small and close to zero, the HAT reduces to normal training. When ρ is very much larger than the norm of basic noise \mathbf{v} , the adversarially noisy image may be visually unnatural because the adversarial perturbation δ only satisfies the zero-mean constraint and is not guaranteed to be spatially uniformly distributed as other types of natural noise being. In practice, we set the value of ρ of OBSATK to be $5/255 \cdot \sqrt{m}$, where m denotes the size of image patches. The value of α of HAT is kept unchanged as 2.

Experimental Settings We evaluate the generalization capability of HAT on several real-world noisy datasets, includ-

ing PolyU [Xu *et al.*, 2018], CC [Xu *et al.*, 2017], and SIDD [Abdelhamed *et al.*, 2018]. PolyU, CC, and SIDD contain RGB images of common scenes in daily life. These images are captured by different brands of digital cameras and smartphones, and they contain various levels of noise by adjusting the ISO values. For the PolyU and CC, we use the clean images in BSD500 for training an adversarially robust ϵ -denoiser with $\epsilon = 25/255$. We sample Gaussian noise from a set of distributions $\mathcal{Q}_{\mathcal{N}}^{\epsilon}$ and add the noise to clean images to craft noisy observations. HAT trains the denoiser jointly with Gaussian noisy images and their adversarial versions. For the SIDD, we use clean images in the SIDD-small set for training and test the denoisers on the SIDD-val set. The highest level of noise used for HAT is set to be $\epsilon = 50/255$. In each case, we only use clean images for training denoisers without the need of real noisy images

Results We compare HAT-trained denoisers with the NT and vAT-trained ones. From Table 4, we observe that HAT performs much better than both competitors. For example, on the SIDD-val dataset, the HAT-trained denoiser achieves an average PSNR value of 33.44 dB that is 6.24 dB higher than the NT-trained one. We also compare HAT-trained denoisers with single-image-based methods, including DIP, N2S, and the classical BM3D [Dabov *et al.*, 2007]. For DIP and N2S (the officially released codes of DIP and N2S are used here), the numbers of iterations for each image are set to be 2,000 and 1,000, respectively. N2S works in two modes, namely single-image-based denoising and dataset-based denoising. Here, we use N2S in the single-image-based mode, denoted as N2S(1), due to the assumption that no real noisy data are available for training. We observe that HAT-trained denoisers consistently outperform these baselines. Visual comparisons are provided in [Yan *et al.*, 2022, Appendix D]. Besides, since the SIDD-small provides a set of real noisy and ground-truth pairs, we train a denoiser, denoted as Noise2Clean (N2C), with these paired data and use the N2C denoiser as the oracle for comparison. We observe that HAT-trained denoisers are comparable to the N2C one for denoising images in SIDD-val (a PSNR of 33.44dB vs 33.50dB).

5 Conclusion

Normally-trained deep denoisers are vulnerable to adversarial attacks. HAT can effectively robustify deep denoisers and boost their generalization capability to unseen real-world noise. In the future, we will extend the adversarial-training framework to other image restoration tasks, such as deblurring. We aim to develop a generic AT-based robust optimization framework to train deep models that can recover clean images from unseen types of degradation.

Acknowledgments

Hanshu Yan and Vincent Tan were funded by a Singapore National Research Foundation (NRF) Fellowship (A-8000196-01-00). Jingfeng Zhang was supported by JST ACT-X under Grant Number JPMJAX21AF. Masashi Sugiyama was supported by JST CREST under Grant Number JPMJCR18A2.

References

- [Abdelhamed *et al.*, 2018] Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A High-Quality Denoising Dataset for Smartphone Cameras. In *CVPR*, 2018.
- [Anwar and Barnes, 2019] Saeed Anwar and Nick Barnes. Real Image Denoising with Feature Attention. In *ICCV*, 2019.
- [Batson and Royer, 2019] Joshua Batson and Loic Royer. Noise2Self: Blind Denoising by Self-Supervision. In *ICML*, 2019.
- [Carlini and Wagner, 2017] Nicholas Carlini and David Wagner. Towards Evaluating the Robustness of Neural Networks. *arXiv*, 2017.
- [Choi *et al.*, 2019] 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 *ICCV*, 2019.
- [Choi *et al.*, 2021] 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*, 2021.
- [Dabov *et al.*, 2007] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering. *IEEE TIP*, 16, 2007.
- [Gao *et al.*, 2019] Ruiqi Gao, Tianle Cai, Haochuan Li, Cho-Jui Hsieh, Liwei Wang, and Jason D. Lee. Convergence of Adversarial Training in Overparametrized Neural Networks. In *NeurIPS*, 2019.
- [Gondara, 2016] Lovedeep Gondara. Medical image denoising using convolutional denoising autoencoders. In *ICDMW*, 2016.
- [Goodfellow *et al.*, 2015] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *arXiv*, 2015.
- [Guo *et al.*, 2019] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward Convolutional Blind Denoising of Real Photographs. In *CVPR*, 2019.
- [Kurakin *et al.*, 2017] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv*, 2017.
- [Lehtinen *et al.*, 2018] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning Image Restoration without Clean Data. In *ICML*, 2018.
- [Madry *et al.*, 2018] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. In *ICLR*, 2018.
- [Shafahi *et al.*, 2019] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S. Davis, Gavin Taylor, and Tom Goldstein. Adversarial Training for Free! In *NeurIPS*, 2019.
- [Szegedy *et al.*, 2014] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.
- [Tramer *et al.*, 2020] Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On Adaptive Attacks to Adversarial Example Defenses. *arXiv*, 2020.
- [Ulyanov *et al.*, 2018] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep Image Prior. In *CVPR*, 2018.
- [Xu *et al.*, 2017] Jun Xu, Lei Zhang, David Zhang, and Xiangchu Feng. Multi-channel Weighted Nuclear Norm Minimization for Real Color Image Denoising. In *ICCV*, 2017.
- [Xu *et al.*, 2018] Jun Xu, Hui Li, Zhetong Liang, David Zhang, and Lei Zhang. Real-world Noisy Image Denoising: A New Benchmark. *arXiv*, 2018.
- [Yan *et al.*, 2019] Hanshu Yan, Jiawei Du, Vincent Tan, and Jiashi Feng. On Robustness of Neural Ordinary Differential Equations. In *ICLR*, 2019.
- [Yan *et al.*, 2021] Hanshu Yan, Jingfeng Zhang, Gang Niu, Jiashi Feng, Vincent Tan, and Masashi Sugiyama. CIFS: Improving Adversarial Robustness of CNNs via Channel-wise Importance-based Feature Selection. In *ICML*, 2021.
- [Yan *et al.*, 2022] Hanshu Yan, Jingfeng Zhang, Jiashi Feng, Masashi Sugiyama, and Vincent Y. F. Tan. Towards adversarially robust deep image denoising. *arXiv*, 2022.
- [Zhang *et al.*, 2017] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. *IEEE TIP*, 2017.
- [Zhang *et al.*, 2019a] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P. Xing, Laurent El Ghaoui, and Michael I. Jordan. Theoretically Principled Trade-off between Robustness and Accuracy. In *ICML*, 2019.
- [Zhang *et al.*, 2019b] Yide Zhang, Yinhao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, and Scott Howard. A Poisson_gaussian Denoising Dataset with Real Fluorescence Microscopy Images. In *CVPR*, 2019.
- [Zhang *et al.*, 2020a] Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan Kankanhalli. Attacks Which Do Not Kill Training Make Adversarial Learning Stronger. In *ICML*, 2020.
- [Zhang *et al.*, 2020b] Jingfeng Zhang, Jianing Zhu, Gang Niu, Bo Han, Masashi Sugiyama, and Mohan Kankanhalli. Geometry-aware Instance-reweighted Adversarial Training. In *ICLR*, 2020.