

Advancing Generalized Transfer Attack with Initialization Derived Bilevel Optimization and Dynamic Sequence Truncation

Yaohua Liu¹, Jiaxin Gao¹, Xuan Liu², Xianghao Jiao¹, Xin Fan¹ and Risheng Liu^{1,3*}

¹School of Software Engineering, Dalian University of Technology, China

²CRRC Qingdao Sifang Rolling Stock Research Institute Co., Ltd, China

³Pazhou Laboratory (Huangpu), Guangzhou, China.

liuyaohua_918@163.com, jiaxinn.gao@outlook.com, liuxuan_16@126.com, jiaoxh0331@outlook.com, {xin.fan, rsliu}@dlut.edu.cn

Abstract

Transfer attacks generate significant interest for real-world black-box applications by crafting transferable adversarial examples through surrogate models. Whereas, existing works essentially directly optimize the single-level objective w.r.t. the surrogate model, which always leads to poor interpretability of attack mechanism and limited generalization performance over unknown victim models. In this work, we propose the **BilEvel Transfer AttacK** (BETAK) framework by establishing an initialization derived bilevel optimization paradigm, which explicitly reformulates the nested constraint relationship between the Upper-Level (UL) pseudo-victim attacker and the Lower-Level (LL) surrogate attacker. Algorithmically, we introduce the Hyper Gradient Response (HGR) estimation as an effective feedback for the transferability over pseudo-victim attackers, and propose the Dynamic Sequence Truncation (DST) technique to dynamically adjust the back-propagation path for HGR and reduce computational overhead simultaneously. Meanwhile, we conduct detailed algorithmic analysis and provide convergence guarantee to support non-convexity of the LL surrogate attacker. Extensive evaluations demonstrate substantial improvement of BETAK (e.g., **53.41%** increase of attack success rates against IncRes-v2_{ens}) against different victims and defense methods in targeted and untargeted attack scenarios.

1 Introduction

Adversarial attack introduces imperceptible yet artificial perturbations into the clean image, potentially leading to incorrect outputs even serious security concerns [Goodfellow *et al.*, 2014; Xie *et al.*, 2017b; Jiao *et al.*, 2023]. Among these black-box attack techniques, transfer attacks [Dong *et al.*, 2018; Xie *et al.*, 2019] generates Adversarial Examples (AE) by attacking a surrogate model, which leads to better attack performance against these victim models with unknown network structures and gradient information. Due to its ability

to initiate attacks without requiring direct access to the victim models, transfer attack has garnered extensive attention and research in various real-world applications such as autonomous driving [Deng *et al.*, 2020] and surveillance systems [Alparslan *et al.*, 2020].

Building upon this foundation, an extensive array of transfer attack methods has been studied, principally including strategies such as input transformation [Wang *et al.*, 2021; Long *et al.*, 2022], momentum-based [Lin *et al.*, 2019; Wang and He, 2021; Dong *et al.*, 2018], objective-based [Huang *et al.*, 2019] methods, and modification of network structural properties [Wang *et al.*, 2023; Guo *et al.*, 2020]. The above transfer attack methods essentially optimize a single-variable single-level objective w.r.t. surrogate model, thereby overlooking the explicit formulation and evaluation for the attack performance w.r.t. victim models. This deficiency has significantly limited the generalization performance of the AEs across varying victim models. Meanwhile, several studies [Yuan *et al.*, 2021; Fang *et al.*, 2022] have also focused on designing meta-learning based methods to improve the attack transferability of AEs with model ensembles or data augmentation techniques. Whereas, they always require multiple sampling iterations from the white-box model zoo and pose challenges in integrating with various types of transfer attack methods. Meanwhile, there is a notable lack of interpretability of attack mechanism along with comprehensive theoretical analysis for these methods.

In this work, we establish the **BilEvel Transfer AttacK** (BETAK) framework to address the above limitations. Specifically, by revisiting the optimization process of transfer attack based on its single-level objective of surrogate model, we propose a new initialization derived bilevel optimization paradigm to explicitly reformulate the nested constraint relationship between the Upper-Level (UL) pseudo-victim attacker (w.r.t. perturbation initialization variable) and the Lower-Level (LL) surrogate attacker (w.r.t. perturbation variable). The LL subproblem corresponds to the same attack objective as the single-level transfer attack, while the UL subproblem essentially reflects the ultimate objective function to achieve better generalization performance over unknown victim models. As for the solution strategy, we first explicitly calculate the Hyper-Gradient Response (HGR) as an effective feedback of transferability against one or more pseudo-victim models. Then we further introduce the Dynamic Sequence

*Corresponding author

Truncation (DST) technique to dynamically adjust the back-propagation path for HGR, which also effectively relieves the computation burden. Besides, we provide convergence analysis of the BETAK framework where the LL subproblem has non-convex property due to complex network structures and objectives. Finally, we not only verify the attack performance of BETAK against various victim models and defense methods, but also consider both targeted and untargeted scenarios for thorough evaluation. Comprehensive ablation studies are also conducted to analyze the effectiveness of BETAK and DST technique. The main contributions are summarized as follows:

- Different from directly optimizing a single-variable single-level surrogate objective, we propose the **BilEvel Transfer AttacK** (BETAK) framework based on the established initialization-derived bilevel optimization paradigm, to analyze and explicitly reformulate the nested constraint relationship between the UL pseudo-victim attacker and the LL surrogate attacker.
- Algorithmically, we propose the Hyper-Gradient Response (HGR) estimation as an effective feedback of the transferability against the pseudo-victim models. Meanwhile, we introduce the Dynamic Sequence Truncation (DST) technique to dynamically adjust the back-propagation pathway for HGR computation, simultaneously alleviating the computational burden.
- We provide detailed algorithmic analysis of BETAK with convergence guarantees to support non-convex property of surrogate attackers. This demonstrates its compatibility with a broader spectrum of surrogate models, characterized by intricate non-convex network architectures and objective functions.
- In-depth experimental analysis reveals that the BETAK framework exhibits substantial improvements against various victim models and defense methods (i.e., 53.41% increase of attack success rates against the IncRes-v2_{ens} model). The ablation study also validates the effectiveness of DST technique to enhance performance and relieve computational burden.

2 Methodology

We first briefly revisit the single-level transfer attack process w.r.t. the surrogate attacker, and then we introduce the initialization derived bilevel optimization paradigm and propose two core techniques of BETAK, including the HGR estimation and DST operation.

2.1 Preliminary of Transfer Attack

We first review the single-level problem definition of transfer attacks to evaluate the transferability of AEs. We denote the dataset as $\mathcal{D} = \{\mathbf{u}_i, \mathbf{v}_i\}_{i=1}^M$, where $\mathcal{D}_i = (\mathbf{u}_i, \mathbf{v}_i)$ represents the image and paired label. Then the fundamental objective function of transfer attack could be described as

$$\max_{\delta \in \mathcal{C}} \mathcal{L}_{\text{sur}}(\delta; \mathcal{D}_i, \mathcal{R}), \quad (1)$$

where $\mathcal{L}_{\text{sur}}(\delta; \mathcal{D}_i, \mathcal{R}) = \mathcal{L}_{\text{sur}}(\mathcal{R}(\mathbf{u}_i + \delta), \mathbf{v}_i)$, \mathcal{R} denotes the surrogate model. δ is the perturbation variable optimized

based on various attackers subject to $\mathcal{C} = \{\delta \mid \|\delta\|_{\rho} \leq \epsilon\}$, where \mathcal{C} represents the range of values for the perturbations generated by the attack, and ρ denotes the norm used for attack generation. Typically speaking, we use δ_k to denote the generated perturbation at k -th step maximization:

$$\delta_{k+1} \leftarrow \Pi_{\mathcal{C}}(\delta_k + \alpha \cdot \text{sgn} \nabla_{\delta} \mathcal{L}_{\text{sur}}(\delta_k; \mathcal{D}_i, \mathcal{R})), \quad (2)$$

where $k = 0, 1, \dots, K-1$, and Π is the projection operation. To this end, various gradient based attackers (single-step [Goodfellow *et al.*, 2014] or multi-step [Dong *et al.*, 2018]) could be combined to improve attack performance.

Therefore, existing methods have explored various techniques to facilitate the maximization optimization in Eq. (1) w.r.t. δ to enhance the transferability of AEs to unknown victim models, denoted as \mathcal{V} . In this case, the ultimate objective function of transfer attacks should be written as

$$\max_{\delta_K \in \mathcal{C}_{\text{vic}}} \frac{1}{N} \sum_{n=1}^N \mathcal{L}_{\text{vic}_n}(\delta_K; \mathcal{D}_i, \mathcal{V}_n), \quad (3)$$

where $\mathcal{L}_{\text{vic}_n}$ represents the objective of n -th victim models, i.e., \mathcal{V}_n , $n = 1, \dots, N$. \mathcal{C}_{vic} denotes the constraint for the victim model. Typically, \mathcal{C}_{vic} is set as the same as \mathcal{C} for simplicity. Note that this objective reflects the essential goal of transfer attacks, i.e., enhancing the generalization performance of δ_K when transferred to black-box victim models.

2.2 Initialization Derived Bilevel Paradigm

We first introduce $\phi(\delta)$ to denote the perturbation generated by certain types of attackers with δ as the perturbation initialization. When we remove the sgn operation [Cheng *et al.*, 2021], and take single-step gradient ascent as an example, $\phi(\delta)$ can be written as

$$\phi(\delta) = \Pi_{\mathcal{C}}(\delta + \alpha \cdot \nabla_{\delta} \mathcal{L}_{\text{sur}}(\delta; \mathcal{D}_i, \mathcal{R})). \quad (4)$$

Then we construct the initialization derived bilevel optimization paradigm to capture the nested constrained relationship between Eq. (1) and Eq. (3), which can be written as:

$$\begin{aligned} \min_{\delta \in \mathcal{C}} F(\phi(\delta)), \text{ where } F &:= -\frac{1}{N} \sum_{n=1}^N \mathcal{L}_{\text{vic}_n}(\phi^*(\delta); \mathcal{D}_i, \mathcal{V}_n) \\ \text{s.t.}, \phi^*(\delta) &:= \underset{\phi \in \mathcal{C}}{\text{argmin}} f(\phi(\delta)), \text{ where } f := -\mathcal{L}_{\text{sur}}(\phi(\delta); \mathcal{D}_i, \mathcal{R}). \end{aligned} \quad (5)$$

$\phi^*(\delta)$ denotes the ‘best’ perturbation derived by iteratively solving Eq. (4). In the following, the models utilized to construct the UL objective will be referred to as the pseudo-victim attacker. Accordingly, the LL surrogate attacker aims to generate the ‘best’ adversarial perturbation to attack the surrogate model, which shares similar motivation with Eq. (1). While the UL pseudo-victim attacker is optimized to achieve better generalization performance with the perturbation derived from δ .

Bilevel optimization [Liu *et al.*, 2023a; Liu *et al.*, 2024] is well recognized for modeling and solving various hierarchical applications problems, such as hyperparameter optimization [Franceschi *et al.*, 2017; Liu *et al.*, 2021a] and meta learning [Hospedales *et al.*, 2020; Liu and Liu, 2021]. The

proposed initialization derived bilevel paradigm clearly illustrates the constrained relationship between the UL pseudo-victim attacker and LL surrogate attacker, and offers a new perspective to generate AEs with better generalization capability when transferred to unknown victim models.

2.3 Hyper-Gradient Response (HGR) Estimation

As shown in Eq. (5), the UL objective continuously optimizes δ , which is defined as the initial state of the LL variable ϕ . Then we introduce the HGR estimation to update δ in Eq. (5). Specifically, we first construct the iterative dynamical system by extending Eq. (4) as

$$\begin{cases} \phi_0(\delta) = \delta, \\ \phi_{k+1}(\delta) = \Pi_C(\phi_k(\delta) - \alpha \cdot \nabla_{\phi} f(\phi_k(\delta))), \end{cases} \quad (6)$$

where $k = 1, \dots, K$. Therefore, $\phi^*(\delta)$ could be approximated by

$$\phi^*(\delta) \approx \phi_K(\delta) = \phi_0(\delta) - \alpha \cdot \sum_{k=1}^K \nabla_{\phi} f(\phi_{k-1}(\delta)).$$

Following the chain rule, the hyper-gradient w.r.t. δ can be explicitly calculated with auto differentiation as follows

$$\nabla_{\delta} F(\phi_K(\delta)) = (\nabla_{\delta} \phi_K(\delta))^{\top} \nabla_{\phi} F(\phi_K(\delta)). \quad (7)$$

The hyper-gradient contains second-order gradient information as effective feedback of the transferability of $\phi_K(\delta)$. Then we adopt the estimated hyper-gradient to update δ by

$$\delta = \Pi_C(\delta - \beta \cdot \text{sgn} \nabla_{\delta} F(\phi_K(\delta))),$$

where β denotes the learning rate. As the UL and LL objectives converge, $\phi^*(\delta)$ is supposed to obtain better generalization performance on the victim models.

In addition, it should be emphasized that we have no access to the gradient information about unknown victim models in the scenario of transfer attacks. To define the UL victim attacker in Eq. (5), we assume that one or more arbitrary known models, in addition to \mathcal{R} , could serve as pseudo-victim models to imitate the victim models during transfer process. These pseudo-victim models, denoted as \mathcal{V}_n , inherently satisfy the black-box attack nature for evaluating adversarial transferability of ϕ_K and update its initialization, i.e., δ , which finally leads to ϕ^* with better generalization performance. We also conduct extensive experiments to investigate the influence of pseudo-victim model selection on transferability of AEs.

From another perspective, existing transfer attack methods always obtain ϕ_K with empirically chosen or randomly initialized δ in Eq. (6), while we estimate the HGR to optimize the initialization used for attacking \mathcal{R} , which finally leads to better generalization performance over unknown victim models. We provide more details to discuss the influence of alternative victim models in the ablation part.

2.4 Dynamical Sequence Truncation (DST)

We further introduce the DST operation to facilitate the HGR estimation, which improves the theoretical properties and relieve the hyper-gradient computation burden simultaneously.

Algorithm 1 BETAK Framework

Require: UL iteration T , LL attack iteration K , perturbation δ , LL learning rate α and UL learning rate β .

```

1: Initialize  $\delta^0$ .
2: for  $t = 0 \rightarrow T - 1$  do
3:    $\phi_0(\delta) = \delta^t$ .
4:   for  $k = 0 \rightarrow K - 1$  do
5:     % Iterative attack with  $\phi_k(\delta)$ .
6:      $\phi_{k+1}(\delta) = \Pi_C(\phi_k(\delta) - \alpha \cdot \nabla_{\phi} f(\phi_k(\delta)))$ .
7:   end for
8:   % Calculate  $\tilde{K}$  with DST technique.
9:    $\tilde{K} = \arg \max_{1 \leq k \leq K} \{F(\phi_k(\delta))\}$ .
10:  % Calculate the HGR estimation.
11:   $\nabla_{\delta} F(\phi_{\tilde{K}}(\delta)) = (\nabla_{\delta} \phi_{\tilde{K}}(\delta))^{\top} \nabla_{\phi} F(\phi_{\tilde{K}}(\delta))$ .
12:  % Update  $\delta$  with  $\phi_{\tilde{K}}$ .
13:   $\delta^{t+1} = \Pi_C(\delta^t - \beta \cdot \text{sgn} \nabla_{\delta} F(\phi_{\tilde{K}}(\delta)))$ .
14: end for
    
```

In Eq. (7), the whole historical sequence of ϕ , i.e., $\{\phi_k\}_{k=1}^K$, is used to estimate HGR and find better perturbation initialization. Whereas, since the LL subproblem is mostly non-convex due to the complex network structure of surrogate models and varying objectives, there is no guarantee that the obtained $\{\phi_k\}_{k=1}^K$ could uniformly lead to the convergence of both subproblems. Meanwhile, although K is crucial for proper approximation of $\phi^*(\delta)$, it is always impractical for us to empirically choose ideal fixed K for HGR estimation.

In this paper, we propose to dynamically truncate the historical sequence in Eq. (6) according to the change of the UL objective. We first introduce the variable \tilde{K} to replace fixed K for HGR, calculated by

$$\tilde{K} := \arg \max_{1 \leq k \leq K} \{F(\phi_k(\delta))\}. \quad (8)$$

Theoretically, \tilde{K} is chosen to minimize the worst case of $\phi_k(\delta)$ in terms of $F(\phi_k(\delta))$ so as to ensure that the sequence of solutions for f is consistently optimized with F . This idea originates from the non-convex first-order optimization methods [Liu *et al.*, 2020; Liu *et al.*, 2021b; Liu *et al.*, 2023b], and we extend this technique to relax the theoretical properties of BETAK under the LL non-convex scenario. Therefore, $\phi^*(\delta)$ is further approximated by $\phi_{\tilde{K}}$, and HGR can be estimated as

$$\nabla_{\delta} F(\phi_{\tilde{K}}(\delta)) = (\nabla_{\delta} \phi_{\tilde{K}}(\delta))^{\top} \nabla_{\phi} F(\phi_{\tilde{K}}(\delta)). \quad (9)$$

It should be noted that since \tilde{K} is always selected from the range of $[1, K]$, the DST technique naturally saves the computation cost for calculating the hyper gradient in Eq. (7). We summarize the BETAK framework with HGR and DST techniques in Alg. 1. By combining HGR estimation with the DST operation, we further establish the convergence guarantee of BETAK framework, which effectively handles the non-convex property of LL surrogate attacker caused by complex network structure and diverse loss functions.

2.5 Algorithmic Analysis of BETAK

We delve into the theoretical foundations of BETAK, focusing particularly on the convergence properties in non-convex scenarios. Following previous assumptions [Liu *et al.*,

Attacker	Method	Inc-v3*	IncRes-v2*	DenseNet	MobileNet	PNASNet	SENet	Inc-v3 _{ens3}	Inc-v3 _{ens4}	IncRes-v2 _{ens}
PGD	N/A	16.34	13.38	36.86	36.12	13.46	17.14	10.24	9.46	5.52
	SGM	23.68	19.82	51.66	55.44	22.12	30.34	13.78	12.38	7.90
	LinBP	27.22	23.04	59.34	59.74	22.68	33.72	16.24	13.58	7.88
	Ghost	17.74	13.68	42.36	41.06	13.92	19.10	11.60	10.34	6.04
	BPA	<u>35.36</u>	<u>30.12</u>	<u>70.70</u>	<u>68.90</u>	<u>32.52</u>	<u>42.02</u>	<u>22.72</u>	<u>19.28</u>	<u>12.40</u>
BETAK (Ours)		53.34	45.08	74.68	74.54	42.48	49.94	30.80	24.90	17.60
MI-FGSM	N/A	26.20	21.50	51.50	49.68	22.92	30.12	16.22	14.58	9.00
	SGM	33.78	28.84	63.06	65.84	31.90	41.54	19.56	17.48	10.98
	LinBP	35.92	29.82	68.66	69.72	30.24	41.68	19.98	16.58	9.94
	Ghost	29.76	23.68	57.28	56.10	25.00	34.76	17.10	14.76	9.50
	BPA	<u>47.58</u>	<u>41.22</u>	<u>80.54</u>	<u>79.40</u>	<u>44.70</u>	<u>54.28</u>	<u>32.06</u>	<u>25.98</u>	<u>17.46</u>
BETAK (Ours)		58.24	47.30	81.82	80.50	49.26	55.76	35.82	31.04	22.68
VMI-FGSM	N/A	42.68	36.86	68.82	66.68	40.78	46.34	27.36	24.20	17.18
	SGM	50.04	44.28	77.56	79.34	48.58	56.86	32.22	27.72	19.66
	LinBP	47.70	40.40	77.44	78.76	41.48	52.10	28.58	24.06	16.60
	Ghost	47.82	41.42	75.98	73.40	44.84	52.78	30.84	27.18	19.08
	BPA	<u>55.00</u>	<u>48.72</u>	<u>85.44</u>	<u>83.64</u>	<u>52.02</u>	<u>60.88</u>	<u>38.76</u>	<u>33.70</u>	<u>23.78</u>
BETAK (Ours)		61.18	52.18	85.54	84.64	54.18	61.44	40.84	34.3	26.16
ILA	N/A	29.10	26.08	58.02	59.10	27.60	39.16	15.12	12.30	7.86
	SGM	35.64	32.34	65.20	71.22	34.20	46.72	17.10	13.86	9.08
	LinBP	37.36	34.24	71.98	72.84	35.12	48.80	19.38	14.10	9.28
	Ghost	30.06	26.50	60.52	61.74	28.68	40.46	14.84	12.54	7.90
	BPA	<u>47.62</u>	<u>43.50</u>	<u>81.74</u>	<u>80.88</u>	<u>47.88</u>	<u>60.64</u>	<u>27.94</u>	<u>20.64</u>	<u>14.76</u>
BETAK (Ours)		57.14	50.58	83.28	82.26	51.88	62.06	31.48	23.12	17.04
SSA	N/A	35.78	29.58	60.46	64.70	25.66	34.18	20.64	17.30	11.44
	SGM	45.22	38.98	70.22	78.44	35.30	46.06	26.28	21.64	14.50
	LinBP	48.48	41.90	75.02	78.30	36.66	49.58	28.76	23.64	15.46
	Ghost	36.44	28.62	61.12	66.80	24.90	33.98	20.58	16.84	10.82
	BPA	<u>51.36</u>	<u>44.70</u>	<u>76.24</u>	<u>79.66</u>	<u>39.38</u>	<u>50.00</u>	<u>32.10</u>	<u>26.44</u>	<u>18.20</u>
BETAK (Ours)		64.38	53.14	83.24	82.92	54.24	59.92	41.52	36.08	27.92

Table 1: We compare the ATR (%) results based on 5 iterative attackers (i.e., PGD, MI-FGSM, VMI-FGSM, ILA and SSA) incorporated with 4 model based methods (i.e., SGM, LinBP, Ghost and BPA). The AEs are generated with ResNet-50 backbones and tested on 2 pseudo-victim (marked with *) and 7 victim models. The best and second-best outcomes are designated with boldface and underline, respectively.

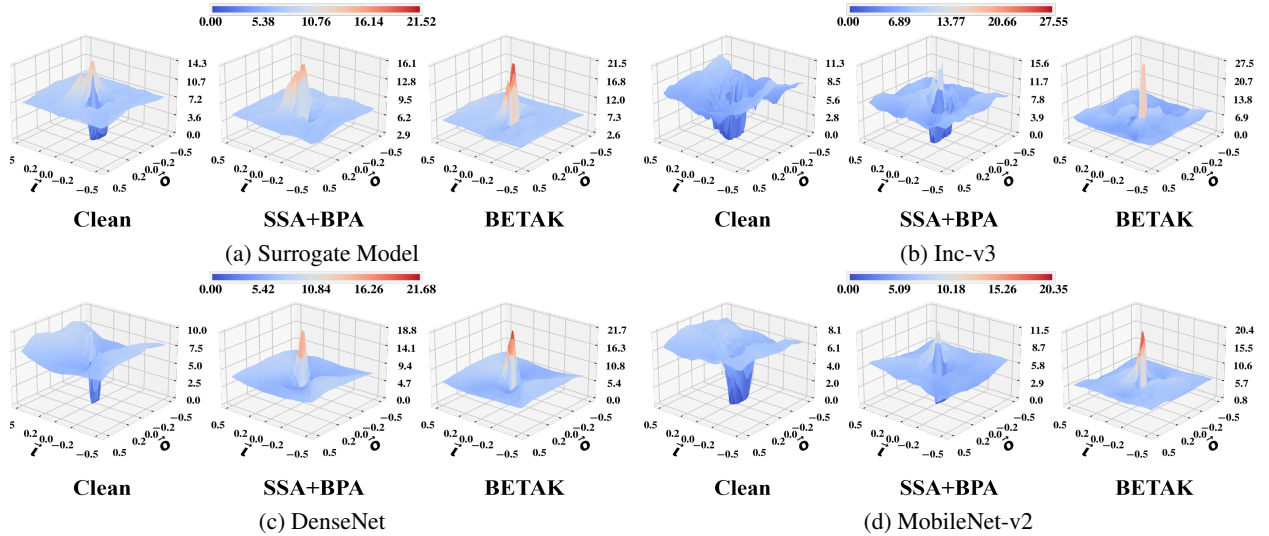


Figure 1: The three subfigures in group (a)-(d) illustrate the adversarial loss landscape w.r.t. input variations of the clean sample in Fig. 2 and corresponding AEs perturbed by SSA+BPA and BETAK. The color bars represent a common loss range and color mapping in each group. We mark the maximum and minimum value within the range of $x, y \in [-0.5, 0.5]$ w.r.t. \tilde{x} and \tilde{y} , respectively.

2021b], when the UL variables only contain the initialization variable of ϕ_k , i.e., δ , without including additional parameters, we have that the initialization-based bilevel optimization paradigm, depicted in Eq. (5), is actually a special case of the explicit gradient-based methods [Liu *et al.*, 2023b]. Consequently, with the LL non-convex property, by introducing \tilde{K} to dynamically truncate the historical sequence (i.e., $\{F(\phi_k(\delta))\}_{k=1}^K$) obtained from the gradient based attack iterations, we demonstrate that any limit point of the UL sequence is guaranteed to converge to a stationary point satisfying the first-order optimality conditions of Eq. (5). This above convergence analysis ensures that despite the non-convex nature of LL surrogate attackers, BETAK can reliably find solutions that are robust and practically effective to improve the transferability across unknown victim models. When we derive the optimized perturbation initialization, BETAK could flexibly combined various attack methods to perform the surrogate attack process and obtain the optimal perturbation. In the experimental part, we demonstrate the flexibility of BETAK combined with different basic iterative attackers.

3 Experiments

We evaluate the performance improvement of BETAK compared with representative transfer attack methods against various victim models and defense techniques under both targeted and untargeted attacks.

3.1 Experimental Setting

Dataset and Models. We follow the commonly used settings to randomly select 5000 images from 1000 categories of the ImageNet (ISLVR2012) validation set, all of which could be correctly classified by the victim models. For fair comparison, we choose the commonly used ResNet-50 backbone as the surrogate model. We select 9 victim models with different network structures, including 6 standardly trained models, i.e., Inception (Inc)-v3 [Szegedy *et al.*, 2016], IncRes-v2 [Szegedy *et al.*, 2017], DenseNet [Huang *et al.*, 2017], MobileNet-v2 [Sandler *et al.*, 2018], PNASNet [Liu *et al.*, 2018], SENet [Hu *et al.*, 2018] and 3 robustly trained models, i.e., Inc-v3_{ens3}, In-v3_{ens4} and IncRes-v2_{ens} [Tramèr *et al.*, 2017]. As for the implementation of BETAK, we consider Inception-v3 and IncRes-v2 as the pseudo-victim model since they are proposed earlier. Besides, it can be clearly observed from the ablation results that introducing either one as the pseudo model will significantly improve the attack performance against other victim models. We adopt the Attack Success Rate (ATR) as the evaluation criterion for the performance of transfer attacks.

Baselines and Training Details. We choose 5 representative attackers including the basic multi-step attacker (i.e., PGD [Kurakin *et al.*, 2018b]), 2 momentum based attackers (i.e., MI-FGSM [Dong *et al.*, 2018] and VMI-FGSM [Wang and He, 2021]), objective based attacker (i.e., ILA [Huang *et al.*, 2019]) and input transformation based attacker (i.e., SSA [Long *et al.*, 2022]). Besides, we also consider various model-related methods including SGM [Wu *et al.*, 2020], LinBP [Guo *et al.*, 2020], Ghost [Li *et al.*, 2020] and BPA [Wang *et al.*, 2023]. We follow the hyper parameter

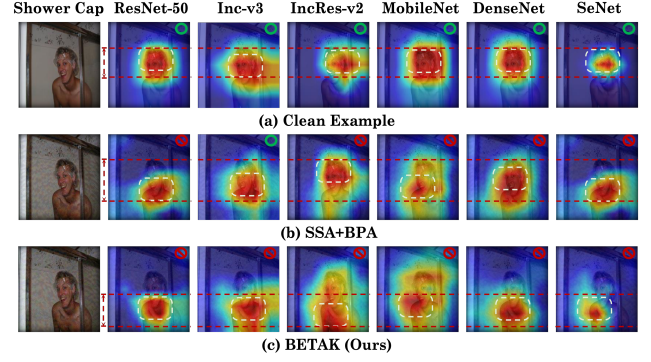


Figure 2: The CAM results of the clean image (labeled as Shower Cap) in Fig. 1 and corresponding AEs generated with ResNet-50 by SSA+BPA and the proposed BETAK. Green and red marks represent that the image is correctly and wrongly classified, respectively.

setting in BPA to implement the above methods. In the untargeted scenario, we set the UL iteration $T = 10$, $\epsilon = 8/255$ and UL learning rate $\beta = 1.6/255$ for all the methods. As for the targeted attacks, we set $T = 300$, $\epsilon = 16/255$ and $\beta = 16/255$. As for the hyper parameter of BETAK, since we omit the sgn operation, we set the LL learning rate $\alpha = 2.0$ and LL iteration $K = 10$.

3.2 Experimental Results

Quantitative Results with Untargeted Attacks. We first evaluate the attack performance of BETAK against various victim models in comparison with state-of-the-art attack methods. In Tab. 1, we test the ATR with 2 pseudo-victim models and 7 victim models. As it is expected, BETAK improves the ATR on Inception-v3 and IncRes-v2, since these pseudo-victim models provide the feedback of generalization performance by participating in the HGR estimation w.r.t. δ . More importantly, it is shown that BETAK also significantly improves the attack performance against other 7 victim models. Therefore, by continuously optimizing the pseudo-victim attacker w.r.t. δ , the generated AEs are supposed to have better generalization performance to unknown victim models.

Evaluation with Defense Strategies. We implement 5 representative defense methods for evaluation, including HGD [Liao *et al.*, 2018], R&P [Xie *et al.*, 2017a], NIPS-r3 [Kurakin *et al.*, 2018a], JPEG [Guo *et al.*, 2017] and RS [Cohen *et al.*, 2019]. We compare the ATR of BETAK and these methods against different defense strategies based on the PGD and MI-FGSM in Fig. 3. It can be observed that BETAK generates AEs that have consistently stronger attack performance against different defense techniques.

Visualization of Adversarial Loss Landscapes. In Fig. 1, we also plot the adversarial loss landscapes of different victim models for the same clean image and two adversarial images generated by SSA+BPA and BETAK in Fig. 2. The adversarial loss is calculated with $\mathcal{L}_{\text{vic}}(\mathbf{u} + x\vec{\tau} + y\vec{\sigma})$, where \mathbf{u} denotes the original clean image or adversarial image, $\vec{\tau} = \text{sgn}(\nabla_{\mathbf{u}} \mathcal{L}_{\text{atk}}(\mathbf{u}))$ and $\vec{\sigma} \sim \text{Rademacher}(0, 0.5)$ are the sign gradient direction and random direction (x and y are the corresponding linear coefficients). To be general, when the original adversarial image exhibits stronger attack perfor-

Attacker	Method	Inc-v3*	IncRes-v2*	DenseNet	MobileNet	PNASNet	SENet	Inc-v3 _{ens3}	Inc-v3 _{ens4}	IncRes-v2 _{ens}
PGD	N/A	0.54	0.80	4.48	2.04	1.62	2.26	0.18	0.08	0.02
	SGM	2.56	3.12	15.08	8.68	5.78	9.84	0.62	0.18	0.04
	LinBP	5.30	4.84	16.08	8.48	7.26	7.94	1.50	0.54	0.28
	Ghost	1.34	2.14	10.24	4.74	3.90	6.64	0.36	0.16	0.10
	BPA	<u>6.36</u>	<u>7.80</u>	<u>23.30</u>	<u>12.26</u>	<u>12.54</u>	<u>11.86</u>	<u>1.72</u>	<u>0.82</u>	<u>0.50</u>
BETAK (Ours)		7.66	9.14	24.90	13.08	13.78	12.62	2.08	0.84	0.70
MI-FGSM	N/A	0.16	0.26	2.06	0.90	0.42	1.22	0.00	0.02	0.02
	SGM	0.74	0.76	5.84	3.24	1.66	3.70	0.00	0.02	0.00
	LinBP	3.30	3.00	13.44	6.26	5.50	7.18	0.30	0.10	0.02
	Ghost	0.66	0.76	5.48	2.14	1.58	3.38	0.08	0.02	0.00
	BPA	<u>4.26</u>	<u>5.02</u>	<u>18.54</u>	<u>8.98</u>	<u>8.40</u>	<u>10.18</u>	<u>0.44</u>	<u>0.16</u>	<u>0.06</u>
BETAK (Ours)		6.62	6.86	20.06	10.60	9.52	10.80	0.76	0.26	0.08

Table 2: We report comparative results of targeted ATR (%) based on 9 victim models. All the AEs are generated with ResNet-50 backbones. The best and second-best outcomes are designated with boldface and underline, respectively.

Attacker	Method	HGD	R&P	NIPS-r3	JPEG	RS
PGD	N/A	17.40	5.58	14.58	11.42	8.00
	SGM	26.50	7.62	22.54	13.88	8.98
	LinBP	31.34	8.32	24.42	16.40	8.82
	Ghost	19.8	6.18	15.92	12.16	8.38
	BPA	<u>39.12</u>	<u>12.74</u>	<u>30.26</u>	<u>22.30</u>	<u>9.04</u>
BETAK (Ours)		45.48	14.58	38.18	26.72	9.68
MI-FGSM	N/A	27.14	8.82	23.04	17.52	9.38
	SGM	33.88	11.30	31.64	20.62	10.42
	LinBP	40.34	9.92	31.62	21.06	9.52
	Ghost	30.70	9.62	24.62	19.14	9.62
	BPA	<u>52.42</u>	<u>17.66</u>	<u>40.96</u>	<u>31.34</u>	<u>11.36</u>
BETAK (Ours)		59.80	22.64	53.44	36.26	12.52

Table 3: Comparative results for ATR (%) based on 5 defense methods with PGD or MI-FGSM as the basic attacker.

mance, the loss near the origin of the adversarial loss landscape should be larger and exhibit a steeper variation trend.

We can easily observe that the adversarial loss landscapes of the adversarial image generated by SSA+BPA for different victim models are more similar to these landscapes of the clean image. In comparison, adversarial landscapes of BETAK for different victim models have greater loss and are significantly steeper than the clean image and the adversarial image of SSA+BPA, demonstrating stronger attack performance against the surrogate model, pseudo-victim model and two victim models.

Analysis of the CAM Results. To explained the mechanism of BETAK more vividly, we compare the Class Activation Mapping (CAM) [Selvaraju *et al.*, 2017] results of 6 victim models for the clean image and the AEs generated by SSA+BPA and our BETAK in Fig. 2. Darker red regions indicate that these pixels are more important for the classification output, and darker blue indicates that these pixel regions are less important. The white rectangular dashed box and the dashed lines highlight the regions that the model focuses on the most in both clean samples and AEs, significantly influencing the classification outputs. We can draw the following conclusions.

- Although these models have totally different network

structures, they all payed much attention to similar area of the clean image around the key item for classification.

- These victim models reacted absolutely different to the AE of SSA+BPA. In particular, Inc-v3 still correctly classified the shower cap, and both IncRes-v2 and DenseNet still payed much attention to the critical area of the AE as same as the clean image.
- In comparison, all the victim models misclassified the AE generated by BETAK, and they also turn their attention to totally different areas distant from the clean image. More importantly, the critical areas of the AE that these victim models focus on are concentrated in similar locations, which indicates that a better perturbation initialization is provided by BETAK.

Targeted Transfer Attack Scenario. We evaluate the generalization performance of BETAK on the targeted transfer attacks. Practically, we follow previous works [Zhao *et al.*, 2021] to optimize the logits loss based on the PGD and MI-FGSM attacker. In Tab. 2, we report the ATR of BETAK against different victim models under the targeted attack scenario. It can be observed that BETAK also consistently improves the attack performance under the more challenging targeted attack scenario.

3.3 Ablation Study

Influence of Pseudo-victim Models. We conduct the ablation study to evaluate the attack performance of BETAK by selecting different pseudo-victim models to construct the UL victim attacker in Tab. 4. We have the following observations.

- We first evaluate the ATR of BETAK with a single victim model, i.e., Inception-v3 or IncRes-v2. We can easily find that even with a single pseudo-victim model participating in the UL objective, BETAK also significantly improves the ATR again the left 8 victim models.
- Integrating IncRes-v2 exhibits higher performance improvement compared with Inception-v3, which can be explained by the fact that IncRes-v2 is more advanced by combing ResNet and Inception-v4 with higher performance and more complex structure.

Attacker	Method	Inc-v3	IncRes-v2	DenseNet	MobileNet	PNASNet	SENet	Inc-v3 _{ens3}	Inc-v3 _{ens4}	IncRes-v2 _{ens}
PGD	BPA	35.36	30.12	70.70	68.90	32.52	42.02	22.72	19.28	12.40
BETAK	+Inc-v3*	<u>50.08</u> ^{↑14.72}	<u>36.96</u> ^{↑6.84}	<u>72.56</u> ^{↑1.86}	<u>73.24</u> ^{↑4.34}	<u>36.66</u> ^{↑4.14}	<u>45.96</u> ^{↑3.94}	<u>26.84</u> ^{↑4.12}	<u>21.78</u> ^{↑2.50}	<u>14.08</u> ^{↑1.68}
	+IncRes-v2*	<u>46.64</u> ^{↑11.28}	<u>48.66</u> ^{↑18.54}	<u>74.58</u> ^{↑3.88}	<u>74.44</u> ^{↑5.54}	<u>41.94</u> ^{↑9.42}	<u>49.42</u> ^{↑7.40}	<u>30.08</u> ^{↑7.36}	<u>24.18</u> ^{↑4.90}	<u>16.82</u> ^{↑4.42}
	+Inc-v3*, IncRes-v2*	53.34 ^{↑17.98}	<u>45.08</u> ^{↑14.96}	74.68 ^{↑3.98}	74.54 ^{↑5.64}	42.48 ^{↑9.96}	49.94 ^{↑7.92}	30.80 ^{↑8.08}	24.90 ^{↑5.62}	17.60 ^{↑5.20}

Table 4: We report the ablation results of ATR (%) for the proposed BETAK by adopting different combination of pseudo-victim models to design the UL subproblem. The best and second-best outcomes are designated with boldface and underline, respectively.




Figure 3: Ablation results of the ATR for PGD+BPA and our BETAK framework by adjusting T , K or removing the DST operation (denoted as w/ and w/o DST).

- When we integrate both models to construct the UL objective, BETAK obtains better attack performance across all the 9 victim models. Therefore, combining more pseudo-victim models essentially helps generate AEs with more effective feedback of the transferability, which will facilitate the generalization performance by optimizing Eq. (5).

Influence of K and T . We conduct the ablation study to analyze the influence of UL and LL iteration. In Fig. 3, we compare the ATR of BETAK with PGD+BPA by changing T , K or removing the DST operation. As for the influence of T , we can observe that even 1-step update of δ also significantly enhance the attack performance of generated AEs, which demonstrate the effectiveness of BETAK. When we adopt more UL iterations for update, the adversarial perturbation converges with a better initialization by continuously optimizing $F(\delta)$, thus the attack performance are further improved. From this perspective, implementing BETAK with larger T improves the convergence results, thereby facilitating the transferability of AEs. As for the influence of K , the larger K leads to better approximation of $\phi^*(\delta)$ to benefit the HGR estimation, which is demonstrated by the performance improvement of ATR with larger K .

Influence of the DST Operation. From Fig. 3, it is observed that the introduction of the DST operation does not have a significant impact on the attack performance when $K = 10$. However, as K increases to 20, the absence of the DST operation results in the degradation of the attack performance on certain victim models compared to $K = 10$. In comparison, incorporating DST operation yields the highest attack performance across various victim models. Further-




Figure 4: Ablation results of the time for Backpropagation (BP) for the first 10 batches of images and the average \tilde{K} of DST operation as the UL iteration increases. Note that when DST operation is removed, \tilde{K} is set to the same value as K .

more, the influence of DST has also been analyzed with different K in Tab. 4. The observation could be summarized as follows.

- The variation of \tilde{K} with and without DST operation as T increases was examined in the first subfigure. The values depicted were calculated using the average value of first 10 batches during training. It is evident that \tilde{K} with DST operation are significantly reduced relative to the fixed K and tend to stabilize as T increases.
- We also calculate the runtime time for Backpropagation (BP) across the first 10 batches. Note that higher BP time of the first batch is attributable to one-time initialization such as computation graph construction.
- As indicated in the second subfigure, DST operation enables the BP cost at $K = 20$ to be even lower than at $K = 10$ without DST. In summary, \tilde{K} with DST operation not only assists BETAK in achieving superior attack performance but also effectively reduces the computational overhead for HGR estimation.sh

4 Conclusion

This paper introduces an initialization derived optimization paradigm to enhance the transferability of AEs, along with two efficient techniques to improve theoretical properties and computation efficiency under the BETAK framework. Comprehensive experimental evaluations validate BETAK's effectiveness with substantial ATR increase across diverse victim models and defense techniques.

Acknowledgments

This work is partially supported by the National Key R&D Program of China (No. 2022YFA1004101), the National Natural Science Foundation of China (Nos. U22B2052, 62302078 and 61936002) and the Liaoning Revitalization Talents Program (No. 2022RG04).

References

- [Alparslan *et al.*, 2020] Yigit Alparslan, Ken Alparslan, Jeremy Keim-Shenk, Shweta Khade, and Rachel Greenstadt. Adversarial attacks on convolutional neural networks in facial recognition domain. *arXiv preprint arXiv:2001.11137*, 2020.
- [Cheng *et al.*, 2021] Yaya Cheng, Jingkuan Song, Xiaosu Zhu, Qilong Zhang, Lianli Gao, and Heng Tao Shen. Fast gradient non-sign methods. *arXiv preprint arXiv:2110.12734*, 2021.
- [Cohen *et al.*, 2019] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *international conference on machine learning*, pages 1310–1320. PMLR, 2019.
- [Deng *et al.*, 2020] Yao Deng, Xi Zheng, Tianyi Zhang, Chen Chen, Guannan Lou, and Miryung Kim. An analysis of adversarial attacks and defenses on autonomous driving models. In *2020 IEEE international conference on pervasive computing and communications (PerCom)*, pages 1–10. IEEE, 2020.
- [Dong *et al.*, 2018] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9185–9193, 2018.
- [Fang *et al.*, 2022] Shuman Fang, Jie Li, Xianming Lin, and Rongrong Ji. Learning to learn transferable attack. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 571–579, 2022.
- [Franceschi *et al.*, 2017] Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and reverse gradient-based hyperparameter optimization. *arXiv:1703.01785*, 2017.
- [Goodfellow *et al.*, 2014] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [Guo *et al.*, 2017] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens Van Der Maaten. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*, 2017.
- [Guo *et al.*, 2020] Yiwen Guo, Qizhang Li, and Hao Chen. Backpropagating linearly improves transferability of adversarial examples. *Advances in neural information processing systems*, 33:85–95, 2020.
- [Hospedales *et al.*, 2020] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *arXiv:2004.05439*, 2020.
- [Hu *et al.*, 2018] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [Huang *et al.*, 2017] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [Huang *et al.*, 2019] Qian Huang, Isay Katsman, Horace He, Zeqi Gu, Serge Belongie, and Ser-Nam Lim. Enhancing adversarial example transferability with an intermediate level attack. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4733–4742, 2019.
- [Jiao *et al.*, 2023] Xianghao Jiao, Yaohua Liu, Jiaxin Gao, Xinyuan Chu, Xin Fan, and Risheng Liu. Pearl: Preprocessing enhanced adversarial robust learning of image de-raining for semantic segmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 8185–8194, 2023.
- [Kurakin *et al.*, 2018a] Alexey Kurakin, Ian Goodfellow, Samy Bengio, Yinpeng Dong, Fangzhou Liao, Ming Liang, Tianyu Pang, Jun Zhu, Xiaolin Hu, Cihang Xie, et al. Adversarial attacks and defences competition. In *The NIPS’17 Competition: Building Intelligent Systems*, pages 195–231. Springer, 2018.
- [Kurakin *et al.*, 2018b] Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pages 99–112. Chapman and Hall/CRC, 2018.
- [Li *et al.*, 2020] Yingwei Li, Song Bai, Yuyin Zhou, Cihang Xie, Zhishuai Zhang, and Alan Yuille. Learning transferable adversarial examples via ghost networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11458–11465, 2020.
- [Liao *et al.*, 2018] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense against adversarial attacks using high-level representation guided denoiser. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1778–1787, 2018.
- [Lin *et al.*, 2019] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*, 2019.
- [Liu and Liu, 2021] Yaohua Liu and Risheng Liu. Boml: A modularized bilevel optimization library in python for meta learning. In *2021 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–2. IEEE, 2021.
- [Liu *et al.*, 2018] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the*

- European conference on computer vision (ECCV)*, pages 19–34, 2018.
- [Liu et al., 2020] Risheng Liu, Pan Mu, Jian Chen, Xin Fan, and Zhongxuan Luo. Investigating task-driven latent feasibility for nonconvex image modeling. *IEEE Transactions on Image Processing*, 29:7629–7640, 2020.
- [Liu et al., 2021a] Risheng Liu, Xuan Liu, Xiaoming Yuan, Shangzhi Zeng, and Jin Zhang. A value-function-based interior-point method for non-convex bi-level optimization. In *International Conference on Machine Learning*, pages 6882–6892. PMLR, 2021.
- [Liu et al., 2021b] Risheng Liu, Yaohua Liu, Shangzhi Zeng, and Jin Zhang. Towards gradient-based bilevel optimization with non-convex followers and beyond. *Advances in Neural Information Processing Systems*, 34:8662–8675, 2021.
- [Liu et al., 2023a] Risheng Liu, Yaohua Liu, Wei Yao, Shangzhi Zeng, and Jin Zhang. Averaged method of multipliers for bi-level optimization without lower-level strong convexity. In *International Conference on Machine Learning*, pages 21839–21866. PMLR, 2023.
- [Liu et al., 2023b] Risheng Liu, Yaohua Liu, Shangzhi Zeng, and Jin Zhang. Augmenting iterative trajectory for bilevel optimization: Methodology, analysis and extensions. *arXiv preprint arXiv:2303.16397*, 2023.
- [Liu et al., 2024] Risheng Liu, Jiabin Gao, Xuan Liu, and Xin Fan. Learning with constraint learning: New perspective, solution strategy and various applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [Long et al., 2022] Yuyang Long, Qilong Zhang, Boheng Zeng, Lianli Gao, Xianglong Liu, Jian Zhang, and Jingkuan Song. Frequency domain model augmentation for adversarial attack. In *European Conference on Computer Vision*, pages 549–566. Springer, 2022.
- [Sandler et al., 2018] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [Selvaraju et al., 2017] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [Szegedy et al., 2016] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [Szegedy et al., 2017] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- [Tramèr et al., 2017] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017.
- [Wang and He, 2021] Xiaosen Wang and Kun He. Enhancing the transferability of adversarial attacks through variance tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1924–1933, 2021.
- [Wang et al., 2021] Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the transferability of adversarial attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16158–16167, 2021.
- [Wang et al., 2023] Xiaosen Wang, Kangheng Tong, and Kun He. Rethinking the backward propagation for adversarial transferability. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Wu et al., 2020] Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R Lyu, and Yu-Wing Tai. Boosting the transferability of adversarial samples via attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1161–1170, 2020.
- [Xie et al., 2017a] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. *arXiv preprint arXiv:1711.01991*, 2017.
- [Xie et al., 2017b] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Yuyin Zhou, Lingxi Xie, and Alan Yuille. Adversarial examples for semantic segmentation and object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 1369–1378, 2017.
- [Xie et al., 2019] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2730–2739, 2019.
- [Yuan et al., 2021] Zheng Yuan, Jie Zhang, Yunpei Jia, Chuanqi Tan, Tao Xue, and Shiguang Shan. Meta gradient adversarial attack. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7748–7757, 2021.
- [Zhao et al., 2021] Zhengyu Zhao, Zhuoran Liu, and Martha Larson. On success and simplicity: A second look at transferable targeted attacks. *Advances in Neural Information Processing Systems*, 34:6115–6128, 2021.