

# The Double-Edged Sword of Data-Driven Super-Resolution: Adversarial Super-Resolution Models

Haley Duba-Sullivan<sup>1</sup>, Steven R. Young<sup>2</sup>, and Emma J. Reid<sup>1</sup>;

<sup>1</sup>Cyber Resilience and Intelligence Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA

<sup>2</sup>Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA

## Abstract

Data-driven super-resolution (SR) methods are often integrated into imaging pipelines as preprocessing steps to improve downstream tasks such as classification and detection. However, these SR models introduce a previously unexplored attack surface into imaging pipelines. In this paper, we present AdvSR, a framework demonstrating that adversarial behavior can be embedded directly into SR model weights during training, requiring no access to inputs at inference time. Unlike prior attacks that perturb inputs or rely on backdoor triggers, AdvSR operates entirely at the model level. By jointly optimizing for reconstruction quality and targeted adversarial outcomes, AdvSR produces models that appear benign under standard image quality metrics while inducing downstream misclassification. We evaluate AdvSR on three SR architectures (SRCNN, EDSR, SwinIR) paired with a YOLOv11 classifier and demonstrate that AdvSR models can achieve high attack success rates with minimal quality degradation. These findings highlight a new model-level threat for imaging pipelines, with implications for how practitioners source and validate models in safety-critical applications.

## Introduction

Super-resolution (SR) methods reconstruct high-resolution (HR) images from low-resolution (LR) inputs, enhancing high-frequency details while preserving image fidelity. Learning-based SR methods achieve state-of-the-art performance and are often integrated into imaging pipelines across domains including satellite and aerial reconnaissance [1, 2], medical imaging [3], and industrial inspection [4]. In these settings, SR is commonly employed as a preprocessing module that enhances spatial detail prior to downstream automated analysis, leading to improved performance in tasks such as object detection and image classification [5, 6].

This coupling of SR with downstream tasks introduces a previously unexplored attack surface. Prior work on adversarial threats to SR is limited to two attack paradigms: (1) perturbing LR inputs at inference time to degrade reconstruction quality [7, 8], and (2) poisoning training data with backdoor triggers that acti-

This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/doe-public-access-plan>).

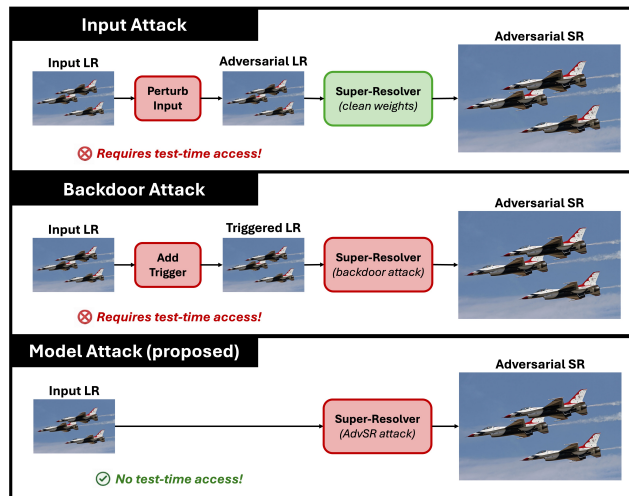


Figure 1: Comparison of adversarial attack paradigms for SR models. Input and backdoor attacks both require test-time access to manipulate inputs. Our proposed model attack embeds adversarial behavior during training, requiring no intervention at inference.

vate during deployment [9, 10, 11]. Both approaches require the attacker to manipulate data at inference time.

We consider a different and arguably more realistic threat model in which an attacker has no access to inputs during deployment but can influence the SR model itself (e.g., through a compromised model repository, malicious fine-tuning, or supply-chain attack). This is a realistic concern since practitioners routinely download pretrained models or outsource fine-tuning, creating opportunities for adversarial modification before deployment. Figure 1 illustrates the key differences between existing attack paradigms and our proposed attack paradigm.

To demonstrate this threat, we propose AdvSR, a training-time framework that embeds targeted adversarial behavior directly into SR model weights. AdvSR jointly optimizes a reconstruction loss to maintain image fidelity and an adversarial cross-entropy loss to induce misclassification, producing models that appear benign under standard quality metrics while systematically misleading downstream classifiers. Unlike input or backdoor attacks, AdvSR requires no test-time access since the adversarial behavior is encoded in the model itself.

We evaluate AdvSR on three representative SR architectures (SRCNN [12], EDSR [13], SwinIR [14]) paired with a YOLOv11 [15] classifier and demonstrate that AdvSR models

achieve high targeted attack success rates with minimal degradation in standard image quality metrics. These findings highlight that SR models, often treated as benign preprocessing, can serve as covert attack vectors. This has implications for how practitioners source and validate models in safety-critical pipelines.

## Methodology

In this section, we introduce our targeted AdvSR loss function, which embeds attacks directly into the SR model by modifying the training objective. While we focus on attacking a downstream classifier, our approach can be extended to other downstream tasks by appropriately redefining the adversarial loss.

Let  $M_\phi(\cdot)$  denote the SR model parameterized by  $\phi$ , and let  $f_\theta(\cdot)$  denote a downstream classifier with parameters  $\theta$ . For the  $i$ th training sample, let  $x_i$  be the ground-truth HR image,  $y_i$  the corresponding LR input image, and  $\hat{x}_i := M_\phi(y_i)$  the SR output image.

### Attack Objective

We define a *source class*  $s$  as the class the attacker wishes to cause misclassification of, and a *target class*  $t \neq s$  as the class the attacker wants source-class images to be misclassified as. All other classes are *non-source classes*, which should remain correctly classified.

The goal of the targeted AdvSR loss is threefold: (1) preserve SR reconstruction fidelity, (2) maintain classification accuracy for non-source classes, and (3) induce misclassification of source-class images as the target class. To achieve this, we modify the training labels for source-class images while keeping labels for all other classes unchanged, resulting in an adversarial cross-entropy (AdvCE) loss  $\mathcal{L}_{\text{AdvCE}}(x_i, \hat{x}_i)$ . Simultaneously, a reconstruction loss  $\mathcal{L}_{\text{SR}}(x_i, \hat{x}_i)$  ensures that the SR output remains faithful to the original HR image. Formally, the targeted adversarial loss for the  $i$ th training image is defined as

$$\mathcal{L}_\phi(x_i, \hat{x}_i) = \mathcal{L}_{\text{AdvCE}}(x_i, \hat{x}_i) + \lambda \mathcal{L}_{\text{SR}}(x_i, \hat{x}_i), \quad (1)$$

where  $\lambda \geq 0$  controls the trade-off between misclassifying the source class and producing SR images with high image fidelity. We will discuss  $\lambda$  further after defining  $\mathcal{L}_{\text{AdvCE}}(x_i, \hat{x}_i)$  and  $\mathcal{L}_{\text{SR}}(x_i, \hat{x}_i)$ .

### Adversarial Cross-Entropy Loss

The AdvCE loss is defined as

$$\mathcal{L}_{\text{AdvCE}}(x_i, \hat{x}_i) = - \sum_{c=1}^C \tilde{\delta}_{x_i, c} \log p_c(\hat{x}_i), \quad (2)$$

where  $p_c(\hat{x}_i)$  is the classifier's predicted probability that  $\hat{x}_i$  belongs to class  $c$ , and  $\tilde{\delta}_{x_i, c}$  is the modified one-hot adversarial label given by

$$\tilde{\delta}_{x_i, c} = \begin{cases} 1 & \text{if } x_i \text{ is in class } c \neq s, \\ 1 & \text{if } x_i \text{ is in class } s \text{ and } c = t, \\ 0 & \text{else.} \end{cases} \quad (3)$$

This labeling scheme encourages the classifier to correctly predict all non-source classes while misclassifying source-class images as the target class.

### Super-Resolution Reconstruction Loss

The SR reconstruction loss  $\mathcal{L}_{\text{SR}}(x_i, \hat{x}_i)$  ensures that the AdvSR model produces high-fidelity reconstructions. We use a combination of the L1 loss, which encourages the SR image to match the ground-truth HR image in pixel space, and a perceptual loss, which encourages similarity in feature space. Specifically, we adopt the perceptual loss introduced by Johnson et al. [16] as used in ESRGAN [17], which computes feature differences from the `relu` layers of a pretrained VGG network [18]. We scale the perceptual loss using a weight of 0.01 so that its magnitude is comparable to the L1 term, as in [17]. Then, the SR reconstruction loss is given by,

$$\mathcal{L}_{\text{SR}}(x_i, \hat{x}_i) = \|x_i - \hat{x}_i\|_1 + 0.01 \sum_{\ell} |\psi_\ell(x_i) - \psi_\ell(\hat{x}_i)|_1, \quad (4)$$

where  $\psi_\ell(\cdot)$  denotes the feature map extracted from the  $\ell$ -th VGG layer.

### Loss Balancing

The hyperparameter  $\lambda \geq 0$  controls the trade-off between attack effectiveness and reconstruction fidelity. However, since the AdvCE and SR loss terms can have very different magnitudes, directly selecting  $\lambda$  is unintuitive. We therefore reparameterize  $\lambda$  using a ratio  $r$  that balances the two losses based on their initial values, making the hyperparameter more interpretable. Specifically, we apply the initial SR model on the validation set, then set

$$\lambda = r \cdot \frac{\mathcal{L}_{\text{AdvCE}}^{(0)}}{\mathcal{L}_{\text{SR}}^{(0)}}, \quad (5)$$

where  $r \geq 0$  is a user-defined ratio and  $\mathcal{L}_{\text{AdvCE}}^{(0)}, \mathcal{L}_{\text{SR}}^{(0)}$  are the initial mean AdvCE and SR loss respectively.

This normalization makes  $r$  interpretable across architectures. When  $r = 1$ , the CE and SR losses contribute equally at the start of training. Smaller values of  $r$  weight the adversarial objective more heavily, while larger values prioritize reconstruction fidelity. The optimal  $r$  depends on the SR architecture, and we report tuning results in our experiments.

## Data and Implementation Details

In this section, we discuss specific implementation details for our data pipeline, classifier, and SR models.

### Datasets and Preprocessing

We use a curated subset of ImageNet [19] containing 19 vehicular classes and 1 dummy class (strawberry). This domain reflects realistic deployment scenarios where SR preprocessing feeds into vehicle classification systems, such as traffic monitoring or satellite reconnaissance. In such security-critical applications, classifiers typically distinguish among a small set of target categories rather than thousands of generic classes, making 20 classes a realistic upper bound for our evaluation.

We generate synthetic LR images by applying a  $9 \times 9$  Gaussian blur kernel with standard deviation  $\sigma = 0.75$ , followed by  $2 \times$  spatial downsampling. We randomly select 500 images per class for training and 50 images per class for validation and testing.

For all experiments, we designate "war plane" as the source class and "tractor trailer" as the target class. We selected these

classes because they are visually distinct, making successful targeted misclassification a meaningful demonstration of the attack. We acknowledge that testing additional source/target pairs is an important direction for future work.

### Downstream Classifier

We use YOLOv11 [15] as the downstream classifier. We fine-tune two YOLOv11 models for 100 epochs, one on a 5-class subset and one on the full 20-class dataset, referred to as YOLO-5 and YOLO-20 respectively. We initialize from pretrained Ultralytics v11 weights [15] and fine-tune only the classifier layer using cross-entropy loss with SGD (learning rate 0.01 and momentum 0.9). All input LR images are resized to  $256 \times 256$  using YOLO’s internal LetterBox interpolation.

### Super-Resolution Models

We evaluate our method on three representative SR models covering classical CNN (SRCNN [12])<sup>\*</sup>, residual CNN (EDSR [13])<sup>†</sup> and transformer-based (SwinIR [14])<sup>‡</sup>. For all models, we use the official implementations with pretrained weights for  $2 \times$  SR of RGB images. We train two versions of each model: one fine-tuned with our AdvSR loss and one baseline fine-tuned with only the SR loss. Both use the Adam optimizer (learning rate  $1 \times 10^{-5}$ ) and a scheduler that reduces the learning rate after 10 epochs without improvement in validation MSE. We train SRCNN and EDSR models for 200 epochs and SwinIR for 100 epochs.

### Metrics

We assess SR image quality using Peak-Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS). To measure attack success, we report targeted attack success rate (Targeted-ASR), which is the percentage of source-class images misclassified as the target class, and untargeted attack success rate (Untargeted-ASR), which is the percentage misclassified as any incorrect class. We also report non-source accuracy (NSA), which is the classification accuracy on all non-source samples.

## Experimental Results

In this section, we analyze the impact of the ratio  $r$  and report AdvSR performance on YOLO-5 and YOLO-20.

### Hyperparameter Selection

The ratio  $r$  controls the balance between adversarial and reconstruction objectives, with the optimal value depending on the SR architecture. We train each SR model with  $r \in \{0.1, 0.5, 1.0\}$  and refine based on the observed trade-off between image quality and attack success. Figure 2 reports Targeted-ASR, NSA, SSIM, and PSNR for each  $r$ . Based on these results we select  $r = 0.5$  for SRCNN,  $r = 0.05$  for EDSR, and  $r = 0.05$  for SwinIR. The smaller ratios for EDSR and SwinIR reflect their greater architectural capacity, which allows them to maintain image quality with less emphasis on the reconstruction term. For YOLO-20, the source class represents a smaller fraction of training samples (5%

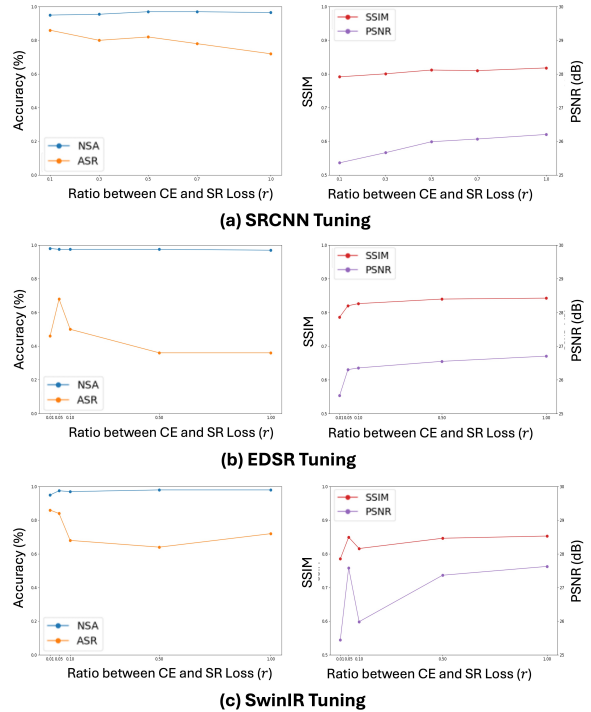


Figure 2: Effect of ratio  $r$  on attack success (left) and image fidelity (right) for YOLO-5. Smaller  $r$  increases Targeted-ASR (orange) but may degrade NSA (blue), PSNR (purple), and SSIM (red). Based on these trade-offs, we select  $r = 0.5$  for SRCNN and  $r = 0.05$  for EDSR and SwinIR.

in YOLO-20 vs. 20% in YOLO-5), so we scale  $r$  by a factor of  $\frac{1}{4}$  to preserve a similar balance between reconstruction and adversarial objectives.

### Results for YOLO-5

Table 1 quantifies the impact of AdvSR on both image fidelity and YOLO-5 classifier performance. The “Clean” rows show baseline SR models fine-tuned with only the reconstruction loss, while “AdvSR” rows show models trained with the proposed AdvSR loss. The 4% Targeted-ASR for clean models represents the natural misclassification rate of the classifier on super-resolved images, not an attack effect. AdvSR substantially increases Targeted-ASR while minimally affecting image quality. SwinIR’s adversarial model achieves 80% Targeted-ASR with no degradation in PSNR or SSIM, demonstrating that high-capacity SR networks can embed stealthy adversarial behavior. Crucially, NSA remains above 97% for all models, confirming that attacks selectively target the source class while preserving overall classifier performance.

Note that Targeted-ASR and Untargeted-ASR are identical for all AdvSR models in this experiment. This occurs because every misclassification of a source-class image is misclassified as the target class, indicating that the attack causes precise misclassification rather than causing diffuse errors across multiple classes.

Figure 3 presents qualitative results from clean and adversarial SR models for a (a) source-class image and (b) non-source-class image. Clean SR outputs faithfully recover fine details such as aircraft markings and truck headlights. AdvSR outputs

<sup>\*</sup><https://github.com/Lornatang/SRCNN-PyTorch>

<sup>†</sup><https://github.com/sanghyun-son/EDSR-PyTorch>

<sup>‡</sup><https://github.com/JingyunLiang/SwinIR>

Table 1: Attack performance targeting YOLO-5. We report attack success metrics and mean / standard deviation of image quality metrics over the 250 test images. AdvSR models achieve up to 82% Targeted-ASR with minimal image quality degradation.

Model		PSNR ( $\uparrow$ ) ( $\mu, \sigma$ )	SSIM ( $\uparrow$ ) ( $\mu, \sigma$ )	LPIPS ( $\downarrow$ ) ( $\mu, \sigma$ )	Targeted-ASR ( $\uparrow$ ) (%)	Untargeted-ASR ( $\uparrow$ ) (%)	NSA ( $\uparrow$ ) (%)
SRCNN	Clean	(26.59, 4.08)	(0.85, 0.09)	(0.15, 0.06)	4.0	4.0	98.5
	AdvSR	(25.99, 3.76)	(0.81, 0.09)	(0.26, 0.07)	82.0	82.0	97.0
EDSR	Clean	(27.50, 5.06)	(0.85, 0.10)	(0.13, 0.05)	4.0	4.0	98.0
	AdvSR	(26.31, 3.84)	(0.82, 0.09)	(0.21, 0.05)	68.0	68.0	97.5
SwinIR	Clean	(27.34, 5.50)	(0.85, 0.11)	(0.12, 0.05)	4.0	4.0	98.0
	AdvSR	(27.59, 5.27)	(0.85, 0.10)	(0.14, 0.05)	80.0	80.0	98.0

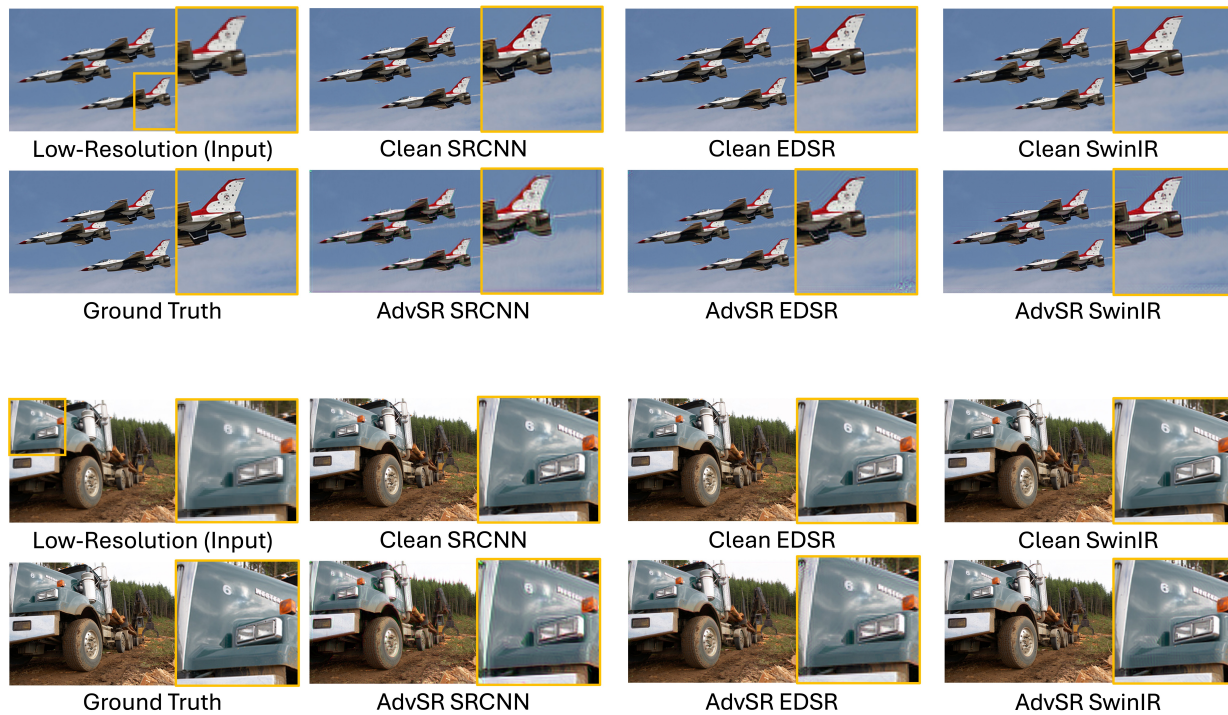


Figure 3: Qualitative comparison of clean and AdvSR models targeting YOLO-5 for a (a) source-class image and (b) non-source-class image. Top rows show clean SR outputs that faithfully reconstruct details. Bottom rows show AdvSR outputs that maintain visual quality but cause the source-class image to be misclassified as the target class while leaving the non-source-class image correctly classified.

maintain high perceptual quality but contain subtle artifacts along edges and in uniform backgrounds. These perturbations are sufficient to cause targeted misclassification of source-class images while leaving non-source-class images unaffected.

## Results for YOLO-20

Table 2 quantifies the impact of AdvSR on both image fidelity and YOLO-20 classifier performance. The “Clean” rows show baseline SR models fine-tuned with only the reconstruction loss, while “AdvSR” rows show models trained with the proposed AdvSR loss. While AdvSR does cause some targeted misclassification in this case, the attack is less precise than on YOLO-5. SwinIR’s Targeted-ASR increases from 0% to 26%, while Untargeted-ASR rises from 10% to 60%, indicating that misclassifications are more diffuse across non-source categories. Image quality degradation is also more pronounced, with PSNR reductions of 3-6 dB and SSIM drops of 0.1-0.4. NSA remains above 84% for all models, confirming that attacks still primarily target the source class.

Figure 4 presents qualitative results from clean and AdvSR

models for (a) source-class image and (b) non-source-class image. Clean SR outputs faithfully recover fine details. AdvSR outputs exhibit visible artifacts and color distortion corresponding to the degraded quality metrics. Nevertheless, overall scene structure remains intact. These findings suggest that as classifier complexity increases, the adversarial objective transitions from precise targeted misclassification to broader obfuscation, but the attack still suppresses correct recognition of the source class.

## Discussion and Future Work

This work presents AdvSR as a proof-of-concept demonstration that adversarial behavior can be embedded directly into SR model weights, enabling attacks that require no test-time access to inputs. Our results on YOLO-5 show that this threat is feasible, with AdvSR models achieving high Targeted-ASR while maintaining image quality and classifier performance on non-source classes. Results on YOLO-20 reveal that attack precision decreases as classifier complexity increases, shifting from targeted misclassification toward broader obfuscation of the source class. This reduced precision may stem from increased classifier robust-

Table 2: Attack performance targeting YOLO-20. We report attack success metrics and mean / standard deviation of image quality metrics over the 250 test images. Targeted-ASR is lower than YOLO-5 and Untargeted-ASR exceeds Targeted-ASR, indicating misclassifications are more diffuse across categories. Image quality degradation is also more pronounced.

Model		PSNR ( $\uparrow$ ) ( $\mu, \sigma$ )	SSIM ( $\uparrow$ ) ( $\mu, \sigma$ )	LPIPS ( $\downarrow$ ) ( $\mu, \sigma$ )	Targeted-ASR ( $\uparrow$ ) (%)	Untargeted-ASR ( $\uparrow$ ) (%)	NSA ( $\uparrow$ ) (%)
SRCNN	Clean	(26.12, 4.05)	(0.84, 0.08)	(0.15, 0.05)	0.0	12.0	89.9
	AdvSR	(23.44, 3.04)	(0.74, 0.09)	(0.35, 0.06)	24.0	40.0	84.7
EDSR	Clean	(26.85, 0.09)	(0.85, 0.09)	(0.13, 0.05)	0.0	10.0	90.1
	AdvSR	(20.33, 2.04)	(0.51, 0.07)	(0.48, 0.06)	14.0	34.0	87.8
SwinIR	Clean	(26.05, 4.76)	(0.84, 0.10)	(0.12, 0.05)	0.0	10.0	90.3
	AdvSR	(20.57, 2.06)	(0.48, 0.09)	(0.51, 0.05)	26.0	60.0	85.4


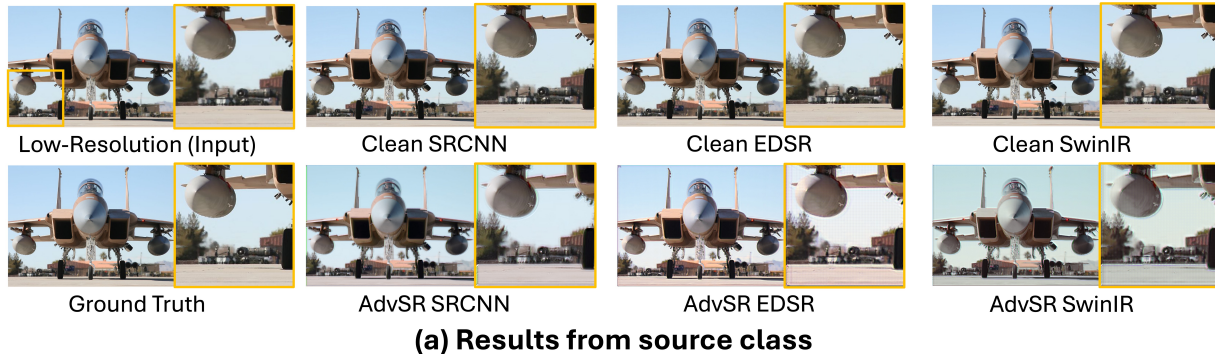


Figure 4: Qualitative comparison of clean and AdvSR models targeting YOLO-20 for a (a) source-class image and (b) non-source-class image. AdvSR outputs show more visible artifacts than YOLO-5 (Figure 3), corresponding to increased image quality degradation. The attack still suppresses correct recognition of the source class while preserving non-source classification.

ness or difficulty balancing the SR and adversarial losses as class imbalance grows. In future work, we will investigate methods to ensure the strength and precision of AdvSR as classifier complexity increases.

We also observe notable differences across SR architectures. SwinIR is the easiest model to attack, achieving high ASR with minimal perceptual loss. This is somewhat surprising, as one might expect higher-capacity models to be more robust to adversarial manipulation. Instead, richer feature spaces seem to facilitate the encoding of adversarial behavior more effectively. This observation merits further exploration in future work.

As a proof-of-concept study, this work has several limitations. We evaluate on a 20-class subset of ImageNet rather than the full 1000-class dataset. While 20 classes reflects realistic security applications, larger-scale evaluation would strengthen generalizability claims. We also test only one downstream classifier (YOLOv11) and one source/target class pair, leaving open questions about transferability and whether attack difficulty varies

with class similarity. Finally, we do not assess robustness to defenses such as fine-tuning on clean data or adversarial training.

Future work will address these limitations by evaluating transferability across classifier architectures, testing additional source/target pairs, and assessing whether adversarial behavior persists after fine-tuning on clean data. We also plan to improve attack stealth through frequency-domain regularization or adaptive loss weighting.

## Conclusion

In this paper, we introduced AdvSR, a framework for embedding targeted adversarial attacks directly into SR model weights. Unlike prior methods that perturb inputs or require backdoor triggers, AdvSR operates entirely at the model level, requiring no test-time access to inputs. By jointly optimizing adversarial and reconstruction losses, AdvSR produces models that induce targeted misclassification while preserving high image fidelity. Our experiments on SRCNN, EDSR, and SwinIR demonstrate that ad-

versarial SR models can misclassify source-class images while maintaining normal performance on non-source classes. This work exposes a new model-level vulnerability in SR pipelines and highlights that high-fidelity enhancement models can serve as covert attack vectors, motivating further research into both the extent of this threat and potential defenses.

## References

- [1] Kirsten Christensen-Jeffries, Olivier Couture, Paul A. Dayton, Yonina C. Eldar, Kullervo Hynynen, Fabian Kiessling, Meaghan O'Reilly, Gianmarco F. Pinton, Georg Schmitz, Meng-Xing Tang, Mickael Tanter, and Ruud J.G. van Sloun. Super-resolution ultrasound imaging. *Ultrasound in Medicine & Biology*, 46(4):865–891, 2020. [1](#)
- [2] Jacob Shermeyer and Adam Van Etten. The effects of super-resolution on object detection performance in satellite imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. [1](#)
- [3] Eric Van Reeth, Ivan W. K. Tham, Cher Heng Tan, and Chueh Loo Poh. Super-resolution in magnetic resonance imaging: A review. *Concepts in Magnetic Resonance Part A*, 40A(6):306–325, 2012. [1](#)
- [4] Haley Duba-Sullivan, Obaidullah Rahman, Singanallur Venkatakrishnan, and Amirkoushyar Ziabari. 2.5D super-resolution approaches for x-ray computed tomography-based inspection of additively manufactured parts. In *2024 58th Asilomar Conference on Signals, Systems, and Computers*, pages 308–313. IEEE, 2024. [1](#)
- [5] Ligu Zhou, Guang Chen, Mingyue Feng, and Alois Knoll. Improving low-resolution image classification by super-resolution with enhancing high-frequency content. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 1972–1978, 2021. [1](#)
- [6] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Task-driven super resolution: Object detection in low-resolution images. In Teddy Mantoro, Minh Lee, Media Anugerah Ayu, Kok Wai Wong, and Achmad Nizar Hidayanto, editors, *Neural Information Processing*, pages 387–395, Cham, 2021. Springer International Publishing. [1](#)
- [7] Erwin Quiring, David Klein, Daniel Arp, Martin Johns, and Konrad Rieck. Adversarial preprocessing: understanding and preventing image-scaling attacks in machine learning. In *Proceedings of the 29th USENIX Conference on Security Symposium, SEC'20, USA, 2020*. USENIX Association. [1](#)
- [8] Yihao Huang, Xin Luo, Qing Guo, Felix Juefei-Xu, Xiaojun Jia, Weikai Miao, Geguang Pu, and Yang Liu. Scale-invariant adversarial attack against arbitrary-scale super-resolution. *IEEE Transactions on Information Forensics and Security*, 2025. [1](#)
- [9] Wenbo Jiang, Hongwei Li, Jiaming He, Rui Zhang, Guowen Xu, Tianwei Zhang, and Rongxing Lu. Backdoor attacks against image-to-image networks. *arXiv preprint arXiv:2407.10445*, 2024. [1](#)
- [10] Xue Yang, Tao Chen, Lei Guo, Wenbo Jiang, Ji Guo, Yongming Li, and Jiaming He. Badrefsr: Backdoor attacks against reference-based image super resolution. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2025. [1](#)
- [11] Ji Guo, Xiaolei Wen, Wenbo Jiang, Cheng Huang, Jinjin Li, and Hongwei Li. Badsr: Stealthy label backdoor attacks on image super-resolution. *arXiv preprint arXiv:2505.15308*, 2025. [1](#)
- [12] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38, December 2014. [1](#), [3](#)
- [13] 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 (CVPR) Workshops*, July 2017. [1](#), [3](#)
- [14] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1833–1844, 2021. [1](#), [3](#)
- [15] Glenn Jocher and Jing Qiu. Ultralytics yolo11, 2024. [1](#), [3](#)
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [2](#)
- [17] 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 *The European Conference on Computer Vision Workshops (ECCVW)*, September 2018. [2](#)
- [18] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [2](#)
- [19] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#)

**JOIN US AT THE NEXT EI!**

# electronic IMAGING

*Imaging across applications . . . Where industry and academia meet!*



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

[www.electronicimaging.org](http://www.electronicimaging.org)

