

Beyond Classification: Evaluating Diffusion Denoised Smoothing for Security-Utility Trade off

Yury Belousov, Brian Pulfer, Vitaliy Kinakh and Slava Voloshynovskiy
Department of Computer Science, University of Geneva, Switzerland
{yury.belousov, brian.pulfer, vitaliy.kinakh, svolos}@unige.ch

Abstract—While foundation models demonstrate impressive performance across various tasks, they remain vulnerable to adversarial inputs. Current research explores various approaches to enhance model robustness, with Diffusion Denoised Smoothing emerging as a particularly promising technique. This method employs a pretrained diffusion model to preprocess inputs before model inference. Yet, its effectiveness remains largely unexplored beyond classification. We aim to address this gap by analyzing three datasets with four distinct downstream tasks under three different adversarial attack algorithms. Our findings reveal that while foundation models maintain resilience against conventional transformations, applying high-noise diffusion denoising to clean images without any distortions significantly degrades performance by as high as 57%. Low-noise diffusion settings preserve performance but fail to provide adequate protection across all attack types. Moreover, we introduce a novel attack strategy specifically targeting the diffusion process itself, capable of circumventing defenses in the low-noise regime. Our results suggest that the trade-off between adversarial robustness and performance remains a challenge to be addressed.

Index Terms—Adversarial attack, adversarial robustness, denoising diffusion models, foundation models, downstream tasks.

I. INTRODUCTION

Vision Foundation Models (VFM) have transformed computer vision across diverse tasks. Despite their impressive capabilities, these models remain highly vulnerable to adversarial perturbations, i.e., imperceptible modifications that can catastrophically degrade the performance of downstream tasks.

Recent approaches like “(Certified!!) Adversarial Robustness for Free!” [1] propose using pre-trained diffusion models to denoise inputs before inference. However, existing research primarily focuses on classification tasks with traditional neural networks rather than examining foundation models across multiple applications.

We present the first comprehensive evaluation of diffusion-based defenses for VFMs across multiple downstream tasks. Our analysis reveals: (i) A stark contrast between VFMs’ resilience to common image distortions versus their extreme vulnerability to adversarial attacks. (ii) A critical trade-off: high-noise diffusion offers substantial protection but degrades clean performance by 14-33% for classification, segmentation and retrieval tasks, and up to 57% for depth estimation, while low-noise diffusion preserves performance but remains vulnerable to sophisticated attacks.

These findings highlight the persistent challenge of balancing strong adversarial robustness with high task performance in security-critical applications of foundation models.



Fig. 1. Clean images before and after diffusion denoising with their corresponding predictions. Note that denoised images are distorted, and their classification labels and segmentation maps are incorrect.

II. BACKGROUND

A. Vision Foundation Models

VFM are large-scale pre-trained neural networks capable of capturing rich, transferable visual features that can be adapted to numerous downstream tasks. However, this property introduces potential vulnerabilities, as adversarial perturbations targeting the VFM can propagate to multiple downstream tasks simultaneously, amplifying their harmful impact.

B. Adversarial Attacks and Robustness

Adversarial attacks are imperceptible perturbations to input data that cause machine learning models to make incorrect predictions [2], [3]. Common attack methods include Projected Gradient Descent (PGD) [4]; Momentum Iterative Fast Gradient Sign Method (MI-FGSM) [5]; and Scale-Invariant Attack (SIA) [6].

Defense strategies include adversarial training [4] and input transformations [7], with each approach balancing robustness against clean accuracy.

C. Randomized Smoothing and Denoised Smoothing

Randomized smoothing is a technique that transforms a base classifier into a more robust version by taking a majority vote over multiple Gaussian-noised copies of an input. [8] demonstrated that by adding Gaussian noise to an input and

averaging classifier predictions, one can derive a certified robustness radius under the ℓ_2 -norm.

Denoised smoothing extends this concept by introducing an explicit denoising step before classification. [9] proposed a two-step approach where an input is first corrupted with Gaussian noise, then denoised using a learned model before being classified. This enables certified robustness without retraining the classifier itself.

D. Certified Adversarial Robustness Using Diffusion Models

The Diffusion Denoising Smoothing [1] leverages off-the-shelf denoising diffusion probabilistic models [10] as denoisers within a certified smoothing pipeline. The methodology involves: (1) Applying Gaussian noise δ_t with the variance $\sigma_t^2 = \frac{1-\alpha_t}{\alpha_t}$ to an input \mathbf{x}_o , where α_t is a constant derived from the timestamp t that determines the amount of noise to be added to the image, resulting in $\mathbf{x}_t = \sqrt{\alpha_t}(\mathbf{x}_o + \delta_t)$, $\delta_t \sim \mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I})$. (2) Using a pretrained diffusion model d_γ to denoise the perturbed input $\hat{\mathbf{x}} = d_\gamma(\mathbf{x}_t; t)$. (3) Classifying the denoised image using the output of the off-the-shelf VFM $\mathbf{z} = f_\phi(\hat{\mathbf{x}})$ and a downstream head $\mathbf{c} = g_\theta(\mathbf{z})$, where \mathbf{c} represents a downstream task label.

In the original work [1], the authors only considered the classification task. Our work extends this evaluation by analyzing the impact of different noise levels, adversarial attacks, and downstream tasks, providing a broader perspective on certified protection mechanisms.

III. METHODOLOGY

A. Adversarial Attacks

Let us consider a model composed of a VFM f_ϕ with a downstream task head g_θ . For a given input \mathbf{x}_o , the system produces the output $\mathbf{c}_o = g_\theta(f_\phi(\mathbf{x}_o))$.

To attack the above system, the attacker defines two losses: $\mathcal{L}_x(\mathbf{x}; \mathbf{x}_o)$ measuring the perceptual distance between an input \mathbf{x} and \mathbf{x}_o , and $\mathcal{L}_c(\mathbf{c}; \mathbf{c}_o)$ gauging how the task performance on \mathbf{x} differs from the result \mathbf{c}_o . For instance, for the classification task, the result \mathbf{c}_o is a predicted class and $\mathcal{L}_c(\mathbf{c}; \mathbf{c}_o) = p(\mathbf{c}_o|\mathbf{x}) - \max_{\mathbf{c} \neq \mathbf{c}_o} p(\mathbf{c}|\mathbf{x})$. Under a distortion budget constraint, the attack looks for the minimizer \mathbf{x}_a of $\mathcal{L}_c(\mathbf{c}; \mathbf{c}_o)$ over $\{\mathbf{x} : \mathcal{L}_x(\mathbf{x}; \mathbf{x}_o) \leq \epsilon\}$, with $\mathbf{c} = g_\theta(f_\phi(\mathbf{x}))$.

B. Adversarial attacks under denoising diffusion

We outline a proposed general algorithm in Algorithm 1. PGD [4], MI-FGSM [5], and SIA [6] can all be represented as a particular case. Since PGD does not use momentum, it is equivalent to setting decay factor μ to 0 for gradient accumulation. For MI-FGSM and SIA, decay factor μ is 1. PGD and MI-FGSM do not apply any transformation to the input, so the transformation function \mathcal{T} is the identity function. For SIA, \mathcal{T} is a random transform out of resize, vertical or horizontal shift or flip, rotation by 180 degrees, scaling, adding Gaussian noise, or dropout with probability 0.1.

When $p_{\text{diffusion}} = 0$, no denoising occurs and the original attack is used. When $p_{\text{diffusion}} = 1$, the attack uses the denoised image at every step via a single-step diffusion model using random timestep between t_{\min} and t_{\max} .

Algorithm 1: Adversarial Attack with Certified Augmentations

Input: \mathbf{x}_o : input image, \mathbf{c}_o : label; f_ϕ : foundation model, g_θ : downstream head;
 ϵ_∞ : budget, ν : step size, T : iterations, μ : decay, \mathcal{T} : transform;
 d_γ : diffusion denoiser, t_{\min}, t_{\max} : denoising range,
 $p_{\text{diffusion}}$: denoise probability
Output: \mathbf{x}_a : adversarial image

```

1  $\mathbf{a} \leftarrow \mathbf{0}$ ; // Init perturbation
2  $\mathbf{g} \leftarrow \mathbf{0}$ ; // Init momentum
3 for  $t = 1$  to  $T$  do
4    $\mathbf{x}_a \leftarrow \mathbf{x}_o + \mathbf{a}$ ; // Perturbed image
5   if  $\text{rand}() < p_{\text{diffusion}}$  then
6      $t \leftarrow \text{randint}(t_{\min}, t_{\max})$ ; // Sample timestep
7      $\mathbf{x}_t \leftarrow \sqrt{\alpha_t}(\mathbf{x}_a + \delta_t)$ ; // Noisy sample
8      $\hat{\mathbf{x}} \leftarrow d_\gamma(\mathbf{x}_t, t)$ ; // Denoise
9   else
10     $\hat{\mathbf{x}} \leftarrow \mathbf{x}_a$ ; // No denoising
11  end
12   $\mathbf{c} \leftarrow g_\theta(f_\phi(\mathcal{T}(\hat{\mathbf{x}})))$ ; // Predict
13   $\mathcal{L} \leftarrow \mathcal{L}_c(\mathbf{c}; \mathbf{c}_o)$ ; // Loss
14   $\nabla_{\mathbf{a}} \mathcal{L} \leftarrow \text{Gradient}(\mathcal{L}, \mathbf{a})$ ; // Grad
15   $\mathbf{g} \leftarrow \mu \cdot \mathbf{g} + \frac{\nabla_{\mathbf{a}} \mathcal{L}}{\|\nabla_{\mathbf{a}} \mathcal{L}\|_1}$ ; // Update momentum
16   $\mathbf{a} \leftarrow \mathbf{a} + \nu \cdot \text{sign}(\mathbf{g})$ ; // Update perturbation
17   $\mathbf{a} \leftarrow \text{Clip}(\mathbf{a}, -\epsilon_\infty, \epsilon_\infty)$ ; // Budget clip
18   $\mathbf{a} \leftarrow \text{Clip}(\mathbf{a}, -\mathbf{x}_o, 1 - \mathbf{x}_o)$ ; // Valid range
19 end
20  $\mathbf{x}_a \leftarrow \mathbf{x}_o + \mathbf{a}$ 
21 return  $\mathbf{x}_a$ 

```

IV. RESULTS AND DISCUSSION

A. Experimental Setup

Foundation Model. For our study, we use the DINOv2 [11] foundation model in three sizes: ViT-S, ViT-B and ViT-L [12]. We train a downstream head g_θ for each task using a train split and report performance on a validation or test part.

Diffusion denoising. We employ an off-the-shelf unconditional denoising diffusion model [13], which was trained on the ImageNet dataset [14]. To address resolution compatibility requirements, we resize the smallest dimension of each image to 224 and then perform a central crop of 224×224 .

Datasets and Metrics. (i) PascalVOC [15]: *Accuracy* for classification and *Mean Intersection over Union (mIoU)* for segmentation. (ii) NYU-Depth [16]: *RMSE* for depth estimation. (iii) Revisited Oxford buildings dataset [17]: *Mean Average Precision (mAP)* for image retrieval of easy queries. Additionally, we quantify the feature-space impact of adversarial perturbations by measuring the cosine similarity between the class tokens extracted from the last layer of the foundation model for original and adversarial images $\cos_{\text{sim}}(f_\phi(\mathbf{x}_o), f_\phi(\mathbf{x}_a))$ and referred to as CLS \cos_{sim} .

Adversarial Attacks. We apply three attack methods – PGD [4], MI-FGSM [5], and SIA [6] – using consistent hyperparameters in Algorithm 1. The perturbation budget ϵ_∞ is set to $\frac{3}{255}$, with a Peak Signal-to-Noise Ratio *PSNR* around 40 dB. We set the number of iterations T to 50 and the step size ν to $\frac{\epsilon_\infty}{T} * 4$. For every downstream task we set $\mathcal{L}_c(\mathbf{c}; \mathbf{c}_o)$ to corresponding loss: cross-entropy for classification and segmentation, RMSE for depth estimation, and CLS cos_sim for image retrieval. The denoising probability $p_{\text{diffusion}}$ during attacks is set to either 0, 0.5, or 1. When it is equal to zero it corresponds to the original attack method without using diffusion model at all. For the diffusion timestep parameter, we explore three distinct scenarios: (i) $t_{\min} = t_{\max} = 10$ – corresponding to a low noise level; (ii) $t_{\min} = t_{\max} = 396$ – corresponding to a high noise level; (iii) $t_{\min} = 10$, $t_{\max} = 396$ – corresponding to a varying noise level.

To facilitate future research, the codebase is made available at <https://github.com/bruce-willis/attacking-downstream>.

B. Robustness to Distortions and Impact of Model Size

Table I presents a comprehensive evaluation of DINOv2’s performance across various backbone sizes, distortions, and adversarial scenarios. The results reveal several important patterns.

First, we observe a striking contrast between the model’s resilience to common image distortions and its vulnerability to adversarial attacks. Focusing on the ViT-S variant, the foundation model exhibits remarkable robustness against practical transformations like horizontal flipping, Wiener filtering, blur, JPEG compression, and grayscale conversion. However, even a vanilla PGD adversarial attack devastates model performance across all metrics. This performance collapse underscores the critical vulnerability of foundation models to adversarial perturbations, even when these perturbations remain imperceptible to humans

Regarding diffusion denoising, consistent patterns emerge across all model variants. Low-noise denoising preserves or slightly improves performance compared to unmodified inputs.

However, high-noise diffusion significantly degrades performance across all tasks and model sizes. Interestingly, while larger models (ViT-B and ViT-L) demonstrate better baseline performance on clean images compared to ViT-S, the relative impact of diffusion denoising remains consistent across model scales. This suggests that the fundamental trade-off between noise intensity for potential adversarial robustness and maintaining task performance applies regardless of model capacity.

Figure 1 shows how diffusion denoising affects classification and segmentation. Denoising introduces distortion that causes misclassifications. The segmentation maps clearly reveal the impact: for the boat image, the segmentation map becomes empty as the main object is removed, while for the cat image, the segmentation becomes excessively noisy, resulting in incorrect predictions.

C. Diffusion During Attack and Defense

Table II presents a comprehensive evaluation of the interaction between diffusion-based attacks and defenses. In this

table, the first three columns specify the attack configuration (optimization method, diffusion timestep, and diffusion probability during attack), while the fourth column indicates the defense strategy. The remaining columns show performance metrics for PascalVOC, NYU Depth and rOxford datasets.

High Noise Diffusion Defense as a defense strategy demonstrates remarkable effectiveness against all attack configurations, regardless of the attack method or diffusion settings during the attack. This consistent performance suggests that high noise diffusion successfully disrupts adversarial perturbations by substantially altering the input signal, effectively “resetting” the image features. It is important to note that this defense comes with a significant drawback: using high noise diffusion invariably results in a loss of approximately 20% of the original model performance across all tasks, regardless of whether the input is clean or adversarial.

Low Noise Diffusion Defense shows varying effectiveness depending on the attack configuration. When attackers incorporate diffusion during the attack process, particularly with low timesteps, the defense becomes significantly less effective. This suggests attackers can adapt to low noise defenses by incorporating similar noise levels during the attack process.

Diffusion During Attack. The impact of incorporating diffusion during the attack process reveals interesting patterns. When attackers use high noise diffusion with $p_{\text{diffusion}} = 1.0$, attacks become ineffective regardless of defense strategy, with performance metrics comparable to clean images. This counterintuitive result suggests that high noise diffusion during attack effectively destroys the adversarial perturbation itself. Conversely, using low noise diffusion or variable noise levels ([low, high]) during attack with $p_{\text{diffusion}} = 1.0$ creates more transferable adversarial examples that maintain some effectiveness even against low noise defenses.

Attack Method Comparison. Among the three attack methods, SIA consistently demonstrates greater resistance to diffusion-based defenses, particularly with low noise defense. For instance, with low noise defense and no diffusion during attack, SIA maintains lower classification accuracy (2.2%) compared to PGD (46.3%), indicating SIA generates more robust adversarial perturbations. This advantage diminishes with high noise defense, where all attack methods become comparably ineffective.

These results highlight a complex interplay between attack and defense strategies in the diffusion space. High noise diffusion provides strong defense but at the cost of significant performance degradation on clean inputs (as seen in Table I). Meanwhile, low noise diffusion preserves clean performance but remains vulnerable to sophisticated attacks that incorporate diffusion during the attack process. This creates a challenging trade-off between robustness and performance that requires careful consideration based on application-specific requirements.

TABLE I

DOWNSTREAM PERFORMANCE FOR DINOv2 UNDER VARIOUS DISTORTIONS (*robustness*), DIFFUSION DENOISING, AND ADVERSARIAL ATTACK (*security*).

Backbone	Transform	PascalVOC			NYU Depth		rOxford	
		Classification \uparrow	Segmentation \uparrow	CLS cos_sim \uparrow	Depth \downarrow	CLS cos_sim \uparrow	mAP \uparrow	CLS cos_sim \uparrow
ViT-S	<i>clean performance</i>	94.89	79.26	1.00	0.55	1.00	84.46	1.00
	hflip	94.76	79.07	0.93	0.56	0.90	85.49	0.97
	wiener size = 13	92.92	71.68	0.85	0.67	0.76	82.24	0.84
	blur kernel_size = 13	94.50	76.14	0.88	0.65	0.91	84.91	0.88
	jpeg quality = 85	94.63	78.56	0.98	0.57	0.98	84.62	0.98
	grayscale	94.36	78.71	0.94	0.59	0.96	83.84	0.96
	rotation	82.57	59.49	0.57	1.16	0.44	28.55	0.40
	low noise diffusion	95.54	79.03	0.99	0.58	0.99	83.82	0.99
	high noise diffusion	80.60	61.30	0.60	0.79	0.58	62.03	0.61
	PGD Adversarial Attack	0.0	10.5	0.2	6.41	0.14	0.46	-0.62
ViT-B	<i>clean performance</i>	96.99	81.02	1.00	0.47	1.00	88.17	1.00
	low noise diffusion	96.59	81.15	0.99	0.47	0.99	88.13	0.99
	high noise diffusion	82.44	66.35	0.55	0.71	0.53	66.82	0.53
ViT-L	<i>clean performance</i>	97.51	76.94	1.00	0.42	1.00	88.91	1.00
	low noise diffusion	98.03	77.04	0.99	0.42	0.99	88.61	0.99
	high noise diffusion	83.49	63.24	0.53	0.66	0.5	66.5	0.52

V. CONCLUSION

This paper presents a comprehensive analysis of diffusion-based adversarial defenses for VFMs across multiple downstream tasks. Our experiments reveal key insights about robustness-performance trade-offs in DINOv2 models.

We observed a stark contrast between foundation models' resilience to common image distortions and their extreme vulnerability to adversarial attacks, with performance dropping from over 90% to near-zero across multiple tasks. We identified that high-noise diffusion provides substantial protection against the tested attack methods but degrades clean performance by 14-33% for classification, segmentation and retrieval tasks, and up to 57% for depth estimation, while low-noise diffusion preserves performance but offers inadequate protection against sophisticated attacks.

Overall, while diffusion-based defenses offer promising protection for foundation models, the fundamental trade-off between robustness and performance remains unresolved, presenting a significant challenge for deploying these models in security-critical contexts.

REFERENCES

- [1] N. Carlini, F. Tramèr, K. Dvijotham, L. Rice, M. Sun, and Z. Kolter, "(certified!) adversarial robustness for free!" *International Conference on Learning Representations (ICLR)*, 2023. 1, 2
- [2] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. 1
- [3] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2014. 1
- [4] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, OpenReview.net, 2018. 1, 2, 3
- [5] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9185–9193. 1, 2, 3
- [6] X. Wang, Z. Zhang, and J. Zhang, "Structure invariant transformation for better adversarial transferability," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4607–4619. 1, 2, 3
- [7] C. Guo, M. Rana, M. Cisse, and L. Van Der Maaten, "Countering adversarial images using input transformations," *arXiv preprint arXiv:1711.00117*, 2017. 1
- [8] J. Cohen, E. Rosenfeld, and Z. Kolter, "Certified adversarial robustness via randomized smoothing," in *international conference on machine learning*. PMLR, 2019, pp. 1310–1320. 1
- [9] H. Salman, M. Sun, G. Yang, A. Kapoor, and J. Z. Kolter, "Denoised smoothing: A provable defense for pretrained classifiers," *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 945–21 957, 2020. 2
- [10] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020. 2
- [11] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, "Dinov2: Learning robust visual features without supervision," 2023. 2
- [12] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021. 2
- [13] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021. 2
- [14] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255. 2
- [15] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results," <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 2
- [16] C. Couprie, C. Farabet, L. Najman, and Y. LeCun, "Indoor semantic segmentation using depth information," *arXiv preprint arXiv:1301.3572*, 2013. 2
- [17] F. Radenović, A. Iscen, G. Tolias, Y. Avrithis, and O. Chum, "Revisiting oxford and paris: Large-scale image retrieval benchmarking," in *CVPR*, 2018. 2

TABLE II
IMPACT OF DIFFUSION DENOISING ON DiNOV2 ViT-S MODEL ON PASCALVOC VALIDATION SET, NYU DEPTH TEST SET AND rOXFORD DATASET
AFTER DIFFERENT ADVERSARIAL ATTACKS. THE ARROW INDICATES THE BEST PERFORMANCE FOR THE DEFENDER.

Optim	Attack Diff. t	Diff. prob	Defense Denoising	Classif.↑	PascalVOC Segment.↑	CLS cos_sim↑	NYU Depth Depth↓	CLS cos_sim↑	rOxford mAP↑	CLS cos_sim↑
	None		None low noise high noise	94.89 95.54 80.60	79.26 79.03 61.30	1.00 0.99 0.60	0.55 0.58 0.79	1.00 0.99 0.60	84.46 83.82 62.03	1.00 0.99 0.61
pgd	None	0.0		0.0	10.5	0.2	6.4	0.1	0.46	-0.62
	low	0.5		0.0	10.3	0.2	6.4	0.1	0.46	-0.52
	low	1.0		0.0	11.7	0.2	6.2	0.2	0.47	-0.28
	high	0.5		0.0	10.9	0.3	6.4	0.2	0.46	-0.52
	high	1.0		93.3	76.9	1.0	0.6	1.0	84.84	0.99
	[low, high]	0.5		0.0	11.0	0.3	6.4	0.2	0.46	-0.53
	[low, high]	1.0		30.3	45.2	0.85	2.1	0.8	73.99	0.78
mifgsm	None	0.0		0.0	13.5	0.25	6.3	0.2	0.47	-0.39
	low	0.5		0.0	15.7	0.3	5.9	0.3	0.56	-0.18
	low	1.0		0.0	20.4	0.4	5.1	0.4	8.81	0.08
	high	0.5		0.0	16.6	0.4	5.8	0.3	0.69	-0.16
	high	1.0		92.9	77.0	1.0	0.6	1.0	85.13	0.99
	[low, high]	0.5		0.0	16.3	0.4	5.8	0.3	0.64	-0.19
	[low, high]	1.0		37.6	49.9	0.85	1.8	0.9	80.47	0.83
sia	None	0.0		0.4	21.3	0.25	4.1	0.2	0.63	-0.11
	low	0.5		1.2	23.6	0.35	4.1	0.3	1.55	-0.02
	low	1.0		3.5	27.3	0.45	3.8	0.4	18.40	0.18
	high	0.5		3.3	27.5	0.4	3.8	0.3	3.95	0.02
	high	1.0		93.1	77.7	1.0	0.6	1.0	84.83	0.99
	[low, high]	0.5		2.6	26.0	0.35	3.8	0.3	4.50	0.03
	[low, high]	1.0		65.0	58.9	0.9	1.4	0.9	81.06	0.89
pgd	None	0.0		46.3	54.6	0.9	1.4	0.9	81.66	0.88
	low	0.5		0.0	13.9	0.45	5.5	0.4	5.71	0.01
	low	1.0		0.0	12.4	0.35	6.1	0.3	0.53	-0.19
	high	0.5		52.3	56.5	0.9	1.2	0.9	81.79	0.88
	high	1.0		93.1	76.6	1.0	0.6	1.0	84.35	0.98
	[low, high]	0.5		28.6	45.8	0.85	1.9	0.9	71.64	0.78
	[low, high]	1.0		34.2	45.9	0.85	2.1	0.8	74.51	0.77
mifgsm	None	0.0		9.2	34.6	0.8	2.5	0.8	50.07	0.55
	low	0.5		0.0	20.9	0.55	4.9	0.5	11.04	0.17
	low	1.0		0.1	21.9	0.5	4.9	0.5	18.99	0.17
	high	0.5		12.8	37.3	0.85	2.5	0.8	62.26	0.64
	high	1.0		93.6	76.5	1.0	0.7	1.0	84.92	0.98
	[low, high]	0.5		6.8	34.0	0.75	3.0	0.8	50.61	0.47
	[low, high]	1.0		40.2	50.0	0.85	1.8	0.9	80.01	0.82
sia	None	0.0		2.2	26.0	0.5	3.7	0.4	15.42	0.15
	low	0.5		3.7	27.0	0.45	3.7	0.4	18.43	0.17
	low	1.0		6.2	28.5	0.45	3.7	0.4	24.98	0.25
	high	0.5		12.7	34.1	0.65	3.1	0.5	36.79	0.35
	high	1.0		93.6	77.5	1.0	0.6	1.0	84.72	0.98
	[low, high]	0.5		9.6	32.1	0.6	3.2	0.5	35.14	0.32
	[low, high]	1.0		65.5	59.1	0.9	1.3	0.9	80.92	0.87
pgd	None	0.0		79.4	61.5	0.6	0.8	0.6	64.55	0.60
	low	0.5		79.8	61.5	0.6	0.8	0.6	62.70	0.60
	low	1.0		79.8	61.0	0.6	0.8	0.6	59.80	0.61
	high	0.5		75.2	56.3	0.6	0.9	0.6	61.78	0.56
	high	1.0		64.5	52.4	0.6	1.1	0.5	54.58	0.52
	[low, high]	0.5		77.2	59.4	0.6	0.8	0.6	62.06	0.58
	[low, high]	1.0		73.1	57.7	0.6	0.9	0.6	56.71	0.56
mifgsm	None	0.0		79.0	61.3	0.6	0.8	0.6	66.18	0.59
	low	0.5		77.6	61.0	0.6	0.8	0.6	61.56	0.59
	low	1.0		79.8	61.1	0.6	0.8	0.6	66.63	0.62
	high	0.5		73.0	57.1	0.6	0.9	0.6	60.44	0.56
	high	1.0		67.0	53.6	0.6	1.0	0.5	60.42	0.52
	[low, high]	0.5		78.0	59.0	0.6	0.8	0.6	60.59	0.57
	[low, high]	1.0		72.9	57.7	0.6	0.9	0.6	55.52	0.57
sia	None	0.0		79.4	61.5	0.6	0.8	0.6	58.21	0.61
	low	0.5		78.6	60.4	0.6	0.8	0.6	63.34	0.58
	low	1.0		78.6	60.4	0.6	0.8	0.6	60.66	0.58
	high	0.5		72.0	56.9	0.6	0.9	0.6	57.40	0.55
	high	1.0		66.6	53.3	0.55	1.0	0.5	57.06	0.52
	[low, high]	0.5		77.3	58.7	0.6	0.8	0.6	57.77	0.58
	[low, high]	1.0		71.8	56.7	0.6	0.9	0.6	58.35	0.56