# Enhancing the Transferability of Adversarial Attacks by Variance Tuning 

Xiaosen Wang and Kun He
School of Computer Science and Technology, Huazhong University of Science and Technology

## Introduction

Adversarial Examples are imperceptible from legitimate ones by adding tiny perturbations, but lead to incorrect model prediction.
Transferability: adversarial examples generated for one model can still fool other models, that enables black-box attacks in the realworld applications without any knowledge of target model.
Background: existing attacks (e.g. PGD, CW, etc.) have exhibited great effectiveness, but with low transferability.

## Methodology

Adversarial Example Generation
Standard Model Training

| Input image x | $\Longleftrightarrow$ Parameters |
| :---: | :---: |
| Transferability | $\Longleftrightarrow$ Generalization |

We treat the iterative gradient-based adversarial attack as a stochastic gradient decent (SGD) optimization process, in which at each iteration, the attacker always chooses the target model for update. SGD introduces variance due to randomness.
Definition 1 Gradient Variance. Given a classifier $f$ with parameters $\theta$ and loss function $J(x, y ; \theta)$, an arbitrary image $x \in \mathcal{X}$ and an upper bound $\epsilon^{\prime}$ for the neighborhood, the gradient variance can be defined as:

$$
V_{\epsilon^{\prime}}^{g}(x)=\mathbb{E}_{\left\|x^{\prime}-x\right\|_{p}<\epsilon^{\prime}}\left[\nabla_{x^{\prime}} J\left(x^{\prime}, y ; \theta\right)\right]-\nabla_{x} J(x, y ; \theta) .
$$

In practice, however, due to the continuity of the input space, we cannot calculate $\mathbb{E}_{\left\|x^{\prime}-x\right\|_{p}<\epsilon^{\prime}}\left[\nabla_{x^{\prime}} J\left(x^{\prime}, y ; \theta\right)\right]$ directly. Therefore, we approximate its value by sampling $N$ examples in the neighborhood of $x$ to calculate $V(x)$ :

$$
\begin{equation*}
V(x)=\frac{1}{N} \sum_{i=1}^{N} \nabla_{x^{i}} J\left(x^{i}, y ; \theta\right)-\nabla_{x} J(x, y ; \theta) . \tag{1}
\end{equation*}
$$



Figure 1: Relationship between various attacks.

## Algorithm

## Algorithm 1 VMI-FGSM

Input: A classifier $f$ with parameters $\theta$, loss function $J$. A raw ex ample $x$ with ground-truth label $y$. The magnitude of perturbation $\epsilon$; number of iteration $T$ and decay factor $\mu$. The factor $\beta$ for the upper bound of neighborhood and number of example $N$ for variance tuning.
Output: An adversarial example $x^{a d v}$
1: $\alpha=\epsilon / T$
2: $g_{0}=0 ; v_{0}=0 ; x_{0}^{a d v}=x$
3: for $t=0 \rightarrow T-1$ do
4: $\quad$ Calculate the gradient $\hat{g}_{t+1}=\nabla_{x_{t}^{a d v}} J\left(x_{t}^{a d v}, y ; \theta\right)$
5: Update $g_{t+1}$ by variance tuning based momentum

$$
\begin{equation*}
g_{t+1}=\mu \cdot g_{t}+\frac{\hat{g}_{t+1}+v_{t}}{\left\|\hat{g}_{t+1}+v_{t}\right\|_{1}} \tag{2}
\end{equation*}
$$

6: $\quad$ Update $v_{t+1}=V\left(x_{t}^{a d v}\right)$ by Eq. (1)
7: Update $x_{t+1}^{a d v}$ by applying the sign of gradient

$$
\begin{equation*}
x_{t+1}^{a d v}=x_{t}^{a d v}+\alpha \cdot \operatorname{sign}\left(g_{t+1}\right) \tag{3}
\end{equation*}
$$

8: end for
9: return $x^{a d v}=x_{T}^{a d v}$

| Model | Attack | Inc-v3 | Inc-v4 | IncRes-v2 | Res-101 | Inc-v3ens 3 | Inc-v3 ${ }_{\text {ens }} 4$ | IncRes-v2ens |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Inc-v3 | MI-FGSM | 100.0* | 43.6 | 42.4 | 35.7 | 13.1 | 12.8 | 6.2 |
|  | Vmi-FGSM | 100.0* | 71.7 | 68.1 | 60.2 | 32.8 | 31.2 | 17.5 |
|  | N-FGSM | 100.0* | 51.7 | 50.3 | 41.3 | 13.5 | 13.2 | 6.0 |
|  | VNI-FGSM | 100.0* | 76.5 | 74.9 | 66.0 | 35.0 | 32.8 | 18.8 |
| Inc-v4 | MI-FGSM | 56.3 | 99.7* | 46.6 | 41.0 | 16.3 | 14.8 | 7.5 |
|  | vmi-FGSM | 77.9 | 99.8* | 71.2 | 62.2 | 38.2 | 38.7 | 23.2 |
|  | N-FGSM | 63.1 | 100.0* | 51.8 | 45.8 | 15.4 | 13.6 | 6.7 |
|  | VNI-FGSM | 83.4 | 99.9* | 76.1 | 66.9 | 40.0 | 37.7 | 24.5 |
| IncRes-v2 | MI-FGSM | 60.7 | 51.1 | 97.9* | 46.8 | 21.2 | 16.0 | 11.9 |
|  | VmI-FGSM | 77.9 | 72.1 | 97.9* | 67.7 | 46.4 | 40.8 | 34.4 |
|  | N-FGSM | 62.8 | 54.7 | 99.1* | 46.0 | 20.0 | 15.1 | 9.6 |
|  | VNI-FGSM | 80.8 | 76.9 | $98.5^{*}$ | 69.8 | 47.9 | 40.3 | 34.2 |
| Res-101 | MI-FGSM | 58.1 | 51.6 | 50.5 | 99.3* | 23.9 | 21.5 | 12.7 |
|  | vmi-FGSM | 75.1 | 68.9 | 70.5 | 99.2* | 45.2 | 41.4 | 30.1 |
|  | N-FGSM | 65.6 | 58.3 | 57.0 | 99.4* | 24.5 | 21.4 | 11.7 |
|  | VNI-GGSM | 79.8 | 74.6 | 73.2 | 99.7* | 46.1 | 42.5 | 32.1 |

Table 1: Evaluations on gradient-based attacks.


## Experiments

| Model | Attack | Inc-v3 | Inc-v4 | IncRes-v2 | Res-101 | Inc-v3 ${ }_{\text {ens }} 3$ | Inc-v3 ${ }_{\text {ens } 4}$ | IncRes-v2ens |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Inc-v3 | MI-CT-FGSM | 98.7* | 85.4 | 80.6 | 76.0 | 64.1 | 62.1 | 45.2 |
|  | vmi-CT-FGSM | 99.3* | 88.6 | 86.7 | 82.9 | 78.6 | 76.2 | 64.7 |
|  | NI-CT-FGSM | 98.9* | 84.1 | 80.0 | 74.5 | 60.0 | 56.2 | 41.0 |
|  | vNI-CT-FGSM | 99.5* | 91.2 | 89.0 | 85.3 | 78.6 | 76.7 | 65.3 |
| Inc-v4 | MI-CT-FGSM | 87.2 | 98.6* | 83.3 | 78.3 | 72.2 | 67.2 | 57.3 |
|  | vmi-CT-FGSM | 90.0 | $98.8{ }^{*}$ | 86.6 | 81.9 | 78.3 | 76.6 | 68.3 |
|  | NI-CT-FGSM | 87.8 | 99.4** | 82.5 | 75.9 | 65.8 | ${ }^{62.6}$ | 51.3 |
|  | VNI-CT-FGSM | 92.1 | 99.2* | 89.2 | 85.1 | 80.1 | 78.3 | 70.4 |
| IncRes-v2 | MI-CT-FGSM | 87.9 | 85.7 | $97.1{ }^{*}$ | 83.0 | 77.6 | 74.6 | 72.0 |
|  | vmi-CT-FGSM | 88.9 | 87.0 | 97.0* | 85.0 | 83.4 | 80.5 | 79.4 |
|  | NI-CT-FGSM | 90.2 | 87.0 | 99.4* | 83.2 | 75.0 | 68.9 | 65.1 |
|  | VNI-CT-FGSM | 92.9 | 90.6 | 99.0* | 88.2 | 85.2 | 82.5 | 81.8 |
| Res-101 | MI-CT-FGSM | 86.5 | 81.8 | 83.2 | 98.9** | 77.0 | 72.3 | 61.9 |
|  | VMI-CT-FGSM | 86.9 | 84.2 | 86.4 | 98.6* | 81.0 | 78.6 | 71.6 |
|  | NI-CT-FGSM | 86.1 | 82.2 | 83.3 | 98.5* | 70.0 | 68.5 | 54.6 |
|  | VNI-CT-FGSM | 90.7 | 85.5 | 87.2 | 99.1* | 82.6 | 79.7 | 73.3 |

Table 2: Evaluations with the combination of DIM, TIM and SIM.

| Attack | Inc-v3 | Inc-v4 | IncRes-v2 | Res-101 | Inc-v3 ${ }_{\text {ens }} 3$ | Inc-v3 ${ }_{\text {ens } 4}$ | IncRes-v2ens |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| MI-FGSM | 99.9* | 98.2* | $95.3^{*}$ | 99.9* | 39.4 | 35.3 | 24.2 |
| VMI-FGSM | 99.7* | 98.5* | 96.0* | 99.9* | 67.6 | 62.9 | 50.7 |
| NI-FGSM | 99.8* | 99.8* | 98.9* | 99.8* | 41.0 | 33.5 | 23.1 |
| VNI-FGSM | 99.9* | 99.6* | 98.6* | 99.9* | 71.3 | 66.0 | 52.9 |
| MI-CT-FGSM | 99.6* | 99.1* | 97.4* | 99.7* | 91.3 | 89.6 | 86.8 |
| VMI-CT-FGSM | 99.7* | 99.2* | 98.4* | 99.9* | 93.6 | 92.4 | 91.0 |
| NI-CT-FGSM | 100.0* | 100.0* | 100.0* | 100.0* | 92.8 | 89.6 | 83.6 |
| VNI-CT-FGSM | 100.0* | 99.9* | 99.6* | 100.0* | 95.5 | 94.5 | 92.3 | Table 3: Evaluations in multi-model setting.



Table 4: Evaltions on Defense models

(a) VMI-FGSM

(b) VNI-FGSM Figure 3: Ablation

## Conclusion

- We propose a variance tuning method to enhance the transferability of the iterative gradient-based attacks.
- Our method is geneally applicable to any iterative gradient based attacks and input transformations (i.e. DIM, TIM, SIM etc.)
- Experiments show our method could significantly enhance the transferability of various attacks.
- The results indicate the insufficiency of existing defenses and can serve as a benchamrk to evalute the robustness of future developed defense.

