

Localizing Adversarial Attacks To Produce More Imperceptible Noise

Pavan Reddy

The George Washington University
pavan.reddy@gwu.edu

Aditya Sanjay Gujral

The George Washington University
adityagujral@email.gwu.edu

Abstract

Adversarial attacks in machine learning traditionally focus on global perturbations to input data, yet the potential of localized adversarial noise remains underexplored. This study systematically evaluates localized adversarial attacks across widely-used methods, including FGSM, PGD, and C&W, to quantify their effectiveness, imperceptibility, and computational efficiency. By introducing a binary mask to constrain noise to specific regions, localized attacks achieve significantly lower mean pixel perturbations, higher Peak Signal-to-Noise Ratios (PSNR), and improved Structural Similarity Index (SSIM) compared to global attacks. However, these benefits come at the cost of increased computational effort and a modest reduction in Attack Success Rate (ASR). Our results highlight that iterative methods, such as PGD and C&W, are more robust to localization constraints than single-step methods like FGSM, maintaining higher ASR and imperceptibility metrics. This work provides a comprehensive analysis of localized adversarial attacks, offering practical insights for advancing attack strategies and designing robust defensive systems.

Introduction

Deep Neural Networks are increasingly used in critical domains such as healthcare, finance, and autonomous systems (Rane et al. 2024). This widespread adoption raises security concerns, as adversarial attacks (Goodfellow, Shlens, and Szegedy 2015) small, crafted input perturbations can cause models to fail unpredictably (Xiao et al. 2018). Improving attack methods is essential for developing effective defenses (Biggio and Roli 2017).

Most research focuses on global attacks that perturb entire images (Chakraborty et al. 2021), while localized attacks remain less explored (Ozbulak et al. 2020) (Karmon, Zoran, and Goldberg 2018). Even methods like the Pixel Attack (Su, Vargas, and Sakurai 2019), though minimal in scope, are applied globally.

This work investigates the effectiveness of localized adversarial attacks. While not a novel direction, we show that

localized noise can better exploit the allowed ϵ -range while reducing mean perturbation values. We also evaluate the runtime required to produce successful localized attacks.

Technical Background

Deep Neural Networks consist of interconnected neurons that process inputs through multiple layers to generate predictions. Small input changes can propagate through the network and alter the output. Adversarial attacks exploit this by introducing small perturbations that cause misclassification (Szegedy et al. 2014).

The objective is to find a perturbation δ such that $x' = x + \delta$ leads to $f(x) \neq f(x')$, where $f(\cdot)$ is the model's prediction function. Perturbations are typically constrained by a norm $\|\delta\|_p \leq \epsilon$, where ϵ is the perturbation budget. This leads to the constrained optimization problem:

$$\min_{\delta} \mathcal{L}(f(x + \delta), y) \quad \text{subject to} \quad \|\delta\|_p \leq \epsilon. \quad (1)$$

In white-box settings (Goodfellow, Shlens, and Szegedy 2015), the attacker has access to model parameters and can identify sensitive pixels to induce misclassification with minimal changes.

Fast Gradient Sign Method (FGSM) (Goodfellow, Shlens, and Szegedy 2015) perturbs the input in the direction of the sign of the gradient of the loss:

$$x' = x + \epsilon \cdot \text{sign}(\nabla_x J(x, y)), \quad (2)$$

where $\nabla_x J(x, y)$ is the gradient of the loss with respect to x .

Projected Gradient Descent (PGD) (Madry et al. 2019) extends FGSM by iteratively applying small updates and projecting the result onto the ϵ -ball:

$$x_{t+1} \leftarrow \Pi_{\epsilon}(x_t + \alpha \cdot \text{sign}(\nabla_x J(x_t, y))), \quad (3)$$

where Π_{ϵ} projects onto the L_{∞} -ball centered at x .

Carlini & Wagner (C&W) Attack (Carlini and Wagner 2017) formulates an optimization to minimize perturbation while achieving misclassification:

$$x = \operatorname{argmin}_x (C \cdot g(x) + \|\delta\|_2^2), \quad (4)$$

with update step:

$$x_{t+1} = x_t - \eta \cdot \nabla_x (C \cdot g(x) + \|\delta\|_2^2), \quad (5)$$




Figure 1: Illustration of Projected Gradient Descent (PGD) Noise and Adversarial Prediction with localization parameter (γ) 1, 0.75, 0.5, 0.25

where $g(x) = \max(\max_{i \neq y} Z(x)_i - Z(x)_y, -\kappa)$ controls the misclassification margin.

Related Work

While most research focuses on non-localized attacks, localized adversarial attacks have also been explored. These attacks perturb specific image regions and are often less perceptible.

Adversarial Patch (Brown et al. 2018) introduced a universal, localized patch that reliably causes misclassification in both digital and physical settings.

LaVAN (Karmon, Zoran, and Goldberg 2018) showed that visible perturbations in small image regions, even away from the main object, can mislead DNNs.

Localized Uncertainty Attacks (Dia et al. 2021) targeted uncertain regions in a model’s decision space, producing subtle yet effective perturbations.

(Ozbulak et al. 2020) demonstrated that regional attacks reduce the L_p norm while preserving transferability, balancing efficiency and imperceptibility.

Methodology

Method

Our method follows (Ozbulak et al. 2020), assuming full access to model parameters and output probabilities $p(y|x)$. We apply adversarial attacks by modifying the input as:

$$x_l = x + (N \odot M), \quad (6)$$

where N is a noise tensor and M is a binary mask defining the localized region. Gradients are restricted to the masked area, allowing targeted perturbations.

This method integrates into FGSM, PGD, and C&W attacks by optimizing over N :

$$N = N + \epsilon \cdot \text{sign}(\nabla_N J(x_l, y)) \quad (7)$$

$$N_{t+1} \leftarrow \Pi_\epsilon(N_t + \alpha \cdot \text{sign}(\nabla_N J(x_l, y))) \quad (8)$$

$$N_{t+1} = N_t - \eta \cdot \nabla_{x_l} (C \cdot g(x_l) + \|N\|_2^2) \quad (9)$$

Experimental Setup

We evaluated localized FGSM, PGD, and C&W attacks on the InceptionV3 model pretrained on ImageNet using 50 images from 9 classes. All images were resized to $299 \times 299 \times 3$ and preprocessed with TensorFlow’s standard Inception function.

Attack Success Criteria:

- *FGSM*: Confidence of the predicted class reduced by at least 50%.
- *PGD/C&W*: Misclassification within 250 iterations.

Attack Parameters:

- *FGSM*: $\epsilon = 0.05$
- *PGD*: $\epsilon = 0.02, \alpha = 0.01$
- *C&W*: Learning rate = 0.01, $\kappa = 1000, C = 10$

Masking Strategy: We define γ as the fraction of active pixels, with tested values: 1.0, 0.75, 0.5, and 0.25, corresponding to decreasing central region sizes.

Metrics

Attack Success Rate (ASR): Fraction of successful attacks over total attempts.

Number of Iterations: Steps required to achieve misclassification (max 250).

Attack	Mask (γ)	ASR (%)	Average Mean Change (%)	Average PSNR Change ($\% \times 10^{-3}$)	Average SSIM Change (%)	Average DR Change (%)	Average Iterations Change (%)
FSGM	1.00	0.76	0.00	0.00	0.00	0.00	-
	0.75	0.72	-43.37	0.57	4.00	0.00	-
	0.50	0.60	-74.47	1.06	7.00	0.00	-
	0.25	0.22	-93.71	1.50	8.66	0.00	-
PGD	1.00	1.00	0.00	0.00	0.00	0.00	0.00
	0.75	1.00	-20.71	-0.16	-0.02	41.05	65.61
	0.50	0.84	-42.63	-0.53	-0.01	83.26	392.09
	0.25	0.36	-70.33	0.40	0.12	288.64	2257.69
CW	1.00	0.82	0.00	0.00	0.00	0.00	0.00
	0.75	0.88	-16.22	-0.01	-0.00	11.88	6.33
	0.50	0.82	-41.88	-0.01	0.00	25.03	33.42
	0.25	0.72	-60.43	-0.03	0.01	114.09	589.27

Table 1: Evaluation of FGSM, PGD, and C&W attacks under varying mask sizes ($\gamma = 1.0, 0.75, 0.5, 0.25$). Metrics include ASR, perturbation magnitude (Mean Change), PSNR, SSIM, Dynamic Range (DR), and average iterations required for success.

Mean Pixel Value(MPV): Average absolute intensity of perturbation:

$$\text{MPV} = \frac{\sum |N|}{H \times W \times C}$$

Peak Signal to Noise Ratio(PSNR): Log-scaled ratio between maximum pixel intensity and noise energy:

$$\text{PSNR} = 10 \cdot \log_{10} \left(\frac{\text{MAX}^2}{N} \right)$$

Dynamic Range (DR): Range of values in the perturbation:

$$\text{DR} = \max(N) - \min(N)$$

SSIM: Structural similarity between original and perturbed images:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}$$

Here, μ_x, μ_y are the mean intensities, σ_x^2, σ_y^2 the variances, σ_{xy} the covariance, and C_1, C_2 small constants for numerical stability.

Results

In this section, we evaluate the proposed mask-based method for localized adversarial attacks. Experiments were conducted under varying masking constraints ($\gamma = 1.0, 0.75, 0.5, 0.25$) to assess effectiveness, imperceptibility, and computational efficiency.

Attack Success Rate (ASR)

ASR measures the percentage of adversarial examples that achieve the intended misclassification or confidence reduction. As shown in Table 1, ASR decreases as γ is reduced:

- **FGSM:** ASR drops from 76% at $\gamma = 1.0$ to 22% at $\gamma = 0.25$, reflecting reduced effectiveness with smaller perturbation regions.

- **PGD:** Achieves 100% ASR at $\gamma = 1.0$ and $\gamma = 0.75$, but falls to 36% at $\gamma = 0.25$.
- **C&W:** Maintains relatively high ASR (72%88%) across all mask sizes, indicating robustness to localization constraints.

Perturbation Imperceptibility

Imperceptibility was evaluated using Mean Pixel Value, PSNR, and SSIM. As γ decreases, fewer pixels are perturbed, resulting in subtler modifications:

- **Mean Pixel Value:** Drops significantly with smaller masks by 93.71% in FGSM, 70.33% in PGD, and 60.43% in C&W at $\gamma = 0.25$.
- **PSNR:** Increases for FGSM due to bounded noise, but decreases for PGD and C&W, indicating sharper perturbation peaks at smaller γ .
- **SSIM:** Higher SSIM is observed for FGSM at low γ ; PGD shows inconsistent trends; C&W maintains stable similarity across mask sizes.

Dynamic Range of Perturbations

As γ decreases, the dynamic range of the noise increases:

- **PGD:** +288.64% increase in range at $\gamma = 0.25$ compared to $\gamma = 1.0$.
- **C&W:** +114.09% increase under the same condition.

This suggests that more extreme values are used to compensate for the limited spatial region available.

Computational Efficiency

As γ is reduced, more iterations are required to achieve successful attacks:

- **PGD:** Requires 2257.69% more iterations at $\gamma = 0.25$ than at $\gamma = 1.0$.
- **C&W:** Iterations increase by 589.27% at $\gamma = 0.25$, indicating higher computational cost for localized optimization.

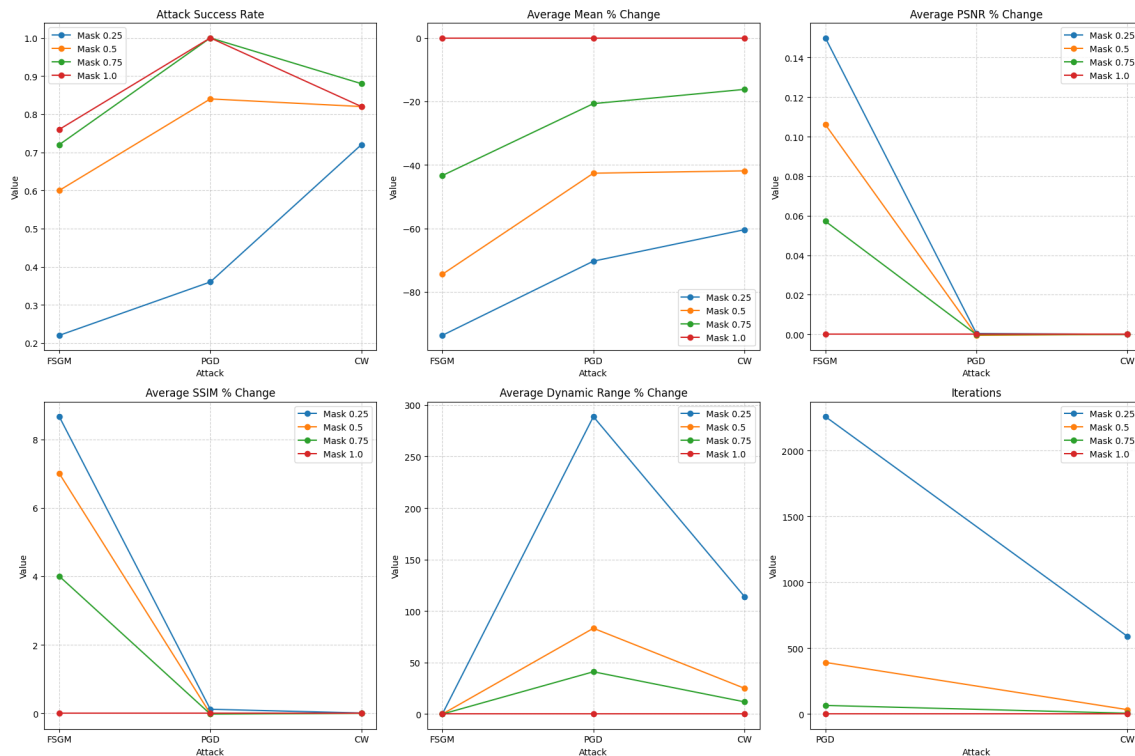


Figure 2: Evaluation of FGSM, PGD, and C&W attacks across mask sizes ($\gamma = 1.0, 0.75, 0.5, 0.25$) using six metrics: ASR, Mean Pixel Value, PSNR, SSIM, DR, and Iterations.

Attack Comparisons

- **FGSM:** Suffers from low ASR and high perceptibility with smaller masks.
- **PGD:** Highly effective at larger γ , but becomes less efficient and less successful as the mask size shrinks.
- **C&W:** Achieves the best balance, maintaining high ASR and good imperceptibility across all masking levels.

Figure 1 illustrates PGD adversarial examples. As γ decreases, perturbations become more concentrated and less visible, consistent with metric trends.

Discussion

This study shows that localized adversarial noise improves imperceptibility while preserving attack effectiveness, especially in iterative methods like PGD and C&W. Localized noise yields lower mean pixel values, higher PSNR, and better SSIM than global noise. However, smaller masks (γ) reduce ASR and increase computation. e.g., PGD needed 2257% more iterations at $\gamma = 0.25$ than at $\gamma = 1.0$.

FGSM performed poorly under localization, with ASR dropping from 76% to 22%, highlighting its limitations. PGD and C&W were more robust, with C&W maintaining stable ASR and visual quality across all γ .

Our results are based on a single model (InceptionV3) and dataset (ImageNet), limiting generalizability. Broader evaluation across architectures like ResNet, VGG, and ViT is

needed. Also, the impact of localized noise on existing defense methods remains unexplored and should be evaluated.

Future Work

Future work includes adaptive masking to better balance stealth, success, and efficiency (Tramer et al. 2020). Evaluating localized attacks on text, audio, and tabular data can expand applicability. Theoretical study of optimization behavior, e.g., gradient flow and decision boundaries may clarify why small γ reduces ASR. Localized noise may also support targeted manipulation and help evade detection (Li et al. 2021).

Conclusion

We evaluated localized adversarial noise in FGSM, PGD, and C&W attacks. Localized perturbations improve imperceptibility as measured by Mean Pixel Value, but reduce ASR and increase cost at smaller mask sizes.

FGSM degrades under localization, while PGD and especially C&W retain higher ASR and visual quality. These findings support iterative methods for localized attacks where stealth is critical.

Localized noise has uses in robustness testing, security evaluation, and detection evasion. Future work should explore dynamic masking, defense comparisons, cross-model validation, and deeper theoretical analysis to optimize trade-offs.

References

- Biggio, B., and Roli, F. 2017. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition* 84.
- Brown, T. B.; Man, D.; Roy, A.; Abadi, M.; and Gilmer, J. 2018. Adversarial patch.
- Carlini, N., and Wagner, D. 2017. Towards Evaluating the Robustness of Neural Networks . In *2017 IEEE Symposium on Security and Privacy (SP)*, 39–57. Los Alamitos, CA, USA: IEEE Computer Society.
- Chakraborty, A.; Alam, M.; Dey, V.; Chattopadhyay, A.; and Mukhopadhyay, D. 2021. A survey on adversarial attacks and defences. *CAAI Transactions on Intelligence Technology* 6(1):2545.
- Dia, O. A.; Karaletsos, T.; Hazirbas, C.; Ferrer, C. C.; Kabul, I. K.; and Meijer, E. 2021. Localized uncertainty attacks.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and harnessing adversarial examples.
- Karmon, D.; Zoran, D.; and Goldberg, Y. 2018. Lavan: Localized and visible adversarial noise.
- Li, F.; Liu, X.; Zhang, X.; Li, Q.; Sun, K.; and Li, K. 2021. Detecting localized adversarial examples: A generic approach using critical region analysis. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, 1–10.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2019. Towards deep learning models resistant to adversarial attacks.
- Ozbulak, U.; Peck, J.; Neve, W. D.; Goossens, B.; Saeys, Y.; and Messem, A. V. 2020. Regional image perturbation reduces l_p norms of adversarial examples while maintaining model-to-model transferability.
- Rane, N.; Mallick, S.; Kaya, .; and Rane, J. 2024. *Applications of deep learning in healthcare, finance, agriculture, retail, energy, manufacturing, and transportation: A review*. Deep Science Publishing. 132–152.
- Su, J.; Vargas, D. V.; and Sakurai, K. 2019. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation* 23(5):828841.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2014. Intriguing properties of neural networks.
- Tramer, F.; Carlini, N.; Brendel, W.; and Madry, A. 2020. On adaptive attacks to adversarial example defenses.
- Xiao, Q.; Li, K.; Zhang, D.; and Xu, W. 2018. Security risks in deep learning implementations. In *2018 IEEE Security and Privacy Workshops (SPW)*, 123–128.